EXHIBIT C

United States Government Accountability Office



Report to the Committee on Homeland Security and Governmental Affairs, U.S. Senate

May 2014

VIRTUAL CURRENCIES

Emerging Regulatory, Law Enforcement, and Consumer Protection Challenges

GAO Highlights

Highlights of GAO-14-496, a report to the Committee on Homeland Security and Governmental Affairs, U.S. Senate

Why GAO Did This Study

Virtual currencies-digital representations of value that are not government-issued-have grown in popularity in recent years. Some virtual currencies can be used to buy real goods and services and exchanged for dollars or other currencies. One example of these is bitcoin, which was developed in 2009. Bitcoin and similar virtual currency systems operate over the Internet and use computer protocols and encryption to conduct and verify transactions. While these virtual currency systems offer some benefits, they also pose risks. For example, they have been associated with illicit activity and security breaches, raising possible regulatory, law enforcement, and consumer protection issues. GAO was asked to examine federal policy and interagency collaboration issues concerning virtual currencies.

This report discusses (1) federal financial regulatory and law enforcement agency responsibilities related to the use of virtual currencies and associated challenges and (2) actions and collaborative efforts the agencies have undertaken regarding virtual currencies. To address these objectives, GAO reviewed federal laws and regulations, academic and industry research, and agency documents; and interviewed federal agency officials, researchers, and industry groups.

What GAO Recommends

GAO recommends that CFPB take steps to identify and participate in pertinent interagency working groups addressing virtual currencies, in coordination with other participating agencies. CFPB concurred with this recommendation.

View GAO-14-496. For more information, contact Lawrance L. Evans, Jr. at (202) 512-8678 or evansl@gao.gov.

VIRTUAL CURRENCIES

Emerging Regulatory, Law Enforcement, and Consumer Protection Challenges

What GAO Found

Virtual currencies are financial innovations that pose emerging challenges to federal financial regulatory and law enforcement agencies in carrying out their responsibilities, as the following examples illustrate:

- Virtual currency systems may provide greater anonymity than traditional payment systems and sometimes lack a central intermediary to maintain transaction information. As a result, financial regulators and law enforcement agencies may find it difficult to detect money laundering and other crimes involving virtual currencies.
- Many virtual currency systems can be accessed globally to make payments and transfer funds across borders. Consequently, law enforcement agencies investigating and prosecuting crimes that involve virtual currencies may have to rely upon cooperation from international partners who may operate under different regulatory and legal regimes.
- The emergence of virtual currencies has raised a number of consumer and investor protection issues. These include the reported loss of consumer funds maintained by bitcoin exchanges, volatility in bitcoin prices, and the development of virtual-currency-based investment products. For example, in February 2014, a Tokyo-based bitcoin exchange called Mt. Gox filed for bankruptcy after reporting that it had lost more than \$460 million.

Federal financial regulatory and law enforcement agencies have taken a number of actions regarding virtual currencies. In March 2013, the Department of the Treasury's Financial Crimes Enforcement Network (FinCEN) issued guidance that clarified which participants in virtual currency systems are subject to antimoney-laundering requirements and required virtual currency exchanges to register with FinCEN. Additionally, financial regulators have taken some actions regarding anti-money-laundering compliance and investor protection. For example, in July 2013, the Securities and Exchange Commission (SEC) charged an individual and his company with defrauding investors through a bitcoin-based investment scheme. Further, law enforcement agencies have taken actions against parties alleged to have used virtual currencies to facilitate money laundering or other crimes. For example, in October 2013, multiple agencies worked together to shut down Silk Road, an online marketplace where users paid for illegal goods and services with bitcoins.

Federal agencies also have begun to collaborate on virtual currency issues through informal discussions and interagency working groups primarily concerned with money laundering and other law enforcement matters. However, these working groups have not focused on emerging consumer protection issues, and the Consumer Financial Protection Bureau (CFPB)—whose responsibilities include providing consumers with information to make responsible decisions about financial transactions—has generally not participated in these groups. Therefore, interagency efforts related to virtual currencies may not be consistent with key practices that can benefit interagency collaboration, such as including all relevant participants to ensure they contribute to the outcomes of the effort. As a result, future interagency efforts may not be in a position to address consumer risks associated with virtual currencies in the most timely and effective manner.

Contents

Letter		1
	Background Federal Agencies Face Emerging Challenges in Carrying Out	4
	Responsibilities Related to the Use of Virtual Currencies Agencies Have Taken Some Actions on Virtual Currencies, but	12
	Risks	24
	Conclusions	39
	Recommendation for Executive Action	40
	Agency Comments	40
Appendix I	How Bitcoins Enter into Circulation and Are Used in Transactions	42
Appendix II	Interagency Working Groups that Have Addressed Virtual Currency	
	Issues	43
Appendix III	Comments from the Consumer Financial Protection Bureau	49
Appendix IV	Comments from the National Credit Union Administration	50
Appendix V	GAO Contact and Staff Acknowledgments	51
Table		
	Table 1: Interagency Working Groups that Have Addressed Virtual Currency Issues, as of April 2014	43
Figures		
	Figure 1: Ways to Obtain and Spend Bitcoins	8
	Figure 2: Bitcoin Price Index in U.S. Dollars, January 1, 2013	40
	Inrougn March 31, 2014 Figure 3: Screen Shot of the Silk Road Website	10 วว
	I Igure J. Ourden Shot of the Slik I todu Website	55

Figure 4: How Bitcoins Enter into Circulation and Are Used in Transactions

Abbreviations		
BSA	Bank Secrecy Act	
BSAAG	Bank Secrecy Act Advisory Group	
CFPB	Consumer Financial Protection Bureau	
CFTC	Commodity Futures Trading Commission	
DATA	Digital Asset Transfer Authority	
DEA	Drug Enforcement Administration	
DHS	Department of Homeland Security	
DOJ	Department of Justice	
ECTF	Electronic Crimes Task Forces	
EFTA	Electronic Fund Transfer Act	
FATF	Financial Action Task Force	
FBI	Federal Bureau of Investigation	
FDIC	Federal Deposit Insurance Corporation	
FFIEC	Federal Financial Institutions Examination Council	
FinCEN	Financial Crimes Enforcement Network	
HSI	Homeland Security Investigations	
ICE	U.S. Immigration and Customs Enforcement	
IOC-2	International Organized Crime Intelligence and Operations Center	
IRS	Internal Revenue Service	
NCUA	National Credit Union Administration	
000	Office of the Comptroller of the Currency	
SEC	Securities and Exchange Commission	
TOR	The Onion Router	
USAID	United States Agency for International Development	
VCET	Virtual Currency Emerging Threats Working Group	

This is a work of the U.S. government and is not subject to copyright protection in the United States. The published product may be reproduced and distributed in its entirety without further permission from GAO. However, because this work may contain copyrighted images or other material, permission from the copyright holder may be necessary if you wish to reproduce this material separately.

42

S. GOVERNMENT ACCOUNTABILITY OFFICE

441 G St. N.W. Washington, DC 20548

May 29, 2014

The Honorable Thomas R. Carper Chairman

The Honorable Tom A. Coburn Ranking Member Committee on Homeland Security and Governmental Affairs United States Senate

While not widely used or accepted, virtual currencies, such as bitcoin, have grown in popularity in recent years and have emerged for some as potential alternatives to traditional currencies issued by governments. Virtual currencies operate over the Internet and, in some cases, may be used to buy real goods and services and exchanged for traditional currencies. They offer potential benefits over traditional currencies, including lower transaction costs and faster funds transfers. Because some virtual currency transactions provide greater anonymity than transactions using traditional payment systems, law enforcement and financial regulators have raised concerns about the use of virtual currencies for illegal activities. Additionally, recent cases involving the loss of funds from virtual currency exchanges have highlighted potential consumer protection issues.

You asked us to examine potential policy issues related to virtual currencies and the status of federal agency collaboration in this area. This report focuses on the federal financial regulatory agencies and selected federal law enforcement agencies that have a role in protecting the U.S. financial system and investigating financial crimes.¹ Specifically, this report addresses (1) agency responsibilities related to the use of virtual currencies and the emerging challenges these currencies pose to the

¹Other federal agencies that were outside the scope of this report, such as the Internal Revenue Service (IRS), have responsibilities related to virtual currencies. For example, as we reported in May 2013, IRS is responsible for ensuring taxpayer compliance for all economic areas, including virtual economies and currencies. For more information, see GAO, *Virtual Economies and Currencies: Additional IRS Guidance Could Reduce Tax Compliance Risks*, GAO-13-516 (Washington, D.C.: May 15, 2013). In March 2014, IRS determined that virtual currencies will be treated as property for purposes of U.S. federal taxes. Therefore, general tax principles that apply to property transactions apply to transactions using virtual currency. See IRS Notice 2014-21.

agencies; and (2) actions the agencies have taken in response to the emergence of virtual currencies, including interagency collaborative efforts. We selected the law enforcement agencies included in our review based on their involvement in investigating virtual-currency-related crimes and participation in interagency collaborative efforts and congressional hearings on virtual currency issues.

To describe agency responsibilities related to the use of virtual currencies and the emerging challenges these currencies pose, we reviewed the following agency information: testimony and written statements from relevant congressional hearings, written responses to congressional questions, unclassified intelligence assessments, financial reports, training presentations, and descriptions of missions and responsibilities from agencies' websites.² We also reviewed prior GAO reports, Congressional Research Service reports, and relevant laws and regulations, including the Bank Secrecy Act (BSA) and related anti-money laundering provisions such as Title III of the USA PATRIOT Act, to gain an understanding of agencies' responsibilities in administering and enforcing anti-money-laundering laws and regulations, as well as in investigating and prosecuting financial and other crimes.³ In addition, we reviewed academic articles and papers from industry stakeholders. Further, we interviewed officials from the following federal financial regulatory and law enforcement agencies:

- The Board of Governors of the Federal Reserve System (Federal Reserve);
- The Bureau of Consumer Financial Protection (also known as the Consumer Financial Protection Bureau or CFPB);
- The Commodity Futures Trading Commission (CFTC);
- The Department of Homeland Security (DHS), including U.S. Immigration and Customs Enforcement–Homeland Security Investigations (ICE-HSI) and the U.S. Secret Service (Secret Service);

²We reviewed testimony and agency statements from two congressional hearings: the November 18, 2013, U.S. Senate Committee on Homeland Security and Governmental Affairs hearing "Beyond Silk Road: Potential Risks, Threats, and Promises of Virtual Currencies," and the November 19, 2013, U.S. Senate Committee on Banking, Housing, and Urban Affairs hearing, "The Present and Future Impact of Virtual Currency."

³Pub. L. No. 91-508, 84 Stat. 1114 (1970) (codified as amended at 12 U.S.C. §§ 1829(b), 1951-1959; 31 U.S.C. §§ 5311-5330); Pub. L. No. 107-56, 115 Stat. 272 (2001) (codified as amended in scattered sections of U.S.C.).

- The Department of Justice (DOJ), including the Criminal Division and two of its components—the Asset Forfeiture and Money Laundering Section and Computer Crime and Intellectual Property Section—and the Federal Bureau of Investigation (FBI);
- The Department of the Treasury (Treasury), including the Financial Crimes Enforcement Network (FinCEN) and the Office of the Comptroller of the Currency (OCC);
- The Federal Deposit Insurance Corporation (FDIC);
- The National Credit Union Administration (NCUA); and
- The Securities and Exchange Commission (SEC).

Additionally, we interviewed an academic whose research focused on virtual currencies and industry stakeholders, including the Bitcoin Foundation, the Digital Asset Transfer Authority (DATA), and the National Money Transmitters Association, which represent the interests of a large number of virtual currency and money transmission businesses.

To examine the actions and collaborative efforts federal agencies have undertaken in response to the emergence of virtual currencies, we reviewed agency information, including FinCEN's regulatory guidance and administrative rulings on the applicability of BSA to virtual currency participants, testimony and written statements from the previously mentioned congressional hearings, written responses to congressional questions, intelligence assessments, a CFPB query of its Consumer Complaint Database, and press releases.⁴ We also interviewed officials from the agencies listed previously to obtain further information on the actions they have taken to address the emergence of virtual currencies and their efforts to collaborate with other federal agencies on this issue. Additionally, we interviewed the academic and industry stakeholders noted previously, as well as the Digital Economy Task Force, to determine the extent to which private sector groups were involved in

⁴FinCEN, Application of FinCEN's Regulations to Persons Administering, Exchanging, or Using Virtual Currencies, FIN-2013-G001, March 18, 2013; FinCEN, Application of FinCEN's Regulations to Virtual Currency Mining Operations, FIN-2014-R001, January 30, 2014; FinCEN, Application of FinCEN's Regulations to Virtual Currency Software Development and Certain Investment Activity, FIN-2014-R002, January 30, 2014; and FinCEN, Application of Money Services Business Regulations to the Rental of Computer Systems for Mining Virtual Currencies, FIN-2014-R007, April 29, 2014.

interagency collaborative efforts.⁵ We reviewed GAO's key practices on collaboration and assessed whether interagency collaborative efforts related to virtual currencies were consistent with practices concerning the inclusion of relevant participants.⁶

We conducted this performance audit from November 2013 to May 2014 in accordance with generally accepted government auditing standards. Those standards require that we plan and perform the audit to obtain sufficient, appropriate evidence to provide a reasonable basis for our findings and conclusions based on our audit objectives. We believe that the evidence obtained provides a reasonable basis for our findings and conclusions based on our audit objectives.

Background

Virtual currencies are financial innovations that have grown in number and popularity in recent years. While there is no statutory definition for virtual currency, the term refers to a digital representation of value that is not government-issued legal tender. Unlike U.S. dollars and other government-issued currencies, virtual currencies do not necessarily have a physical coin or bill associated with their circulation. While virtual currencies can function as a unit of account, store of value, and medium of exchange, they are not widely used or accepted. Some virtual currencies can only be used within virtual economies (for example, within online role-playing games) and may not be readily exchanged for government-issued currencies such as U.S. dollars, euro, or yen. Other virtual currencies may be used to purchase goods and services in the real economy and can be converted into government-issued currencies through virtual currency exchanges. In previous work, we described the

⁵The Digital Economy Task Force was established in 2013 by Thomson Reuters (a multinational media and information firm) and the International Centre for Missing & Exploited Children to explore the benefits and risks of the emerging digital economy, including the use of virtual currency. This task force includes members from both the public and private sectors.

⁶GAO, Managing for Results: Key Considerations for Implementing Interagency Collaborative Mechanisms, GAO-12-1022 (Washington, D.C.: Sept. 27, 2012) and Managing for Results: Implementation Approaches Used to Enhance Collaboration in Interagency Groups, GAO-14-220 (Washington, D.C.: Feb. 14, 2014).

latter type of virtual currencies as "open flow."⁷ Open-flow virtual currencies have received considerable attention from federal financial regulatory and law enforcement agencies, in part because these currencies interact with the real economy and because depository institutions (for example, banks and credit unions) may have business relationships with companies that exchange virtual currencies for government-issued currencies. Throughout the remainder of this report, we use the term virtual currencies to mean open-flow virtual currencies, unless otherwise stated.⁸

Virtual currency systems, which include protocols for conducting transactions in addition to digital representations of value, can either be centralized or decentralized. Centralized virtual currency systems have a single administering authority that issues the currency and has the authority to withdraw the currency from circulation. In addition, the administrating authority issues rules for use of the currency and maintains a central payment ledger. In contrast, decentralized virtual currency systems have no central administering authority. Validation and certification of transactions are performed by users of the system and therefore do not require a third party to perform intermediation activities.

A prominent example of a decentralized virtual currency system is bitcoin. Bitcoin was developed in 2009 by an unidentified programmer or programmers using the name Satoshi Nakamoto. According to industry stakeholders, bitcoin is the most widely circulated decentralized virtual currency. The bitcoin computer protocol permits the storage of unique digital representations of value (bitcoins) and facilitates the assignment of bitcoins from one user to another through a peer-to-peer, Internet-based

⁷GAO-13-516. In that report we described "closed-flow" virtual currencies as those that can be used only within a game or virtual environment and cannot be cashed out for dollars or other government-issued currencies. We also described hybrid virtual currencies as those that have characteristics of both open- and closed-flow currencies—for example, such currencies can be used to buy real goods and services but are not exchangeable for government-issued currencies.

⁸Some stakeholders with whom we spoke said they preferred the term digital currency to virtual currency, due partly to the connotation that something which is virtual cannot be used in the real world. We use the term virtual currency to be consistent with terminology used in prior GAO work and in key federal guidance on participants in virtual currency systems.

network.⁹ Each bitcoin is divisible to eight decimal places, enabling their use in any kind of transaction regardless of the value. Users' bitcoin balances are associated with bitcoin addresses (long strings of numbers and letters) that use principles of cryptography to help safeguard against inappropriate tampering with bitcoin transactions and balances.¹⁰ When users transfer bitcoins, the recipient provides their bitcoin address to the sender, and the sender authorizes the transaction with their private key (essentially a secret code that proves the sender's control over their bitcoin address). Bitcoin transactions are irrevocable and do not require the sender or receiver to disclose their identities to each other or a third party. However, each transaction is registered in a public ledger called the "blockchain," which maintains the associated bitcoin addresses and transaction dates, times, and amounts. Users can define how much additional information they require of each other to conduct a transaction.

Because peer-to-peer bitcoin transactions do not require the disclosure of information about a user's identity, they give the participants some degree of anonymity. In addition, computer network communication can be encrypted and anonymized by software to further hide the identity of the parties in transactions.¹¹ However, the transactions are not completely anonymous because the time and amount of each transaction and the associated bitcoin addresses are permanently recorded in the blockchain. As a result, peer-to-peer bitcoin transactions are sometimes described as "pseudonymous." The anonymity of bitcoin is also limited by data analysis techniques that can potentially link bitcoin addresses to personal identities. For example, information about a customer's identity may be recorded when an individual exchanges dollars for bitcoins, and this information may be combined with data from the blockchain to determine

⁹A peer-to-peer network allows users to share data directly and conduct permitted activities without a central server.

¹⁰Cryptography is a branch of mathematics that is based on the transformation of data and can be used to provide security services such as confidentiality and authentication. Bitcoin and other virtual currencies that use cryptography are sometimes called cryptocurrencies.

¹¹According to industry observers, examples of technologies used to increase the privacy of participants in virtual currency transactions include (1) anonymizing networks, which use a distributed network of computers to conceal the real Internet address of users, such as The Onion Router (TOR); (2) "tumblers" such as BitcoinBath and BitLaundry that combine payments from multiple users to obstruct identification through the blockchain; and (3) alternative virtual currencies such as Zerocoin and Anoncoin that aim to make transactions fully anonymous.

the identities of participants in bitcoin transactions. In addition, researchers have developed methods to determine identities of parties involved in some bitcoin transactions by analyzing clusters of transactions between specific addresses.¹²

Bitcoins are created and entered into circulation through a process called mining. Bitcoin miners download free software that they use to solve complex math problems. Solving these problems verifies the validity of bitcoin transactions by grouping several transactions into a block and mathematically proving that the transactions occurred and did not involve double spending of a bitcoin. On average, this process takes about 10 minutes. When a miner or group of miners (mining pools) solves a problem, the bitcoin network accepts the block of transactions as valid and creates new bitcoins and awards them to the successful miner or mining pool.¹³ (For a diagram on how bitcoins enter into circulation through mining, how transactions are conducted, and how miners verify transactions, see app. I.) Over time, the computer processing power needed to mine new bitcoins has increased to the point where mining requires specialized computer hardware and has become increasingly consolidated into large mining pools.

In addition to mining new bitcoins, users can also acquire bitcoins already in circulation by accepting bitcoins as gifts or payments for goods or services, purchasing them at bitcoin kiosks (sometimes referred to as bitcoin automated teller machines), or purchasing them on third-party exchanges. These exchanges allow users to exchange traditional currencies such as U.S. dollars for bitcoins, and exchange bitcoins back to traditional currencies. Individuals may store their bitcoins in a "virtual wallet" (a program that saves bitcoin addresses) on their computer or other data storage device, or use an online wallet service provided by an exchange or third-party virtual wallet provider. To spend their bitcoins, individuals can buy goods or services from other bitcoin users. They may also make purchases from online businesses that either accept bitcoins

¹²See Sarah Meiklejohn, et al, "A Fistful of Bitcoins: Characterizing Payments Among Men with No Names," *;Login:*, vol. 38 no. 6 (2013), available at https://www.usenix.org/system/files/login/articles/03 meiklejohn-online.pdf.

¹³By design, there will be a maximum of 21 million bitcoins in circulation once all bitcoins have been mined, which is projected to occur in the year 2140. Once all bitcoins have been mined, miners will be rewarded for solving the math problems that verify the validity of bitcoin transactions through fees rather than bitcoins.

directly or use third-party payment processors that take payments in bitcoins from buyers and provide businesses the payments in the form of a traditional currency or a combination of bitcoins and traditional currency. Figure 1 shows various ways that individuals can obtain and spend bitcoins.

Figure 1: Ways to Obtain and Spend Bitcoins Obtaining bitcoins Spending bitcoins Exchanges: Virtual currency Other bitcoin users: Bob transfers bitcoins directly to exchange converts Bob's traditional currency, such as bitcoin addresses in the virtual U.S. dollars, into bitcoins and wallets of other bitcoin users transfers them to a bitcoin as a gift or as payment for address in his virtual wallet. goods or services. Exchanges also convert bitcoins into traditional currencies. Other bitcoin users: Bitcoin **Businesses accepting** users transfer bitcoins directly bitcoins directly: To pay for to a bitcoin address in Bob's goods or services, Bob virtual wallet as a gift or for transfers bitcoins directly to payment of goods or services. bitcoin addresses of Bob's businesses that accept virtual payment in bitcoins. wallet Bitcoin kiosks: Bob deposits **Businesses accepting bitcoins** traditional currency into a through payment processors: bitcoin kiosk. The kiosk sends To pay for goods or services, bitcoins from its operator's Bob transfers bitcoins to a bitcoin address to a bitcoin business's payment processor. address in Bob's virtual wallet. The processor converts the bitcoins into traditional currency and remits the traditional currency to the business. In some cases, the processor Mining: Bob installs bitcoin converts only a portion of the mining software on his bitcoins into traditional computer, which is used to currencies. solve complex math problems for the bitcoin network. If Bob successfully solves the problems, he receives newly created bitcoins.

Source: GAO.

Due to limitations in available data, the size of the bitcoin market is unclear.¹⁴ Nonetheless, some data exist that may provide some context for the size of this market:

- According to statistics from the bitcoin blockchain, as of March 31, 2014, approximately 12.6 million bitcoins were in circulation.¹⁵
- At exchange rates as of March 31, 2014 (about \$458 per bitcoin), the total value of the approximately 12.6 million bitcoins in circulation was about \$5.6 billion.¹⁶ For perspective, the total amount of U.S. currency held by the public and in transaction deposits (mainly checking accounts) at depository institutions was about \$2.7 trillion as of March 2014.¹⁷
- Bitcoin exchange rates against the U.S. dollar have changed dramatically over time (see fig. 2). According to one bitcoin price index, the price was about \$13 per bitcoin in the beginning of January 2013 and rose to more than \$1,100 by the beginning of December 2013. Prices subsequently fell to about \$522 in mid-December 2013 and have fluctuated between roughly \$450 and \$950 since then.¹⁸
- From April 2013 through March 2014, the number of bitcoin transactions per day ranged from about 29,000 to 102,000.¹⁹ In comparison, the Federal Reserve Banks processed an average of 44

¹⁷See Federal Reserve Statistical Release H.6 "Money Stock Measures" (Apr. 10, 2014) at http://www.federalreserve.gov/releases/h6/current/H6.pdf.

¹⁸https://www.coindesk.com. (Accessed on Apr.1, 2014.) This index is a composite price calculated as the simple average of bitcoin prices across leading global exchanges that meet certain criteria.

¹⁹https://blockchain.info. (Accessed on Apr. 1, 2014.)

¹⁴Given these limitations, we did not test the reliability of data, such as the data generated from the bitcoin network, but we are providing some figures to provide context for the possible size of the bitcoin market and other virtual currency markets.

¹⁵http://blockchain.info. (Accessed on Mar. 31, 2014.) Due to data limitations, it is difficult to calculate the velocity, or the rate at which bitcoins are spent, and the number of transactions between unique users in a given time period.

¹⁶For data on bitcoin price, see https://www.coindesk.com. (Accessed on Apr. 1, 2014.) For data on the total value and number of bitcoins in circulation, see https://blockchain.info. (Accessed on Mar. 31, 2014.)

million commercial Automated Clearing House (a traditional payment processor) transactions per day in 2013.²⁰



Source: GAO analysis of data from http://www.coindesk.com/price/ (accessed on Apr. 1, 2014).

Note: The index is a composite price calculated as the simple average of bitcoin prices across leading global exchanges that meet certain criteria. The values are expressed in current U.S. dollars.

While bitcoin is the most widely used virtual currency, numerous others have been created. For example, dozens of decentralized virtual currencies are based on the bitcoin protocol such as Litecoin, Auroracoin, Peercoin, and Dogecoin. Similar to the bitcoin market, the size of the market for these virtual currencies is unclear. However, as of March 31, 2014, the total reported value of each of these currencies was less than \$400 million (ranging from about \$33 million for Dogecoin to about \$346 million for Litecoin).²¹ Other virtual currencies that have been created are

²⁰Federal Reserve. See

http://www.federalreserve.gov/paymentsystems/fedach_yearlycomm.htm. (Accessed on Apr. 1, 2014.)

²¹https://coinmarketcap.com. (Accessed on Apr. 1, 2014.)

not based on the bitcoin protocol. One of the more prominent examples is XRP, which is used within a decentralized payment system called Ripple. Ripple allows users to make peer-to-peer transfers in any currency. A key function of XRP is to facilitate the conversion from one currency to another. For example, if a direct conversion between Mexican pesos and Thai baht is not available, the pesos can be exchanged for XRP, and then the XRP for baht. As of March 31, 2014, the total value of XRP was \$878 million.²²

Virtual currencies have drawn attention from federal agencies with responsibilities for protecting the U.S. financial system and its participants and investigating financial crimes. These include, but are not limited to, CFPB, CFTC, DHS, DOJ, SEC, Treasury, and the prudential banking regulators. The prudential banking regulators are the FDIC, Federal Reserve, NCUA, and OCC. Within Treasury, FinCEN has a particular interest in the emergence of virtual currencies because of concerns about the use of these currencies for money laundering and FinCEN's role in combating such activity.²³ Additionally, because virtual currencies (like government-issued currencies) can play a role in a range of financial and other crimes, including cross-border criminal activity, key components of DOJ and DHS have an interest in how virtual currencies are used. Relevant DOJ components include the Criminal Division (which oversees the Computer Crime and Intellectual Property Section and the Asset Forfeiture and Money Laundering Section), the FBI, and the Offices of the U.S. Attorneys (U.S. Attorneys). Relevant DHS components include the Secret Service and ICE-HSI.

²²https://coinmarketcap.com.(Accessed on Apr. 1, 2014.)

²³Money laundering is the process of disguising or concealing the source of funds acquired illicitly to make the acquisition appear legitimate.

Federal Agencies Face Emerging Challenges in Carrying Out Responsibilities Related to the Use of Virtual Currencies	While federal agencies' responsibilities with respect to virtual currency are still being clarified, some virtual currency activities and products have implications for the responsibilities of federal financial regulatory and law enforcement agencies. Virtual currencies have presented these agencies with emerging challenges as they carry out their different responsibilities. These challenges stem partly from certain characteristics of virtual currency systems, such as the higher degree of anonymity they provide compared with traditional payment systems and the ease with which they can be accessed globally to make payments and transfer funds across borders.
Some Virtual Currency Activities and Products May Have Implications for Federal Agencies' Responsibilities	Although virtual currencies are not government-issued and do not currently pass through U.S. banks, some activities and products that involve virtual currencies have implications for the responsibilities of federal financial regulatory and law enforcement agencies. These activities and products encompass both legitimate and illegitimate uses of virtual currencies. Examples of legitimate uses include buying virtual currencies and registered virtual-currency-denominated investment products. Examples of illegitimate uses include money laundering and purchasing illegal goods and services using virtual currencies.
FinCEN	FinCEN administers BSA and its implementing regulations. ²⁴ The goal of BSA is to prevent financial institutions from being used as intermediaries for the transfer or deposit of money derived from criminal activity and to provide a paper trail to assist law enforcement agencies in their money laundering investigations. To the extent that entities engaged in money transmission conduct virtual currency transactions with U.S. customers or become customers of a U.S. financial institution, FinCEN has

²⁴Pub. L. No. 91-508, 84 Stat. 1114 (1970) (codified as amended at 12 U.S.C. §§ 1829(b), 1951-1959; 31 U.S.C. §§ 5311-5330); 31 C.F.R. chap. X. In 1994, the Secretary of the Treasury delegated overall authority for enforcement of, and compliance with, BSA and its implementing regulations related to money laundering to the Director of FinCEN. In the same year, the Secretary also delegated BSA examination authority to the prudential banking regulators. 31 C.F.R. § 1010.810(b)(1)-(5).

responsibilities for helping ensure that these entities comply with BSA and anti-money-laundering regulations.²⁵

FinCEN regulations set forth requirements for money services businesses, which include financial institutions and other entities engaged in money transmission.²⁶ FinCEN guidance states that the agency's regulations regarding money services businesses apply to virtual currency exchangers and administrators.²⁷ FinCEN applies its regulations to "convertible virtual currency," which either has an equivalent value in real currency or acts as a substitute for real currency. FinCEN regulations require money services businesses to assess their exposure to money laundering and terrorist financing and establish risk mitigation plans in the form of anti-money-laundering programs.²⁸ Additionally, money services businesses are required to maintain transaction records. For example, for money transfers that are \$3,000 or more, money services businesses must obtain information on the transmitter, the recipient, and the transaction itself, and pass on such information to other intermediary financial institutions in any subsequent fund transmissions. Money

²⁷ FinCEN, *Application of FinCEN's Regulations to Persons Administering, Exchanging, or Using Virtual Currencies*, FIN-2013-G001, March 18, 2013. FinCEN defines an exchanger as a person engaged as a business in the exchange of virtual currency for real currency, funds, or other virtual currency. *Id.* FinCEN defines an administrator as a person engaged as a business in issuing (putting into circulation) a virtual currency, and who has the authority to redeem (to withdraw from circulation) such virtual currency. *Id.* An administrator or exchanger that (1) accepts and transmits a convertible virtual currency, or (2) buys or sells convertible virtual currency for any reason is a money transmitter under FinCEN's regulations.

²⁸31 C.F.R. § 1022.210, subpart C.

²⁵FinCEN shares this responsibility with IRS, to which FinCEN has delegated examination authority for money services businesses. See 31 C.F.R. § 1010. 810(b)(8). IRS activities were outside the scope of our review. FinCEN has also delegated examination authority for BSA compliance to a number of other federal agencies, including the prudential banking regulators, CFTC, and SEC. See 31 C.F.R. § 1010.810(b). These agencies can also use their independent authorities to examine entities under their supervision for compliance with applicable BSA and anti-money-laundering requirements and regulations.

²⁶Under 31 C.F.R. § 1010.100(ff)(1)-(7), money services businesses are generally defined as any of the following: (1) currency dealer or exchanger, (2) check casher, (3) issuer or seller of traveler's checks or money orders, (4) provider or seller of prepaid access, (5) money transmitter, and (6) the U.S. Postal Service. FinCEN's regulations define a money transmitter as a person that provides money transmission services, or any other person engaged in the transfer of funds. 31 C.F.R. § 1010.100(ff)(5)(i).The term money transmission services means the "acceptance of currency, funds, or other value that substitutes for currency to another location or person by any means." Id.

	services businesses are also required to monitor transactions and file reports on large currency transactions and suspicious activities. In addition, certain financial institutions must establish a written customer identification program that includes procedures for obtaining minimum identification information from customers who open an account, such as date of birth, a government identification number, and physical address. ²⁹ Further, financial institutions must file currency transaction reports on customer cash transactions exceeding \$10,000 that include information about the account owner's identity and occupation. ³⁰
	FinCEN also supports the investigative and prosecutive efforts of multiple federal and state law enforcement agencies through its administration of the financial transaction reporting and recordkeeping requirements mandated or authorized under BSA. In addition, FinCEN has the authority to take enforcement actions, such as assessing civil money penalties, against financial institutions, including money services businesses, that violate BSA requirements.
Prudential Banking Regulators	The prudential banking regulators—FDIC, Federal Reserve, NCUA, and OCC—provide oversight of depository institutions' compliance with BSA and anti-money-laundering requirements. Therefore, these regulators are responsible for providing guidance and oversight to help ensure that depository institutions that have opened accounts for virtual currency exchanges or other money services businesses have adequate anti-money-laundering controls for those accounts. ³¹ In April 2005, FinCEN and the prudential banking regulators issued joint guidance to banking organizations (depository institutions and bank holding companies) to clarify BSA requirements with respect to money services businesses and to set forth the minimum steps that banking organizations should take

³⁰31 U.S.C. § 5313(a); 31 C.F.R. § 1010.311.

²⁹31 C.F.R. § 1020.220(a)(2)(i). Under the USA PATRIOT Act, financial institutions also must implement appropriate, specific, and, where necessary, enhanced, due diligence for correspondent accounts and private banking accounts established in the United States for non-U.S. persons. 31 U.S.C. § 5318(i).

³¹In addition, officials from the prudential banking regulators either stated or acknowledged that they would have authority to regulate a supervised entity that issued virtual currency, or cleared or settled transactions related to virtual currency.

	when providing banking services to these businesses. ³² As part of safety and soundness or targeted BSA compliance examinations of depository institutions, the prudential banking regulators assess compliance with BSA and related anti-money-laundering requirements using procedures that are consistent with their overall risk-focused examination approach. ³³ In examining depository institutions for BSA compliance, the regulators review whether depository institutions (1) have developed anti-money- laundering programs and procedures to detect and report unusual or suspicious activities possibly related to money laundering; and (2) comply with the technical recordkeeping and reporting requirements of BSA. ³⁴ While most cases of BSA noncompliance are corrected within the examination framework, regulators can take a range of supervisory actions, including formal enforcement actions, against the entities they supervise for violations of BSA and anti-money-laundering requirements. These formal enforcement actions can include imposing civil money penalties and initiating cease-and-desist proceedings. ³⁵
Consumer Financial Protection Bureau	CFPB is an independent entity within the Federal Reserve that has broad consumer protection responsibilities over an array of consumer financial products and services, including taking deposits and transferring money. CFPB is responsible for enforcing federal consumer protection laws, and it is the primary consumer protection supervisor over many of the
	³² FinCEN, Interagency Interpretive Guidance on Providing Banking Services to Money Services Businesses Operating in the United States, April 26, 2005. FinCEN concurrently issued guidance to money services businesses that identified and explained the types of information and documentation that money services businesses were expected to have and provide to banking organizations. Bank holding companies are companies that own or control one or more banks. In the United States, most banks insured by FDIC are owned or controlled by a bank holding company.
	³³ Under the risk-focused approach, those activities judged to pose the highest risk to an institution are to receive the most scrutiny by examiners.
	³⁴ See 12 U.S.C. § 1786(q), § 1818(s) (federal banking agencies must promulgate regulations requiring insured depository institutions and credit unions to establish procedures regarding BSA compliance; regulators' examinations must include review of BSA compliance procedures); see also procedures for monitoring BSA compliance: 12 C.F.R. § 208.63 (Federal Reserve), 12 C.F.R. § 326.8 (FDIC), 12 C.F.R. § 748.2 (NCUA), and 12.C.F.R. § 21.21 (OCC).
	³⁵ A civil money penalty is a punitive fine assessed for the violation of a law or regulation or for other misconduct. A cease-and-desist proceeding is a formal process that may result in an order that a party halt certain activities or practices; the order may also require the party to take affirmative action to correct the conditions resulting from the practices. See 12 U.S.C. § 1786(e), § 1818(b).

institutions that offer consumer financial products and services. CFPB also has authority to issue and revise regulations that implement federal consumer financial protection laws, including the Electronic Fund Transfer Act³⁶ and title X of the Dodd-Frank Wall Street Reform and Consumer Protection Act (Dodd-Frank Act).³⁷ CFPB officials stated that they are reviewing how these responsibilities are implicated by consumer use (or potential consumer use) of virtual currencies.

Other relevant CFPB responsibilities concerning virtual currencies include accepting and handling consumer complaints, promoting financial education, researching consumer behavior, and monitoring financial markets for new risks to consumers. For example, under authorities provided by the Dodd-Frank Act, CFPB maintains a Consumer Complaint Database and helps monitor and assess risks to consumers in the offering or provision of consumer financial products or services.³⁸ CFPB also issues consumer advisories to promote clarity, transparency, and fairness in consumer financial markets.

SEC regulates the securities markets—including participants such as securities exchanges, broker-dealers, investment companies, and investment advisers—and takes enforcement actions against individuals and companies for violations of federal securities laws. SEC's mission is to protect investors; maintain fair, orderly, and efficient markets; and facilitate capital formation. Virtual currencies may have implications for a number of SEC responsibilities. For example, SEC has enforcement

³⁷Pub. L. No. 111-203, § 1021(c)(5), 124 Stat. 1376, 1980 (2010) (codified at 12 U.S.C. § 5511(c)(5)). For example, section 1032(a) of the Dodd-Frank Act confers authority on CFPB "to prescribe rules to ensure that the features of any consumer financial product or service, both initially and over the term of the product or service, are fully, accurately, and effectively disclosed to consumers in a manner that permits consumers to understand the costs, benefits, and risks associated with the product or service, in light of the facts and circumstances." 12 U.S.C. § 5532(a). In prescribing such disclosure rules, section 1032 requires the Bureau to "consider available evidence about consumer awareness. understanding of, and responses to disclosures or communications about the risks, costs, and benefits of consumer financial products or services." 12 U.S.C. § 5532(c).

³⁸Pub. L. No. 111-203, § 1013(b)(3), § 1021(c), 124 Stat. 1376, 1969, 1980 (2010) (codified at 12 U.S.C. §§ 5493(b)(3), 5511(c)).

Securities and Exchange Commission

³⁶Pub. L. No. 90-321, 92 Stat. 3728 (1978) (codified as amended at 15 U.S.C. §§ 1693-1693r). CFPB issues and enforces Regulation E, which implements the Electronic Fund Transfer Act (EFTA). EFTA establishes basic rights, liabilities, and responsibilities of consumers who use electronic fund transfer services and of financial institutions that offer these services.

authority for violations of federal securities laws prohibiting fraud by any person in the purchase, offer, or sale of securities. SEC enforcement extends to virtual-currency-related securities transactions. Additionally, when companies offer and sell securities (including virtual-currencyrelated securities), they are subject to SEC requirements to either register the offering with SEC or qualify for a registration exemption. SEC reviews registration statements to ensure that potential investors receive adequate information about the issuer, the security, and the offering. Further, if a registered national securities exchange wanted to list a virtual-currency-related security, it could only do so if the listing complied with the exchange's existing rules or the exchange had filed a proposed rule change with SEC to permit the listing.

Virtual currencies may also have implications for other SEC responsibilities, as the following examples illustrate:

- SEC has examination authority for entities it regulates, including registered broker-dealers, to ensure compliance with federal securities laws, SEC rules and regulations, and BSA requirements. According to SEC officials, if a broker-dealer were to accept payments in virtual currencies from customers, this could raise potential antimoney-laundering issues that the broker-dealer would have to account for.
- SEC also regulates and has examination authority over investment advisers subject to its jurisdiction.³⁹ Under the Investment Advisers Act of 1940, investment advisers are fiduciaries.⁴⁰ To the extent that an investment adviser recommends virtual currencies or virtualcurrency-related securities, the investment adviser's federal fiduciary duty would govern this conduct.
- If registered broker-dealers held virtual currencies for their own account or an account of a customer, SEC would have to determine how to treat the virtual currencies for purposes of its broker-dealer financial responsibility rules, including the net capital rule.⁴¹

³⁹15 U.S.C. §§ 80b-2(a)(11), 80b-11(g)-(h).

⁴⁰See 15 U.S.C. § 80b-6(1)-(2); *SEC v. Capital Gains Research Bureau, Inc.,* et al., 375 U.S. 180 (1963).

⁴¹17 C.F.R. § 240.15c3-1. SEC's net capital rule requires all broker-dealers to maintain a minimum level of net capital consisting of highly liquid assets. Assets that are not liquid are deducted in full when computing net capital.

CFTC has the authority to regulate financial derivative products and their Commodity Futures Trading markets, including commodity futures and options.⁴² In addition, CFTC Commission investigates and prosecutes alleged violations of the Commodity Exchange Act and related regulations.⁴³ CFTC's mission is to protect market users and the public from fraud, manipulation, abusive practices, and systemic risk related to derivatives subject to the Commodity Exchange Act. CFTC's responsibilities with respect to virtual currencies depend partly on whether bitcoin or other virtual currencies meet the definition of a commodity under the Commodity Exchange Act.⁴⁴ CFTC officials said the agency would not make a formal determination on this issue until market circumstances require one. According to CFTC, such circumstances could include virtual-currency derivatives emerging or being offered in the United States or CFTC becoming aware of the existence of fraud or manipulative schemes involving virtual currencies. The officials said that if prospective derivatives that are backed by or denominated in virtual currencies that CFTC determines to be commodities emerge, CFTC's regulatory authorities would apply to those derivatives just as they would for any other derivative product subject to CFTC's jurisdiction. To carry out its regulatory responsibilities, CFTC would, among other things, evaluate the derivatives to ensure they were not susceptible to manipulation, review applications for new exchanges wishing to offer such derivatives, and examine exchanges offering these derivatives to ensure compliance with the applicable commodity exchange laws. Similar to SEC, CFTC has examination authority for BSA compliance-in this case directed at futures commission merchants and other futures

market intermediaries—and acceptance of virtual currency payments by

⁴²7 U.S.C. § 2. Financial derivatives are financial instruments whose value is based on one or more underlying reference items. They are used to hedge risk or to exchange a floating rate of return for a fixed rate of return. In the virtual currency context, a derivative might be used to reduce exposure to volatility in virtual currency exchange rates.

⁴³7 U.S.C. §§ 1-26; 17 C.F.R. chap. I.

⁴⁴The Commodity Exchange Act defines a commodity as certain agricultural goods and "all services, rights, and interests (except motion picture box office receipts, or any index, measure, value or data related to such receipts) in which contracts for future delivery are presently or in the future dealt in." 7 U.S.C. § 1a(9).

these entities could raise BSA compliance concerns.⁴⁵ Like SEC, CFTC would also have to make determinations about the capital treatment of virtual currencies if these entities held virtual currencies for their own account or an account of a customer.

Departments of Homeland Security and Justice Law enforcement agencies, including but not limited to DHS and DOJ component agencies and offices, have responsibilities to investigate a variety of federal crimes that may involve the use of virtual currencies and to support the prosecution of those who commit these crimes. Like traditional currencies, virtual currencies can facilitate a range of criminal activities, including fraud schemes and the sale of illicit goods and services, that may fall under the purview of federal law enforcement agencies.

The emergence of virtual currencies has had particular significance for financial crimes. According to DOJ officials, the main law enforcement interests with respect to virtual currencies are to (1) deter and prosecute criminals who use virtual currency systems to launder money (that is, move or hide money that either facilitates or is derived from criminal or terrorist activities); and (2) investigate and prosecute virtual currency services that themselves violate money transmission and money laundering laws.⁴⁶ A number of DOJ and DHS components, including the FBI, ICE-HSI, and Secret Service, investigate financial crimes as part of their broader responsibilities. In addition, DOJ's Asset Forfeiture and Money Laundering Section prosecutes money laundering violations, and DOJ and DHS manage the seizure and forfeiture of assets that represent the proceeds of, or were used to facilitate, federal crimes. Key laws that may apply to the use of virtual currencies in financial crimes include BSA,

⁴⁵Futures commission merchants are entities that solicit or accept orders for the purchase or sale of a commodity for future delivery on or subject to the rules of any exchange and that accept payment from or extend credit to those whose orders are accepted.

⁴⁶One example would be a centralized virtual currency system that allowed users to make untraceable funds transfers.

	as amended by Title III of the USA PATRIOT Act, and anti-money- laundering statutes. ⁴⁷
	Additionally, because virtual currencies operate over the Internet, they have implications for agency components that investigate and prosecute computer crimes (also called cybercrimes). For example, DOJ's Computer Crime and Intellectual Property Section stated that virtual currencies can be attractive to entities that seek to facilitate or conduct computer crimes over the Internet, such as computer-based fraud and identity theft. The section's responsibilities include improving legal processes for obtaining electronic evidence and working with other law enforcement agencies in improving the technological and operational means for gathering and analyzing electronic evidence. The FBI, Secret Service, and ICE-HSI also investigate computer crimes.
Virtual Currencies Present Regulatory, Law Enforcement, and Consumer Protection Challenges	The emergence of virtual currencies presents challenges to federal agencies responsible for financial regulation, law enforcement, and consumer and investor protection. These challenges stem partly from certain characteristics of virtual currencies, such as the higher degree of anonymity they provide and the ease with which they can be sent across borders. In addition, the growing popularity of virtual currencies has highlighted both risks and benefits for agencies to consider in carrying out their responsibilities.
Greater Anonymity	As previously noted, some virtual currency systems may provide a higher degree of anonymity than traditional payment systems because they do not require the disclosure of personally identifiable information (that is, information that can be used to locate or identify an individual, such as names or Social Security numbers) to transfer funds from one party to another. When transferring funds in the amount of \$3,000 or more between the bank accounts of two individuals, the banks involved are required by FinCEN regulations to obtain and keep the names and other
	 ⁴⁷Pub. L. No. 91-508, 84 Stat. 1114 (codified as amended at 12 U.S.C. §§ 1829(b), 1951-1959; 31 U.S.C. §§ 5311-5330); Pub. L. No. 107-56, tit. III, 115 Stat. 272, 296-342 (International Money Laundering Abatement and Anti-Terrorist Financing Act of 2001) (codified at 31 U.S.C. §§ 5301-5318A) (to prevent, detect, and prosecute international

⁽international Money Laundering Abatement and Anti-Terrorist Financing Act of 2001) (codified at 31 U.S.C. §§ 5301-5318A) (to prevent, detect, and prosecute international money laundering); see also Money Laundering Suppression Act of 1994, Pub. L. No. 103-325, §§ 401-413, 108 Stat. 2160, 2243-2255 (codified at 31 U.S.C. § 5330 and scattered sections of U.S.C.) (requires money transmitting businesses to register with Treasury).

information of the individuals, as well as information on the transaction itself.⁴⁸ The customer identification information collected by the banks helps create a paper trail of financial transactions that law enforcement agencies can use to detect illegal activity, such as money laundering or terrorist financing, and to identify and apprehend criminals.⁴⁹ However, in a transfer between two individuals using bitcoins (or a similar type of decentralized virtual currency) no personally identifiable information is necessarily disclosed either to the two individuals or a third-party intermediary.⁵⁰ As a result, virtual currencies may be attractive to parties seeking to protect personally identifiable information, maintain financial privacy, buy or sell illicit goods and services, or move or conceal money obtained by illegal means. Further, virtual currency exchangers or administrators may be used to facilitate money laundering if they do not collect identifying information from customers and retain other transaction information. For these reasons, law enforcement and federal financial regulatory agencies have indicated that virtual currencies can create challenges for agencies in detecting unlawful actions and the entities that carry them out. For example, the FBI has noted that because bitcoin does not have a centralized entity to monitor and report suspicious activity and process legal requests such as subpoenas, law enforcement agencies face difficulty in detecting suspicious transactions using bitcoins and identifying parties involved in these transactions.

Cross-Jurisdictional Nature Because they operate over the Internet, virtual currencies can be used globally to make payments and funds transfers across borders. In addition, according to agency officials, many of the entities that exchange traditional currencies for virtual currencies (or vice versa) are located outside of the United States. If these exchangers have customers located in the United States, they must comply with BSA and anti-moneylaundering requirements. Due to the cross-jurisdictional nature of virtual

⁴⁸31 C.F.R. § 1020.410.

⁴⁹Financial institutions are also required to obtain customer information to satisfy "knowyour-customer" or "customer due diligence" identification programs as part of their antimoney laundering obligations, and financial institutions must subject certain bank accounts held by non-U.S. persons to enhanced due diligence procedures. See 31 U.S.C. § 5318(i).

⁵⁰However, in a virtual currency transfer between individuals through a third-party intermediary (such as a virtual currency exchange), personally identifiable information is required to be collected if the transaction is for \$3,000 or more. This requirement became effective in 2011. We discuss this requirement in the next section of this report.

currency systems, federal financial regulatory and law enforcement agencies face challenges in enforcing these requirements and investigating and prosecuting transnational crimes that may involve virtual currencies. For example, law enforcement may have to rely upon cooperation from international partners to conduct investigations, make arrests, and seize criminal assets. Additionally, violators, victims, and witnesses may reside outside of the United States, and relevant customer and transaction records may be held by entities in different jurisdictions, making it difficult for law enforcement and financial regulators to access them. Further, virtual currency exchangers or administrators may operate out of countries that have weak legal and regulatory regimes or that are less willing to cooperate with U.S. law enforcement.

Balancing Risks and Benefits Virtual currency industry stakeholders have noted that virtual currencies present both risks and benefits that federal agencies need to consider in regulating entities that may be associated with virtual-currency-related activities. As previously noted, the risks include the attractiveness of virtual currencies to those who may want to launder money or purchase illicit goods and services. Another emerging set of risks involves consumer and investor protection—in particular, whether consumers and investors understand the potential drawbacks of buying, holding, and using virtual currencies or investing in virtual-currency-based securities. Consumers may not be aware of certain characteristics and risks of virtual currencies, including the following:

- Lack of bank involvement. Virtual currency exchanges and wallet providers are not banks. If they go out of business, there may be no specific protections like deposit insurance to cover consumer losses.⁵¹
- Stated limits on financial recourse. Some virtual currency wallet providers purport to disclaim responsibility for consumer losses associated with unauthorized wallet access. In contrast, credit and debit card networks state that consumers have no liability for fraudulent use of accounts.
- *Volatile prices.* The prices of virtual currencies can change quickly and dramatically (as shown previously in fig. 2).

Additionally, an SEC official told us that virtual-currency-based securities may be attracting individuals who are younger and less experienced than typical investors. The official expressed concern that younger investors

⁵¹We discuss examples of such losses in the next section of this report.

may lack the sophistication to properly assess the risks of such investments and the financial resources to recover from losses on the investments, including losses resulting from fraud schemes.⁵²

While virtual currencies present risks to consumers and investors, they also provide several potential benefits to consumers and business.

- Cost and speed. Decentralized virtual currency systems may, in some circumstances, provide lower transaction costs and be faster than traditional funds transfer systems because the transactions do not need to go through a third-party intermediary. The irrevocable feature of virtual currency payments may also contribute to lower transaction costs by eliminating the costs of consumer chargebacks.⁵³ Industry stakeholders have noted that cost and time savings may be especially significant for international remittances (personal funds immigrants send to their home countries), which sometimes involve sizeable fees and can take several days. In addition, industry stakeholders have indicated that the potentially lower costs of virtual currency transactions—for example, relative to credit and debit cards—may facilitate the use of micropayments (very small financial transactions) as a way of selling items such as online news articles, music, and smartphone applications.
- Financial privacy. To the extent that bitcoin (or other virtual currency) addresses are not publicly associated with a specific individual, peerto-peer virtual currency transactions can provide a greater degree of financial privacy than transactions using traditional payment systems, because no personally identifiable information is exchanged.⁵⁴
- Access. Because virtual currencies can be accessed anywhere over the Internet, they are a potential way to provide basic financial services to populations without access to traditional financial

⁵²The next section of this report discusses an example of a fraud scheme involving a virtual-currency-based security.

⁵³A chargeback is a payment reversal initiated by a consumer due, for example, to nondelivery of a purchased product.

⁵⁴As previously noted, that privacy may be lost if a connection is established between a bitcoin address and its owner.

	institutions, such as rural populations in developing countries. ⁵⁵ However, the potential benefit hinges on access to the Internet, which these populations may not have, and may be offset by the lack of protections against losses noted previously.
	Federal agency officials have acknowledged the need to consider both the risks and benefits of virtual currencies in carrying out their responsibilities. For example, the Director of FinCEN has testified that the emergence of virtual currencies has prompted consideration of vulnerabilities that these currencies create in the financial system and how illicit actors will take advantage of them. However, she also noted that innovation is an important part of the economy and that FinCEN needs to have regulation that mitigates concerns about illicit actors while minimizing regulatory burden. Similarly, the former Acting Assistant Attorney General for DOJ's Criminal Division has testified that law enforcement needs to be vigilant about the criminal misuse of virtual currency systems while recognizing that there are many legitimate users of those services. Balancing concerns about the illicit use of virtual currencies against the potential benefits of these technological innovations will likely be an ongoing challenge for federal agencies.
Agencies Have Taken Some Actions on Virtual Currencies, but Interagency Working Groups Have Not Focused on Consumer Risks	Federal financial regulators and law enforcement agencies have taken a number of actions related to the emergence of virtual currencies, including providing regulatory guidance, assessing anti-money-laundering compliance, and investigating crimes and violations that have been facilitated by the use of virtual currencies. However, interagency working groups addressing virtual currencies have not focused on consumer protection and have generally not included CFPB.

⁵⁵Some industry observers have suggested that virtual currency system protocols may have applications beyond financial transactions. For example, just as the bitcoin protocol transfers and records ownership rights to currency, it could, in theory, be used to transfer and record ownership rights to stocks, among other things.

FinCEN Has Issued Rules, Guidance, and Administrative Rulings Regarding Virtual Currencies

FinCEN has taken a number of actions in recent years to establish and clarify requirements for participants in virtual currency systems. For example, in July 2011, FinCEN finalized a rule that modified the definitions of certain money services businesses.⁵⁶ Among other things, the rule states that persons who accept and transmit currency, funds, or "other value that substitutes for currency," are considered to be money transmitters.⁵⁷ Additionally, in March 2013, FinCEN issued guidance that clarified the applicability of BSA regulations to participants in certain virtual currency systems.⁵⁸ The FinCEN guidance classified virtual currency exchangers and administrators as money services businesses and, more specifically, as money transmitters.⁵⁹ The guidance also specified that virtual currency users are not money services businesses.⁶⁰ As a result, the guidance clarified that virtual currency exchangers and administrators must follow requirements to register with FinCEN as money transmitters; institute risk assessment procedures and antimoney-laundering program control measures; and implement certain recordkeeping, reporting, and transaction monitoring requirements, unless an exception to these requirements applies.⁶¹ According to FinCEN officials, as of December 2013, approximately 40 virtual currency exchangers or administrators had registered with FinCEN.

⁵⁶Bank Secrecy Act Regulations; Definitions and Other Regulations Relating to Money Services Businesses, 76 Fed. Reg. 43585 (July 21, 2011).

⁵⁸FinCEN, *Application of FinCEN's Regulations to Persons Administering, Exchanging, or Using Virtual Currencies*, FIN-2013-G001, March 18, 2013. This guidance addresses convertible virtual currency—that is, virtual currency which either has an equivalent value in real currency or acts as a substitute for real currency.

⁵⁹According to FinCEN, virtual currency exchangers and administrators with U.S. customers must comply with BSA requirements, such as instituting anti-money-laundering controls, even if they are based outside of the United States.

⁶⁰FinCEN's guidance defines a virtual currency user as "a person who obtains convertible virtual currency and uses it to purchase real or virtual goods or services on the user's own behalf." Although a user is not considered to be a money transmitter, FinCEN warns that a user's activities must still comply with other federal and state laws and regulations.

⁶¹Most states also regulate money services businesses and some have taken steps to address virtual currencies. For example, New York is developing licensing and regulatory requirements specific to virtual currency exchanges and Texas has issued a memorandum describing how current licensing requirements apply to virtual currency exchanges. FinCEN coordinates with its state counterparts to encourage application of FinCEN's guidance on virtual currencies as part of this process.

⁵⁷31 C.F.R. § 1010.100(ff)(5)(i)(A).

In 2014, in response to questions from industry stakeholders, FinCEN issued administrative rulings to clarify the types of participants to which the March 2013 guidance applies.⁶² In January 2014, FinCEN issued rulings stating that the way in which a virtual currency is obtained is not material, but the way in which a person or corporation uses the virtual currency is. As a result, the rulings specify that two kinds of users are not considered money transmitters subject to FinCEN's regulations: miners who use and convert virtual currencies exclusively for their own purposes and companies that invest in virtual currencies exclusively as an investment for their own account.⁶³ However, the rulings specify that these two kinds of users may no longer be exempt from FinCEN's money transmitter requirements if they conduct their activities as a business service for others. The rulings also note that transfers of virtual currencies from these types of users to third parties should be closely scrutinized because they may constitute money transmission. In April 2014, FinCEN issued another administrative ruling, which states that companies that rent computer systems for mining virtual currencies are not considered money transmitters subject to FinCEN's regulations.⁶⁴

FinCEN has also taken additional steps to help ensure that companies required to register as money services businesses under FinCEN's March 2013 virtual currency guidance have done so. According to FinCEN officials, FinCEN has responded to letters from companies seeking clarification about their requirements. Also, officials told us that FinCEN has proactively informed other companies that they should register as money services businesses.

⁶⁴FinCEN, Application of Money Services Business Regulations to the Rental of Computer Systems for Mining Virtual Currencies, FIN-2014-R007, April 29, 2014.

⁶²FinCEN, Application of FinCEN's Regulations to Virtual Currency Mining Operations, FIN-2014-R001, January 30, 2014, and FinCEN, Application of FinCEN's Regulations to Virtual Currency Software Development and Certain Investment Activity, FIN-2014-R002, January 30, 2014.

⁶³For example, a company that purchases and sells virtual currencies whenever such purchases and sales make investment sense according to the company's business plan is acting as a virtual currency user, not a virtual currency exchange.

Some Financial Regulators Have Taken Actions Concerning Anti-Money-Laundering and Securities Law Compliance

As part of their oversight activities, NCUA and SEC have addressed situations involving virtual currencies, and other federal financial regulators have had internal discussions regarding virtual currencies. NCUA has had two supervisory situations in which credit unions were involved with activity related to virtual currencies. These situations emerged after reviews of credit unions found that their anti-moneylaundering and antifraud measures needed to be revised in light of activity involving virtual currency exchanges.

- In 2013, NCUA issued a preliminary warning letter to a federal credit union that provided account services to money services businesses that also served as bitcoin exchanges. The warning letter was based on various conditions that NCUA determined could undermine the credit union's stability. For example, the credit union did not have adequate anti-money-laundering controls in place for its money services business accounts. Further, the letter stated that the credit union should not have served money services businesses that were not part of the credit union's strategic plan, and that serving these businesses was not consistent with the credit union's charter, which called for serving the local community. The warning letter required the credit union to immediately cease all transactions with these money services business accounts and establish an appropriate BSA and anti-money-laundering infrastructure. As a result, the credit union ceased such activity and strengthened its BSA and anti-moneylaundering compliance program.
- In 2012, NCUA provided support to a state regulator's review of a credit union's commercial customer. The state regulator found that this commercial customer was a payment processor—that is, a payment network that allows any business or person to send, request, and accept money—that had customers that were bitcoin exchanges. According to NCUA, the state regulator worked with the credit union to ensure that its BSA compliance program was adequate to monitor and address the risks associated with payment processors that serve bitcoin exchanges. The state regulator also worked to ensure that the payment processor's risk management practices included sufficient antifraud and anti-money-laundering measures. The payment processor subsequently suspended all accounts that served virtual currency exchanges.

In addition, SEC has taken enforcement action against an individual and entity that are alleged to have defrauded investors through a bitcoindenominated Ponzi scheme.⁶⁵ The agency has also issued related investor alerts, has begun to review a registration statement from an entity that wants to offer virtual-currency-related securities, and is monitoring for potential securities law violations related to virtual currencies.

- In July 2013, SEC charged an individual and his company, Bitcoin Savings and Trust, with offering and selling securities in violation of the antifraud and registration provisions of securities laws.⁶⁶ Specifically, SEC alleges that the founder and operator defrauded investors through a bitcoin-denominated Ponzi scheme. The founder and operator allegedly promised investors up to 7 percent weekly interest. However, he allegedly used bitcoins from new investors to make purported interest payments and cover investor withdrawals on outstanding trust investments, diverted investors' bitcoins for day trading in his personal account on a bitcoin currency exchange, and exchanged investors' bitcoins for U.S. dollars to pay for personal expenses. SEC also alleges that Bitcoin Savings and Trust raised at least 700,000 bitcoins in investor funds, which amounted to more than \$4.5 million based on the average price of bitcoin in 2011 and 2012 when the investments were offered and sold. This case was still unresolved as of April 14, 2014.
- SEC's Office of Investor Education and Advocacy has issued two investor alerts on virtual currencies.⁶⁷ The first alert, issued in July 2013, warned about fraudulent investment schemes that may involve bitcoin and other virtual currencies.⁶⁸ The second alert, issued in May

http://www.nasaa.org/30631/informed-investor-advisory-virtual-currency.

⁶⁸http://www.investor.gov/news-alerts/investor-alerts/investor-alert-ponzi-schemes-using-virtual-currencies.

⁶⁵A Ponzi scheme is a type of investment fraud that involves the payment of purported returns to existing investors from funds contributed by new investors.

⁶⁶Securities and Exchange Commission v. Shavers, No. 413-CV-416 (E.D. Texas Aug. 6, 2013).

⁶⁷In addition, in March 2014, the Financial Industry Regulatory Authority, a self-regulatory organization for the securities industry, issued an investor alert about the risks of buying, using, and speculating in virtual currencies and the potential for related scams. See http://www.finra.org/Investors/ProtectYourself/InvestorAlerts/FraudsAndScams/P456458. Also, in April 2014, the North American Securities Administrators Association issued an investor advisory on virtual currencies, related investment risks, and the types of investments that might involve virtual currencies. See

2014, addressed fraud and other investment risks related to virtual currencies.⁶⁹

- SEC staff have begun to review a registration statement from a company that wants to conduct a public offering of virtual-currencyrelated securities and has received notice of a company offering a private virtual-currency-related security, relying upon an exemption from registration. In July 2013, the Winklevoss Bitcoin Trust filed a registration statement for an initial public offering of its securities. The Trust is structured similarly to an exchange-traded fund and will hold bitcoins as its only assets.⁷⁰ The Trust filed amended registration statements in October 2013 and February 2014, but the registration statement remains pending as of April 14, 2014, meaning that the Trust is not yet permitted to sell its securities in a public offering. Also, in October 2013, Bitcoin Investment Trust, a bitcoin-denominated pooled investment fund affiliated with SecondMarket, Inc. and available only to accredited investors, filed a notice with SEC indicating that it had sold securities in an exempt offering in reliance on Rule 506(c) of the Securities Act.⁷¹ Rule 506(c) allows an issuer to raise an unlimited amount of money, but imposes restrictions on who can invest in the offering and requires the issuer to take reasonable steps to verify that those investing are accredited investors.⁷²
- SEC staff are also monitoring the Internet and other sources, such as referrals from other agencies, for potential securities law violations involving bitcoin and other virtual currencies.

⁶⁹http://www.investor.gov/news-alerts/investor-alerts/investor-alert-bitcoin-other-virtual-currency-related-investments.

⁷⁰Exchange-traded funds are commonly structured as open-end investment companies and offer investors a proportionate share in a pool of stocks, bonds, and other assets.

⁷¹Rule 506(c) is one of the exemptive rules under Regulation D that allow some businesses to offer and sell their securities without having to register the offer and sale of securities with SEC. Regulation D is designed to (1) simplify the previously existing rules and regulations, (2) eliminate any unnecessary restrictions that those rules and regulations place on small business issuers, and (3) achieve uniformity between state and federal exemptions to facilitate capital formation consistent with protecting investors.

⁷²17 C.F.R. § 230.506(c). Accredited investors include, among others, individuals whose net worth is more than \$1 million (not including the value of their primary residence) or whose individual income exceeds at least \$200,000 for the most recent 2 years (or joint income with a spouse exceeding \$300,000 for those years) and a reasonable expectation of the same income level in the current year. It also includes certain types of entities, such as insurance companies, banks, and corporations with assets exceeding \$5 million. 17 C.F.R. § 230.501(a).

Further, all of the federal financial regulatory agencies we interviewed have had internal discussions on how virtual currencies work and what implications the emergence of virtual currencies might have for their responsibilities. While agencies generally told us that their conversations have been informal and ad hoc, some efforts have been more organized:

- In 2013, the Federal Reserve took several steps to share information on virtual currencies among the Board of Governors and the 12 Federal Reserve Banks. Among other things, the Board of Governors' BSA and anti-money-laundering specialist conference included a session focused on FinCEN's virtual currency guidance and recent law enforcement actions. The Board of Governors also circulated general information about virtual currencies within the Federal Reserve System to use in answering questions from media and the public about virtual currencies and federal financial regulatory actions to date.
- In 2013, SEC formed an internal Digital Currency Working Group, which aims to foster information sharing internally and externally. According to SEC, the working group consists of approximately 50 members from among SEC's divisions and offices.
- In 2012, FinCEN held three internal information-sharing events on virtual currencies. These events covered issues including how virtual currencies compare to traditional currencies and risks related to emerging payment systems such as virtual currencies.

Law Enforcement Agencies Have Taken Actions against Parties Alleged to Have Used Virtual Currencies to Facilitate Crimes Law enforcement agencies have taken actions against parties involved in the illicit use of virtual currencies to facilitate crimes. These parties have included administrators and users of centralized virtual currency systems designed to facilitate money laundering or other crimes, parties who have used virtual currencies to buy or sell illicit goods and services online, and virtual currency exchanges and online payment processors operating without the proper licenses.

 In 2013 and 2014, law enforcement agencies took actions against Silk Road, a black market website that allegedly accepted bitcoin as the sole payment method for the purchase of illegal goods and services. The website contained over 13,000 listings for controlled substances as well as listings for malicious software programs, pirated media content, fake passports, and computer hacking services (see fig.3). The FBI; Drug Enforcement Administration (DEA); IRS; ICE-HSI; the Bureau of Alcohol, Tobacco, Firearms, and Explosives; the Secret Service; the U.S. Marshals Service; and Treasury's Office of Foreign Assets Control investigated the case together, along with officials from New York as well as Australia, Iceland, Ireland, and France. In September and October 2013, law enforcement shut down the Silk Road website and seized approximately 174,000 bitcoins, which the FBI reported were worth approximately \$34 million at the time of seizure.⁷³ In February 2014, DOJ indicted Silk Road's alleged owner and operator on charges including narcotics conspiracy, engaging in a continuing criminal enterprise, conspiracy to commit computer hacking, and money laundering conspiracy.

- In May 2013, law enforcement agencies seized the accounts of a U.S.-based subsidiary of Mt. Gox, a now-defunct Tokyo-based virtual currency exchange with users from multiple countries including the United States, on the basis that the subsidiary was operating as an unlicensed money services business. The seizure included U.S. bank accounts of Mt. Gox that were held by a private bank and Dwolla, an online payment processor that allegedly allowed users to buy and sell bitcoins on Mt. Gox. According to ICE-HSI, Mt. Gox had moved funds into numerous online black markets, the bulk of which were associated with the illicit purchase of drugs, firearms, and child pornography. At the direction of the U.S. Attorney's office, ICE-HSI ordered Dwolla to stop all payments to Mt. Gox and seized \$5.1 million from the Mt. Gox subsidiary's U.S. accounts.
- Also in May 2013, law enforcement agencies shut down Liberty Reserve, a centralized virtual currency system that was allegedly designed and frequently used to facilitate money laundering and had its own virtual currency. Secret Service, ICE-HSI, and IRS investigated the case together, along with officials from 16 other countries. To shut down the site, FinCEN identified Liberty Reserve as a financial institution of primary money laundering concern under section 311 of the USA PATRIOT Act, effectively cutting it off from the U.S. financial system.⁷⁴ DOJ then charged Liberty Reserve with operating an unlicensed money transmission business and with money laundering for facilitating the movement of more than \$6 billion

⁷³As of March 31, 2014, these bitcoins were worth about \$80 million, according to bitcoin prices from https://www.coindesk.com.

⁷⁴31 U.S.C. § 5318A. Section 311 of the USA PATRIOT Act grants the Secretary of the Treasury the authority, upon finding that reasonable grounds exist for concluding that a foreign jurisdiction, institution, class of transaction, or type of account is of "primary money laundering concern," to require domestic financial institutions and financial agencies to take certain "special measures" to address the primary money laundering concern.
in illicit proceeds.⁷⁵ As of April 2014, this investigation had produced \$40 million in seizures and had resulted in the arrests of five individuals.

- In April 2013, law enforcement agencies filed a civil asset forfeiture complaint against Tcash Ads Inc., an online payment processor that allegedly enabled users to make purchases anonymously from virtual currency exchanges, with operating an unlicensed money services business. Additionally, law enforcement agencies seized the bank accounts of Tcash Ads Inc. The Secret Service worked on the case with FinCEN and DOJ's Asset Forfeiture and Money Laundering Section.
- From October 2010 through November 2012, law enforcement agencies convicted three organizers of a worldwide conspiracy to use a network of virus-controlled computers that deployed e-mail spam designed to manipulate stock prices. The organizers paid the spammers \$1.4 million for their illegal services via the centralized virtual currency e-Gold and wire transfers. Charges included conspiring to further securities fraud using spam, conspiring to transmit spam through unauthorized access to computers, and four counts of transmission of spam by unauthorized computers.

⁷⁵This case is being prosecuted jointly by the DOJ Criminal Division's Asset Forfeiture and Money Laundering Section and the U.S. Attorney's Office for the Southern District of New York.



Figure 3: Screen Shot of the Silk Road Website

Source: U.S. Immigration and Customs Enforcement.

Law enforcement agencies have also taken other actions to help support investigations involving the illicit use of virtual currencies, including the following examples.

- The FBI has produced numerous criminal intelligence products addressing virtual currencies. These intelligence products have generally focused on cases involving the illicit use of virtual currencies, ways in which virtual currencies have been or could be used to facilitate crimes, and the related challenges for law enforcement. The FBI shares these products with foreign, state, and local law enforcement partners as appropriate.
- Through standing bilateral agreements governing the exchange of law enforcement information, ICE-HSI is arranging meetings with various international partners to exchange intelligence and garner operational support on virtual currency issues.

 ICE-HSI also developed the Illicit Digital Economy Program, which aims to target the use of virtual currencies for money-laundering purposes by defining and organizing the primary facets of the digital economy, building internal capacity, training and developing agents and analysts, engaging other agencies, and promoting public-private partnerships.

Interagency Working Groups Have Begun to Address Virtual Currencies, but Have Not Emphasized Consumer Risks or Generally Included CFPB

Federal agency efforts to collaborate on virtual currency issues have involved creating a working group specifically focused on virtual currency, leveraging existing interagency mechanisms, and sharing information through informal interagency channels. For example, in 2012, the FBI formed the Virtual Currency Emerging Threats Working Group (VCET), an interagency working group that includes other DOJ components, FinCEN, ICE-HSI, SEC, Secret Service, Treasury, and other relevant federal partners. The purpose of VCET is to leverage members' expertise to address new virtual currency trends, address potential implications for law enforcement and the U.S. intelligence community, and mitigate the cross-programmatic threats arising from illicit actors' use of virtual currency systems. The VCET meets about once every 3 months.

Federal agencies have also begun to discuss virtual currency issues in existing interagency working groups that address broader topics such as money laundering, electronic crimes, and the digital economy, as follows:

- The BSA Advisory Group—which is chaired by FinCEN and includes the prudential banking regulators, Treasury, federal and state law enforcement and regulatory agencies, and industry representatives has addressed virtual currency issues in a number of ways. In May 2013, FinCEN provided a briefing on bitcoin, and in December 2013 three stakeholders from the virtual currency industry gave presentations on their business models and regulatory challenges. In addition, the BSA Advisory Group invited a representative of the virtual currency industry to join the group in 2014.
- The Federal Financial Institutions Examination Council (FFIEC) Bank Secrecy Act/Anti-Money-Laundering Working Group—which is currently chaired by OCC and includes the prudential banking regulators and CFPB—is in the process of revising the current (2010)

FFIEC BSA/Anti-Money Laundering Examination Manual.⁷⁶ The revisions related to virtual currencies may include information on FinCEN's March 2013 guidance and regulatory expectations that depository institutions should undertake a risk assessment with a particular focus on the money laundering risks posed by new products and services.

- The Secret Service-sponsored Electronic Crimes Task Forces (ECTF) includes 35 Secret Service field offices; federal law enforcement agencies such as ICE-HSI; and members of the private sector, academia, and state and local law enforcement.⁷⁷ This group's mission is to prevent, detect, and investigate electronic crimes, including those involving virtual currency. This group has conducted computer forensics and other investigative activity on various virtual currencies and made arrests of individuals who have used virtual currencies as part of their criminal activities. This group has also held quarterly meetings on virtual currencies to discuss legal and regulatory issues and trends in crimes involving virtual currencies.
- The Digital Economy Task Force was established in 2013 by Thomson Reuters (a multinational media and information firm) and the International Centre for Missing & Exploited Children.⁷⁸ This task force includes members from both the public and private sectors. Task force members from the federal government include representatives from the FBI, ICE-HSI, Secret Service, the Department of State, and the United States Agency for International Development. This group published a report in March 2014 on the benefits and challenges of

⁷⁶FFIEC is a formal interagency body empowered to prescribe uniform principles, standards, and report forms for the federal examination of financial institutions by the Federal Reserve, FDIC, NCUA, OCC, and CFPB, and to make recommendations to promote uniformity in the supervision of financial institutions.

⁷⁷The Secret Service was mandated by the USA PATRIOT Act to establish a nationwide network of Electronic Crimes Task Forces. Pub. L. 107-56, § 105, 115 Stat 272, 277 (2001) (codified at 18 U.S.C. § 3056 note). The goal of the network is to bring together federal, state, and local law enforcement, as well as prosecutors, private industry, and academia to prevent, detect, and investigate various forms of electronic crime.

⁷⁸The International Centre for Missing & Exploited Children is a nonprofit corporation that leads a movement to protect children from sexual exploitation and abduction. The Centre is involved in virtual currency issues because of connections between digital technologies that facilitate anonymity and commercial child pornography, sexual exploitation, and sex trafficking.

the digital economy.⁷⁹ Among other things, the report recommended continuing private and public research into the digital economy and illegal activities, investing in law enforcement training, rethinking investigative techniques, fostering cooperation between agencies, and promoting a national and global dialogue on policy related to virtual currencies.

A number of other existing interagency working groups have discussed or addressed virtual currency issues to some extent. See appendix II for more information on these groups.

Federal agencies have also started to collaborate outside of these working groups to help improve their knowledge of issues related to the emergence of virtual currencies and share pertinent information with various agencies.

- FinCEN and SEC have hosted meetings with industry representatives and consultants to discuss how virtual currency systems such as bitcoin and Ripple work and what legal, regulatory, technology, and law enforcement issues they present. These agencies have invited officials from other federal agencies to these sessions.
- FinCEN consulted with financial regulators and law enforcement agencies as it was formulating its March 2013 guidance on virtual currencies. These agencies included CFPB, CFTC, DEA, FBI, ICE-HSI, IRS, the prudential banking regulators, SEC, and the Secret Service.
- SEC notified CFTC of its review of the Winklevoss Bitcoin Trust registration statement.
- FinCEN issued a Networking Bulletin on cryptocurrencies in March 2013 to provide details to law enforcement agencies and assist them in following money moving between virtual currency channels and the traditional U.S. financial system. Among other things, the bulletin addressed the role of entities that facilitate the purchase and exchange of virtual currencies and the types of records these entities maintain that could be useful to investigative officials. Also, the Networking Bulletin elicited information from its recipients, which in turn helped FinCEN issue additional analytical products of a tactical nature to inform law enforcement operations. FinCEN has also shared

⁷⁹Digital Economy Task Force, *The Digital Economy: Potential, Perils, and Promises* (March 2014).

this information with several regulatory and foreign financial intelligence unit partners.

 CFPB officials said they had recently conferred on virtual currency issues with a number of domestic and international regulators, including the Federal Reserve Bank of San Francisco, the Federal Trade Commission, NCUA, OCC, Treasury, New York State's Department of Financial Services, and the European Banking Authority. In addition, the officials said they had met with industry participants on these issues and conferred with interested academic and consumer group stakeholders, as well as law firms, consultancies, and industry associations.

Although there are numerous interagency collaborative efforts that have addressed virtual currency issues in some manner, interagency working groups have not focused on consumer protection issues. Rather, as previously discussed, these efforts have focused on BSA and anti-moneylaundering controls and investigations of crimes in which virtual currencies have been used. In addition, CFPB's involvement in interagency working groups that address virtual currencies has been limited. GAO's key practices on collaboration state that it is important to include relevant participants in interagency collaborative efforts in order to ensure, among other things, that these participants contribute knowledge, skills, and abilities to the outcomes of the effort.⁸⁰ In addition, these key practices state that once an interagency group has been established, it is important to reach out to potential participants who may have a shared interest in order to ensure that opportunities for achieving outcomes are not missed.⁸¹ CFPB might be a relevant participant in a broader set of collaborative efforts on virtual currencies because virtual currency systems provide a new way of making financial transactions, and CFPB's responsibilities include ensuring that consumers have timely and understandable information to make responsible decisions about financial transactions.⁸² Further, CFPB's strategic goals include helping consumers

⁸⁰GAO, *Managing for Results: Key Considerations for Implementing Interagency Collaborative Mechanisms*, GAO-12-1022 (Washington, D.C.: Sept. 27, 2012).

⁸¹GAO, Managing for Results: Implementation Approaches Used to Enhance Collaboration in Interagency Groups, GAO-14-220 (Washington, D.C.: Feb. 14, 2014).

⁸²CFPB (via the Office of Financial Education) is responsible for educating and empowering consumers to make better-informed financial decisions. Pub. L. No. 111-203, § 1013(d), 124 Stat. 1376, 1970 (2010).

understand the costs, risks, and tradeoffs of financial decisions and surfacing financial trends and emergent risks relevant to consumers.

Although interagency working groups addressing virtual currencies have not focused on consumer protection issues, recent events have highlighted the risks individuals face in buying and holding these currencies. For example, notable examples of bitcoin thefts by computer hackers have occurred in the past few years, including the theft of more than 35,000 bitcoins from a virtual wallet provider in April 2013 and 24,000 bitcoins from a bitcoin exchange in September 2012.⁸³ More recently, in February 2014, Mt. Gox filed for bankruptcy, stating that a security breach resulted in the loss of 850,000 bitcoins, the vast majority of which belonged to its customers. These bitcoins were worth more than \$460 million when Mt. Gox filed for bankruptcy.⁸⁴ Mt. Gox subsequently reported that it had found 200,000 of these bitcoins in an unused virtual wallet.

Certain parties have taken actions to inform consumers about the potential risks associated with virtual currencies, but these actions have occurred outside of federal interagency efforts and have not included CFPB. In April 2014, the Conference of State Bank Supervisors and the North American Securities Administrators Association issued joint model consumer guidance to assist state regulatory agencies in educating consumers about virtual currencies and the risks of purchasing, exchanging, and investing in virtual currencies.⁸⁵ Additionally, from February through April 2014, a number of states issued consumer alerts about virtual currencies.⁸⁶ On the international front, the European

⁸⁶These states include Alabama, California, Florida, Hawaii, Maryland, Massachusetts, Nevada, Washington, and Wisconsin.

⁸³Congressional Research Service, *Bitcoin: Questions, Answers, and Analysis of Legal Issues* (Washington, D.C.: Dec. 20, 2013).

⁸⁴Data from Coindesk.com. These bitcoins were worth approximately \$390 million as of March 31, 2014. https://www.coindesk.com.

⁸⁵For the Conference of State Bank Supervisors and the North American Securities Administrators Association joint model consumer guidance, see http://www.csbs.org/legislative/testimony/Documents/ModelConsumerGuidance---Virtual%20Currencies.pdf.

Banking Authority issued a warning to consumers in December 2013 about the risks involved in buying or holding virtual currencies.⁸⁷

Federal interagency working groups addressing virtual currency issues have not focused on consumer protection, and CFPB has generally not participated in these groups, for a number of potential reasons. For example, the extent to which individuals using virtual currencies are speculative investors or ordinary consumers is unclear, and CFPB has received few consumer complaints about these currencies.⁸⁸ In addition, incidents involving the use of virtual currencies for illicit purposes have made money laundering and other law enforcement issues primary concerns, and existing interagency working groups are primarily composed of agencies that share responsibilities for these matters. However, emerging consumer risks indicate that interagency collaborative efforts may need to place greater emphasis on consumer protection issues in order to address the full range of challenges posed by virtual currencies. Additionally, without CFPB's participation, interagency working groups are not fully leveraging the expertise of the lead consumer financial protection agency, and CFPB may not be receiving information that it could use to assess the risks that virtual currencies pose to consumers.

Conclusions

Bitcoin and other virtual currencies are technological innovations that provide users with certain benefits but also pose a number of risks. Because virtual currencies touch on the responsibilities of multiple federal agencies, addressing these risks will require effective interagency collaboration. Thus far, interagency efforts have had a law enforcement focus, reflecting the attractiveness of virtual currencies to those who may want to launder money or purchase black market items. If virtual currencies become more widely used, other types of regulatory and enforcement issues may come to the forefront. For example, recent events suggest that consumer protection is an emerging risk, as

⁸⁷European Banking Authority, *Warning to Consumers on Virtual Currencies*, EBA/WRG/2013/01, Dec. 12, 2013. See http://www.eba.europa.eu/-/eba-warns-consumers-on-virtual-currencies.

⁸⁸CFPB's complaint intake system is not specifically geared towards virtual currency complaints. However, in February 2014, CFPB ran a query of its Consumer Complaint Database to determine the number of complaints that had mentioned virtual currency or bitcoin and found that only 14 out of about 290,000 complaints met that condition.

	evidenced by the loss or theft of bitcoins from exchanges and virtual wallet providers and consumer warnings issued by nonfederal and non-U.S. entities. However, federal interagency working groups addressing virtual currencies have thus far not emphasized consumer-protection issues, and participation by the federal government's lead consumer financial protection agency, CFPB, has been limited. Therefore, these efforts may not be consistent with key practices that can benefit interagency collaboration, such as including all relevant participants to ensure that their knowledge, skills, and abilities contribute to the outcomes of the effort. As a result, future interagency efforts may not be in a position to address consumer risks associated with virtual currencies in the most timely and effective manner.
Recommendation for Executive Action	To help ensure that federal interagency collaboration on virtual currencies addresses emerging consumer protection issues, we recommend that the Director of CFPB (1) identify which interagency working groups could help CFPB maintain awareness of these issues or would benefit from CFPB's participation; and (2) decide, in coordination with the agencies already participating in these efforts, which ones CFPB should participate in.
Agency Comments	We provided a draft of this report to CFPB, CFTC, DOJ, DHS, FDIC, the Federal Reserve, NCUA, OCC, SEC, and Treasury for review and comment. CFPB and NCUA provided written comments, which are reprinted in appendixes III and IV. In addition, CFPB, CFTC, DHS, DOJ, the Federal Reserve, NCUA, OCC, SEC, and Treasury provided technical comments, which we incorporated into the report where appropriate.
	In its letter, CFPB concurred with our recommendation to identify and participate in pertinent interagency working groups addressing virtual currencies. CFPB stated that, to date, these groups have primarily focused on BSA concerns, anti-money-laundering controls, and the investigation of crimes involving virtual currencies. CFPB said that, as a result, its participation in these working groups has been limited. CFPB also stated that as consumer protection concerns have increased in recent months, its own work on virtual currencies and the work of other financial regulators in this area could benefit from a collaborative approach.
	In its letter, NCUA said that the report provides a clear discussion of the risks related to virtual currencies as well as a survey of current efforts in the regulatory community to address the related policy issues. NCUA also

expressed support for increasing emphasis on consumer protection issues pertaining to virtual currencies.

As agreed with your offices, unless you publicly announce the contents of this report earlier, we plan no further distribution until 30 days from the report date. At that time, we will send copies of this report to CFPB, CFTC, DOJ, DHS, FDIC, the Federal Reserve, NCUA, OCC, SEC, Treasury, interested congressional committees and members, and others. This report will also be available at no charge on our website at http://www.gao.gov.

If you or your staff have any questions concerning this report, please contact me at (202) 512-8678 or evansl@gao.gov. Contact points for our Offices of Congressional Relations and Public Affairs may be found on the last page of this report. GAO staff who made major contributions to this report are listed in appendix V.

. L Erono, Jr.

Lawrance L. Evans, Jr. Director, Financial Markets and Community Investment

Appendix I: How Bitcoins Enter into Circulation and Are Used in Transactions

This appendix shows how bitcoins enter into circulation through "mining," how transactions are conducted, and how miners verify transactions (see fig. 4).

Figure 4: How Bitcoins Enter into Circulation and Are Used in Transactions

Bitcoin Miners

Bitcoin miners essentially serve two purposes: 1) generating new bitcoins to enter into circulation and 2) verifying transactions by ensuring that they occurred and did not involve double spending of a bitcoin. Over time, the computer processing power needed to mine new bitcoins has increased to the point where mining requires specialized computer hardware and has become increasingly consolidated into large mining pools.



Mining

Bitcoins are created and first enter into circulation through a process known as mining. Bitcoin miners install software on their computers, which they use to solve complex math problems that verify transactions for the bitcoin network. The miner or mining pool that successfully solves the problems is rewarded with newly created bitcoins



Addresses and Wallets

Bill's bitcoin balances are associated with his bitcoin addresses (long strings of numbers and letters). Bill stores his bitcoin addresses in his virtual wallet (a program that saves bitcoin addresses on a user's computer or other data storage device, or online via a wallet service provided by an exchange or third-party virtual wallet provider). Bitcoin users can have multiple wallets, and each wallet can hold multiple bitcoin addresses.



Making a Peer-to-Peer Purchase with Bitcoins

Bill wants to buy a t-shirt from Carol, who accepts bitcoins. To conduct the transaction, Carol provides her bitcoin address to Bill, and Bill authorizes the transaction with his private key (essentially a secret code that proves Bill's control over his bitcoin address).



Verifying the Transaction

Bill and Carol's transaction is bundled into blocks with other transactions and verified by bitcoin miners. Within minutes, Bill's bitcoins are assigned to Carol's address and the transaction is registered in a public ledger called the "blockchain." The miner or mining pool that successfully solved the math problems to verify the block containing Bill and Carol's transaction is rewarded with newly created bitcoins.

Source: GAO

Appendix II: Interagency Working Groups that Have Addressed Virtual Currency Issues

In this appendix, we present some of the interagency working groups (including task forces and other interagency collaborative bodies) that have discussed virtual currency issues, and in some cases, taken specific actions. This list is based on information we obtained from the federal financial regulatory and law enforcement agencies we met with and is not intended to be an exhaustive list.

Table 1: Interagency Working Groups that Have Addressed Virtual Currency Issues, as of April 2014

Working group	Participating agencies	Mission and goals	Ways in which group addressed virtual currencies
Bank Secrecy Act Advisory Group (BSAAG)	FinCEN (lead); CFTC; DEA; DOJ Criminal Division; FBI; FDIC; Federal Reserve; ICE- HSI; IRS; NCUA; OCC; Office of National Drug Control Policy; SEC; Secret Service; and U.S. Postal Service; as well as representatives of financial institutions; trade groups; self-regulatory organizations; and state regulatory agencies.	This public-private group serves as a means by which the Secretary of the Treasury receives advice on the manner in which reporting requirements in BSA should be modified to enhance the ability of law enforcement agencies to use the information. It also informs private sector representatives of law enforcement's uses of BSA reports provided by financial institutions.	 Meetings have covered issues related to virtual currencies: The May 2013 meeting included a briefing on the bitcoin virtual currency system. The December 2013 meeting included a panel of virtual currency industry representatives who discussed business models and regulatory compliance challenges. In April 2014, a meeting of the BSAAG Illicit Finance Committee included a presentation on vulnerabilities and challenges related to virtual currencies, as well as opportunities to enhance collective anti-money-laundering efforts and information sharing. In addition, BSAAG invited a representative of the virtual currency industry to join the group in 2014.
Digital Economy Task Force	Thomson Reuters and the International Centre for Missing & Exploited Children (lead); FBI; ICE-HSI; Secret Service; Department of State; and United States Agency for International Development (USAID); as well as members of the private sector and academia.	This group's mission is to educate the public, work collaboratively across stakeholder groups, and balance the convenience of the digital currencies with controls to combat illegal activity.	Created in September 2013, this task force has formed working groups on such issues as safeguarding human rights; regulation; interagency coordination; and law enforcement. In March 2014, the task force published a report on the benefits and challenges of the digital economy. ^a Among other things, the report recommended private and public sector efforts to continue research into the digital economy and illegal activities; investing in law enforcement training; rethinking investigative techniques; fostering cooperation between agencies; and promoting a national and global dialogue on policy.

Working group	Participating agencies	Mission and goals	Ways in which group addressed virtual currencies
Electronic Crimes Task Forces (ECTF) and Working Groups	35 Secret Service field offices (lead) and federal law enforcement agencies such as ICE-HSI, as well as members of the private sector, academia, and state and local law enforcement.	The mission of these groups is to prevent, detect, and investigate various forms of electronic crime, including potential terrorist attacks against critical infrastructure and financial payment systems.	 ECTFs address issues concerning virtual currencies as one of a variety of subjects related to the investigations into electronic crime. Specifically, ECTFs have: conducted computer forensics and other investigative activity concerning various virtual currencies; made arrests of individuals who have used virtual currencies as part of their criminal activities; and discussed virtual currencies at quarterly meetings, covering topics such as types of virtual currencies and related legal and regulatory issues, trends in criminal uses, and methods for conducting investigations.
Federal Financial Institutions Examination Council (FFIEC) BSA/Anti- Money-Laundering Working Group ^b	OCC (rotating chair), CFPB; FDIC; Federal Reserve; NCUA; and the State Liaison Committee are voting members. ^c	FFIEC prescribes uniform principles, standards, and report forms for the federal examination of financial institutions by the prudential banking regulators—FDIC, Federal Reserve, NCUA, and OCC—and makes recommendations to promote uniformity in the supervision of financial institutions. Within this context, the FFIEC BSA/Anti-Money- Laundering Working Group's mission is to enhance coordination of BSA/anti- money-laundering training, guidance, and policy.	The BSA/Anti-Money-Laundering Working Group is leading the revision of the current (2010) FFIEC BSA/Anti-Money Laundering Examination Manual. Revisions related to virtual currencies may include information on FinCEN's March 2013 guidance; a brief note describing Internet-based electronic cash, which includes virtual currency; and regulatory expectations that banks should undertake a risk assessment with a particular focus on the money-laundering risks posed by new products, services, and technologies.

Working group	Participating agencies	Mission and goals	Ways in which group addressed virtual currencies
Financial Action Task Force (FATF)	FATF is an international intergovernmental organization with 36 member countries, including the U.S. Treasury as the lead agency of the U.S. delegation. Other U.S. delegation participants include DOJ's Asset Forfeiture and Money Laundering Section; DHS (including ICE-HSI); SEC; IRS; and the Department of State.	This group sets standards and promotes effective implementation of legal, regulatory, and operational measures for combating money laundering, and the financing of terrorism and proliferation.	 In February 2014, FATF developed a discussion paper on virtual currencies, which described virtual currency systems, participants, and some of the major virtual currencies such as bitcoin, and proposed a common set of terms and conceptual framework for analyzing virtual currencies. The paper also discussed the potential legitimate uses of virtual currencies, the risks these currencies may pose, and the different regulatory approaches countries are taking to address virtual currencies. The U.S. delegation prepared the paper together with delegations from Australia, Canada, Russia, and the United Kingdom. As of April 2014, the discussion paper was not yet public. In March 2014, FATF included a discussion of virtual currencies as part of the Private Sector Consultative Forum, which included experts on virtual currencies. The group discussed how virtual currencies and their exchangers operate; the associated money laundering and terrorist financing risks; what measures countries and financial institutions are taking to assess and mitigate those risks; and what regulatory approaches are currently being taken.

Working group	Participating agencies	Mission and goals	Ways in which group addressed virtual currencies
Interagency Bank Fraud Enforcement Working Group	DOJ (Criminal Division lead, as well as the Asset Forfeiture and Money Laundering Section, Executive Office for U.S. Attorneys, Executive Office for U.S. Trustees, and FBI); CFPB; CFTC; Department of Housing and Urban Development; DHS (ICE-HSI and Secret Service); Export- Import Bank; Farm Credit Administration; FDIC; Federal Housing Finance Agency; Federal Reserve; IRS; NCUA; OCC; SEC; Treasury (Bureau of Public Debt, FinCEN, Office of Inspector General, Office of Critical Infrastructure Protection, Office of Financial Stability, and Special Inspector General for the Troubled Asset Relief Program); U.S. Postal Inspection Service; and the District of Columbia Department of Insurance, Securities, and Banking.	This group's mission is to share information on significant trends, developments, and other issues in financial institution fraud and, as appropriate, identify and carry out projects of common interest to the working group's members.	The working group has occasionally discussed virtual currencies in the past year. Discussions to date have aimed to educate and inform members about virtual currencies. Planned activities include a presentation on the IRS notice addressing the status of virtual currencies under federal tax law. Within the Interagency Bank Fraud Working Group, the Payments Fraud Working Group has also addressed virtual currencies. The June 2013 meeting included presentations on e-Gold, the Liberty Reserve indictment, and FinCEN's guidance on how BSA regulations apply to participants in certain virtual currency systems.

Working group	Participating agencies	Mission and goals	Ways in which group addressed virtual currencies
International Organized Crime Intelligence and Operations Center (IOC-2)	DOJ (lead, including the Bureau of Alcohol, Firearms and Explosives; Criminal Division, DEA, and FBI); DHS (ICE-HSI and Secret Service); IRS-Criminal Investigation; Department of Labor (Office of Inspector General); Department of State (Bureau of Diplomatic Security); and U.S. Postal Inspection Service.	This group's mission is to significantly disrupt and dismantle transnational criminal organizations posing the greatest threat to the United States. The group does so by (1) deconflicting and analyzing transnational organized crime information and intelligence; (2) disseminating information and intelligence to support law enforcement operations, investigations, prosecutions, and forfeiture proceedings; and (3) coordinating jurisdictional and multiagency operations, investigations and prosecutions.	 IOC-2 supports member-agency investigations of both virtual currency administrators that are suspected of violating U.S. law and individuals who are suspected of using virtual currencies to commit crimes. Specifically, IOC-2 assists its member agencies by: sharing investigative details that will serve to deconflict current investigative and prosecutorial targets; identifying current trends in the illicit use of virtual currencies; sharing best practices in developing investigative and prosecutorial strategies; discussing investigative challenges and solutions; identifying tools, points of contact, and other areas of interest that offer assistance and serve as force multipliers in supporting virtual currency investigations and prosecutions; and creating cross-agency relationships for future cooperation and coordination on virtual currency issues, investigations, and prosecutions.

Working group	Participating agencies	Mission and goals	Ways in which group addressed virtual currencies
Terrorist Finance Working Group's New Payments Systems Ad Hoc Working Group	Department of State (lead, including the Bureaus of Economic and Business Affairs, Counterterrorism, and International Narcotics and Law Enforcement Affairs); Department of Defense; DOJ (Asset Forfeiture and Money Laundering Section; Criminal Division; DEA; FBI; National Security Division; and Office of Overseas Prosecutorial Development, Assistance and Training); FDIC; Federal Trade Commission; ICE-HSI; IRS-Criminal Investigation; Treasury (FinCEN, Office of Terrorism and Financial Intelligence, and Office of Technical Assistance), and USAID.	The larger working group's mission is to coordinate counter-terrorism-financing and anti-money-laundering training and technical assistance programs to countries deemed most vulnerable to terrorist financing. Within this context, the New Payments Ad Hoc Working Group's mission is two-fold: (1) to help ensure that foreign partners providing assistance and capacity building have a baseline understanding of new payment systems and the counter-terrorism-financing and anti-money-laundering risks and vulnerabilities that they may pose, and (2) to collaborate with other federal agencies and appropriate public and private sector entities to develop training and technical assistance programs in line with international standards set by groups such as FATF.	 The New Payments Ad Hoc Working Group, which formed in 2013 and meets every two to three months, has addressed the use of virtual currencies at several meetings. Topics have included: briefings on virtual currencies, how they operate, and risks; the set of common virtual currency vocabulary terms proposed in the FATF's discussion paper on virtual currencies; trainings that ad hoc working group participants plan to offer through 2015 on counter-terrorism-financing and antimoney-laundering risks associated with virtual currencies. workshops that the Department of State, USAID, and other ad hoc working group participants offered in 2013 and 2014 on new payment systems—including virtual currencies. the ways in which other interagency collaborative groups—such as the Egmont Group, which is composed of FinCEN and financial intelligence units from other countries.
Virtual Currencies Emerging Threats Working Group	DOJ (FBI lead and other DOJ components); FinCEN; ICE- HSI; SEC; Treasury; Secret Service; and other relevant federal partners.	To address the illicit use of virtual currencies.	This group leverages members' expertise to address new virtual currency trends, address potential implications for law enforcement and the U.S. intelligence community, and mitigate the cross-programmatic threats arising from illicit actors' use of virtual currency systems.

Source: GAO analysis of agency interviews and documents, as well as websites of interagency collaborative efforts.

^aDigital Economy Task Force, *The Digital Economy: Potential, Perils, and Promises* (Mar. 2014).

^bFDIC, the Federal Reserve, and NCUA told us that the FFIEC Taskforce on Supervision, and the Taskforce's Information Technology Subgroup, have also discussed virtual currencies.

^cThe FFIEC State Liaison Committee includes representatives from the Conference of State Bank Supervisors, the American Council of State Savings Supervisors, and the National Association of State Credit Union Supervisors. Other FFIEC BSA/Anti-Money-Laundering Working Group non-voting members include CFTC; FinCEN; IRS; SEC; Treasury's Office of Foreign Assets Control; and Treasury's Office of Terrorist Financing and Financial Crimes.

Appendix III: Comments from the Consumer Financial Protection Bureau

	- C - L
	CTOD Consumer Financial
	Protection Bureau
	1700 G Street, N.W., Washington, DC 20552
	May 6, 2014
	Lawrence Evans Jr. Director Financial Markets and Community Invectment
	U.S. Government Accountability Office
	441 G. Street NW
	Washington, DC 20548
	Dean Mr. Evens
	Dear Mr. Evans,
	Thank you for the opportunity to review and comment on the report: <i>Virtual Currencies</i> –
	Emerging Regulatory, Law Enforcement, and Consumer Protection Challenges, covering policy
	issues related to virtual currencies and the status of federal agency collaboration in this area.
	As you note in the report federal agencies have begun to collaborate on virtual currency issues
	through informal discussions and formal interagency working groups. In that regard, the
	Consumer Financial Protection Bureau (the "CFPB" or the "Bureau") has conferred on virtual
	currency issues with a number of domestic and international regulators, including the Federal
	Reserve Bank of San Francisco, the Federal Trade Commission, the National Credit Union
	Administration, the Office of the Comptroller of the Currency, New York State's Department of
	Financial Services, the European Banking Authority, and the U.S. Department of the Treasury. We
	have similarly met with academic and consumer group stakeholders, law firms, consultancies,
	industry associations, and industry participants.
	To date, formal interagency working groups addressing virtual currencies have focused primarily
	on Bank Secrecy Act concerns, anti-money-laundering controls, and the investigation of crimes in
	which virtual currencies may have been used. Accordingly, the CFPB's participation in these
	formal working groups has necessarily been limited, and our work has focused on more informal
	consultations with a consumer protection perspective.
	As noted in GAO's report, attention to potential consumer protection concerns in the virtual
	currency space has intensified in recent months. The Bureau believes that its own work on virtual
	currency and the work of other financial regulators will benefit from a collaborative response to
	these concerns, thus we concur with the report's recommendation that the Bureau identify
	interagency working groups addressing virtual currencies where CFPB's participation could
	enhance its own work in this area and could contribute valuable consumer protection expertise to
	these efforts. We look forward to increasing our involvement in formal working groups as they
	engage on specific issues relating to consumer protection.
	Sincerely,
	h. ((- h-)e- hy
	William Wade-Gery
	Acting Assistant Director Card and Payment Markets
	Card and Fayment Markets
C C	ionsumerimance.gov

Appendix IV: Comments from the National Credit Union Administration



Appendix V: GAO Contact and Staff Acknowledgments

GAO Contact	Lawrance L. Evans, Jr. (202) 512-8678 or evansl@gao.gov.
Staff Acknowledgments	In addition to the contact named above, Steve Westley (Assistant Director), Bethany Benitez, Chloe Brown, Anna Chung, Tonita Gillich, José R. Peña, and Robert Pollard made key contributions to this report. Also contributing to this report were Jennifer Schwartz, Jena Sinkfield, Ardith Spence, Andrew Stavisky, and Sarah Veale.

GAO's Mission	The Government Accountability Office, the audit, evaluation, and investigative arm of Congress, exists to support Congress in meeting its constitutional responsibilities and to help improve the performance and accountability of the federal government for the American people. GAO examines the use of public funds; evaluates federal programs and policies; and provides analyses, recommendations, and other assistance to help Congress make informed oversight, policy, and funding decisions. GAO's commitment to good government is reflected in its core values of accountability, integrity, and reliability.	
Obtaining Copies of GAO Reports and Testimony	The fastest and easiest way to obtain copies of GAO documents at no cost is through GAO's website (http://www.gao.gov). Each weekday afternoon, GAO posts on its website newly released reports, testimony, and correspondence. To have GAO e-mail you a list of newly posted products, go to http://www.gao.gov and select "E-mail Updates."	
Order by Phone	The price of each GAO publication reflects GAO's actual cost of production and distribution and depends on the number of pages in the publication and whether the publication is printed in color or black and white. Pricing and ordering information is posted on GAO's website, http://www.gao.gov/ordering.htm.	
	Place orders by calling (202) 512-6000, toll free (866) 801-7077, or TDD (202) 512-2537.	
	Orders may be paid for using American Express, Discover Card, MasterCard, Visa, check, or money order. Call for additional information.	
Connect with GAO	Connect with GAO on Facebook, Flickr, Twitter, and YouTube. Subscribe to our RSS Feeds or E-mail Updates. Listen to our Podcasts. Visit GAO on the web at www.gao.gov.	
To Report Fraud.	Contact:	
Waste, and Abuse in Federal Programs	Website: http://www.gao.gov/fraudnet/fraudnet.htm E-mail: fraudnet@gao.gov Automated answering system: (800) 424-5454 or (202) 512-7470	
Congressional Relations	Katherine Siggerud, Managing Director, siggerudk@gao.gov, (202) 512- 4400, U.S. Government Accountability Office, 441 G Street NW, Room 7125, Washington, DC 20548	
Public Affairs	Chuck Young, Managing Director, youngc1@gao.gov, (202) 512-4800 U.S. Government Accountability Office, 441 G Street NW, Room 7149 Washington, DC 20548	

EXHIBIT D





© 2021 Twitter, Inc About Help Terms Privacy Cookies Blog Advertise Businesses Media Developers TweetDeck

EXHIBIT E

Contents lists available at ScienceDirect





journal homepage: www.elsevier.com/locate/ipm

Taxonomy of centralization in public blockchain systems: A systematic literature review



^a Lero, Tierney Building, University of Limerick, Ireland

^b Horizon Globex Ireland DAC, T1-017, Tierney Building, Nexus Center, University of Limerick, Ireland

ARTICLE INFO

Keywords: Decentralized blockchain Centralization Classification Measurement Taxonomy Security

ABSTRACT

Bitcoin introduced delegation of control over a monetary system from a select few to all who participate in that system. This delegation is known as the decentralization of controlling power and is a powerful security mechanism for the ecosystem. After the introduction of Bitcoin, the field of cryptocurrency has seen widespread attention from industry and academia, so much so that the original novel contribution of Bitcoin, i.e., decentralization, may be overlooked, due to decentralizations' assumed fundamental existence for the functioning of such crypto-assets. However, recent studies have observed a trend of increased centralization in cryptocurrencies such as Bitcoin and Ethereum. As this increased centralization has an impact the security of the blockchain, it is crucial that it is measured, towards adequate control. This research derives an initial taxonomy of centralization present in decentralized blockchains through rigorous synthesis using a systematic literature review. This is followed by iterative refinement through expert interviews. We systematically analyzed 89 research papers published between 2009 and 2019. Our study contributes to the existing body of knowledge by highlighting the multiple definitions and measurements of centralization in the literature. We identify different aspects of centralization and propose an encompassing taxonomy of centralization concerns. This taxonomy is based on empirically observable and measurable characteristics. It consists of 13 aspects of centralization, classified over six architectural layers: Governance, Network, Consensus, Incentive, Operational, and Application. We also discuss how the implications of centralization can vary depending on the aspects studied. We believe that this review and taxonomy provides a comprehensive overview of centralization in decentralized blockchains involving various conceptualizations and measures.

1. Introduction

Since the introduction of Bitcoin in 2009, blockchain technology has seen a proliferation of scholarly articles investigating the potential and limitations of the technology (Androulaki et al., 2018; Beck, Avital, Rossi, & Thatcher, 2017; Beck, Müller-Bloch, & King, 2018; Davidson, De Filippi, & Potts, 2016; He, Yu, Zhang, & Bao, 2017; Mattila, 2016; Walport, 2016; Wüst & Gervais, 2018; Yli-Huumo, Ko, Choi, Park, & Smolander, 2016; Zheng, Xie, Dai, Chen, & Wang, 2017). Control over the system is a focal point in a significant proportion of these studies, as this either enhances or restricts the usability of blockchain in a wider information system context (Alzahrani & Bulusu, 2018; Azouvi, Maller, & Meiklejohn, 2018; Baliga, 2017; Beck et al., 2017; Beikverdi & Song, 2015; Cong, He, & Li, 2019; Gencer, Basu, Eyal, Van Renesse, & Sirer, 2018; Gervais, Karame, Capkun, & Capkun, 2014; Judmayer, Stifter,

* Corresponding author. *E-mail address:* 17053145@studentmail.ul.ie (A.R. Sai).

https://doi.org/10.1016/j.ipm.2021.102584

Received 25 June 2020; Received in revised form 4 February 2021; Accepted 7 March 2021 Available online 31 March 2021

0306-4573/© 2021 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/bv/4.0/).



Krombholz and Weippl, 2017; Kwon, Liu, Kim, Song, & Kim, 2019; Mattila, 2016; Sai, Buckley, & Le Gear, 2019a; Wang et al., 2018; Zhang, Xue, & Liu, 2019; Zheng et al., 2017). Indeed, removing central control from the monetary system while continuing to ensure security has been considered as a core novel contribution of Bitcoin (Bonneau et al., 2015). There are three main types of blockchain solutions based on the type of control mechanism used: public, private, and consortium (Zheng et al., 2017).

In public blockchains, such as Bitcoin and Ethereum, every participant in the network contributes to the control mechanism, agreeing on a single state of the data without the need for a trusted third party. All participants can read and write to this single state without any authorization (Guegan, 2017). This consensus is achieved under the assumption of delegation of power of control, and the assumption that the majority of the network participants remain honest i.e., non-malicious. This delegation of power of control is often referred to as decentralization.

Contrary to the decentralized nature of a public blockchain, private and consortium blockchains, such as Hyperledger, tend to impose constraints on participants by including trusted entities in the system (Androulaki et al., 2018; Xu et al., 2021). These constraints can also include limitations on read and write permissions of participants (Guegan, 2017). Based on the sensitivity of the information processed by the blockchain, practitioners may decide to adopt one of these controlling mechanisms (Meijer & Ubacht, 2018; Peck, 2017; Wüst & Gervais, 2018). As reported by Berdik, Otoum, Schmidt, Porter, and Jararweh (2020), the sheer number and complexity of various types of blockchain and their attributes can make it difficult to specifically address the benefits and shortcomings of blockchain as a service for applications within today's information systems. This decision is potentially problematic, e.g., in the case of a practitioner who decides to use a public blockchain for decentralizing the control. As reported by Sai et al. (2019a), decentralization in public blockchain is not a fundamental given by design, but a non-deterministic and probabilistic guarantee provided by clever integration of cryptography, distributed systems, and incentive engineering.

The removal of trusted entities from a distributed system makes a public blockchain attractive to numerous potential users in academia and industry (Mattila, 2016). Public blockchain-based cryptocurrencies have a market capitalization of over \$1.05 trillion (Sai, Buckley, & Le Gear, 2019b), making the platform a lucrative target for malicious actors. The majority of these blockchains use decentralization as a security mechanism. In a decentralized system, the malicious actor would need to compromise half of the consensus power before causing significant harm to the system (Karame, Androulaki, & Capkun, 2012). Because of this interplay between decentralization and security, it is highly desirable to have a high degree of decentralization in public blockchains. The security of a public blockchain has been thoroughly investigated in research (Bonneau et al., 2015; Halpin & Piekarska, 2017; Karame, 2016; Karame & Androulaki, 2016). For example, Bitcoin has been reported as secure, subject to its adherence to the honest majority assumption, with notable exceptions such as selfish mining attacks (Sapirshtein, Sompolinsky, & Zohar, 2016) where the attacker only needs to control over 26% of the network.

Even though the initial implementation of Bitcoin was able to circumvent the need for centralization in the system, new avenues of centralization are surfacing (Gervais et al., 2014). Numerous studies have reported various forms of centralization in Bitcoin and other decentralized cryptocurrency systems (Azouvi et al., 2018; Beikverdi & Song, 2015; Gencer et al., 2018; Gervais et al., 2014). These reports of a trend towards centralization have raised security concerns as the security guarantee of a public blockchain is inherently dependent on the honest majority assumption (Sai et al., 2019a). As reported by Gencer et al. (2018), Bitcoin's network is dominated by consortiums of participants working together in groups known as mining pools. We report that the top 4 mining pools constitute 50.36% of controlling power in Bitcoin with our analysis in Section 5.3. This power accumulation trend is also evident in Ethereum, where the top 4 mining pools aggregate 63% of controlling power (Section 5.3). Given that successful attacks on these networks are much more feasible when 50% of the network chose to carry out such an attack, this mining-centralization implies that only a relatively small number of participants (the heads of these mining pools) need to adopt a dishonest approach to threaten these Blockchains. This illustrates the importance of mining-centralization to the cryptocurrency-ecosystems and this research expands that focus to look at the wider implications of centralization in general in Blockchain systems.

Trusting the probabilistic security guarantees of a public blockchain has often been identified as a barrier to entry in the ecosystem (Iansiti & Lakhani, 2017). Security of Blockchain is considered central to the adoption (Akram, Malik, Singh, Anita, & Tanwar, 2020). The security of prominent blockchains seem to depend on the appropriate decentralization (Sai et al., 2019a). Thus deeply understanding the interplay between security and centralization is an important endeavor. The threats of centralization range well beyond security into adoption, and even crypto-economics (Conti, Kumar, Lal, & Ruj, 2018). The decentralized nature of bitcoin permits the uncensored execution of transactions in the payment system irrespective of political or geographical associations. Centralization may threaten the uncensored nature of the decentralized blockchain. Thus, it is crucial for the security and, consequently, the utility of public blockchain systems that they remain adequately decentralized.

Given the significance of decentralization, several studies have analyzed technical aspects (Beikverdi & Song, 2015; Gencer et al., 2018; Gervais et al., 2014) as well as social constructs of decentralization (Azouvi et al., 2018). By far, the most commonly measured aspects of centralization is the consensus power concentration (Azouvi et al., 2018; Beikverdi & Song, 2015; Gencer et al., 2018; Gervais et al., 2014; Kwon et al., 2019). In a Proof-of-Work based blockchain solution, the individual participants' consensus power is defined by their computational power in proportion to the total computational power of the network. However, this measurement mechanism is only useful in determining the present state of the computational power portions of the network. It fails to capture the multitude of factors that may constitute the overall centralization of the system, such as system governance (Beck et al., 2018), wealth concentration (Chohan, 2019), and geographic distribution of participants (Gencer et al., 2018).

To better understand the semantics of decentralization in blockchain, we intend to measure it on all building blocks of the public blockchain. As reported by Wang, Vergne, and Hsieh (2017), the governance structure of the blockchain can have a profound impact on the operations of a public blockchain but is often overlooked as a potential source of centralization. The issues caused by centralization of governance include the long-discussed issue of block size in Bitcoin (Caffyn, 2015) and specific instances



Fig. 1. Methodology.

of unilateral decision making regarding forks in Ethereum (Wirdum, 2016). Bitcoin and other similar cryptocurrencies rely on improvement protocols to dictate the changes in the core system. According to the empirical analysis of Azouvi et al. (2018), the authors report that the vast majority of the improvement proposal in Ethereum are authored by a single user, Vitalik Buterin, the founder of Ethereum. They also report a similar trend for Bitcoin, where a handful of users contribute to the improvement protocol. This observation has been cited as a potential source of centralization in the governance of these cryptocurrencies (Gervais et al., 2014) and may serve to stifle innovation. Alternatively, they may serve to promote high-quality changes from a pool of proposers that know the Bitcoin/Ethereum ecosystems intimately. But the first step in studying this phenomenon is to acknowledge the potential for centralization in improvement protocol and to formulate a measurement to assess it.

In this study, we identify other forms of centralization, including end-user application centralization. According to Böhme, Christin, Edelman, and Moore (2015), 95% of all Bitcoin trades are processed by seven centralized organizations known as exchanges. Thus, the presence of centralized exchanges may be a contributing factor to the wealth concentration on the blockchain. Based on the analysis of Srinivasan (2017), Bitcoin and Ethereum have a wealth inequality greater than the worst real-world economy. This wealth centralization has been linked to severe security threats (Section 4) (see Fig. 1).

Consequently, we reason that we need a vocabulary to discuss and measure centralization in a more holistic manner. To allow for such modular measurement of centralization, we review the generic architecture (Zhang et al., 2019) of blockchain and use it to identify potential avenues of centralization, via a literature review of the field. Focusing on the generic architecture enables us to capture centralization-causing factors that are not implementation-specific, i.e., the same model may be used for both Bitcoin and Ethereum. We also use the generic architecture to partition different centralization concerns into architectural categories such as consensus, network, and application. This abstraction allows us to organize and observe centralization holistically. Thus, in this work, we present the first in-depth analysis of centralization in blockchains to assess the following questions:

RQ1: What are the different aspects of centralization in public blockchains?

RQ2: How can centralization be adequately measured in a decentralized blockchain instance?

To study decentralization in blockchain, we coded and analyzed the content of relevant blockchain literature (see Fig. 1). We chose ten years subsequent to the publication of the original Bitcoin white paper (Nakamoto, 2008). The survey process was primarily driven by the guidelines provided by Kitchenham (2004). In adherence to the guidelines, we conducted a five-step systematic literature review consisting of *Search, Selection, Quality Assessment, Data Extraction, and Analysis.* This systematic literature review produced the final article pool of 89 articles. These final articles, partitioned by architectural components, form the basis of the taxonomy proposed in this review.

Following the development of the taxonomy, we interviewed industrial and academic experts in the blockchain domain to establish the completeness of the taxonomy and to assess any redundant or less relevant components of the taxonomy. This consisted of ten expert interviews: four academic researchers and six industry experts. It resulted in an iterative refinement of the taxonomy.

We believe that the taxonomy presented in this article can assist in better understanding the socio-technical nature of blockchainbased information systems. The taxonomy focuses on reporting the security and performance implications of centralization systematically to reduce the complexities involved in understanding the benefits and shortcomings of blockchain for information systems, as reported by Berdik et al. (2020). We also highlight the issues associated with managing a decentralized blockchain system in the form of governance and protocol improvements. The paper makes the following contributions:

- We systematically review the existing literature to document the different aspects of centralization in public blockchains (Section 3).
- We outline the different techniques employed in the literature to measure centralization (Section 4).
- We manifest the findings of our review in a conceptual taxonomy that encompasses both categorization and measurement of different aspects of centralization in public blockchains (Section 4).
- We illustrate the relevance and utility of this taxonomy by presenting the centralization state of the two most prominent blockchain instances: Bitcoin and Ethereum, based on this taxonomy (Section 5). We also discuss how the adverse impact of centralization varies depending on aspects (Section 6).
- We identify research gaps specifically with regards to the lack of non-Bitcoin-specific centralization investigations. We also report on the lack of objective metrics for some centralization causing factors.

2. Background

The term blockchain is often used as a generic descriptor for the broader field of Distributed Ledger Technologies (Great Britain. Government Office for Science, 2016). Distributed ledger technology refers to the distributed computing networks that record, share, and synchronize data across many participants. More specifically, Blockchain is a type of data structure used to record data on these distributed computing networks. It is a chronologically linked list of data packets received by the participants within a predefined time period. These blocks are connected in a chronological order to form a chain of blocks. The link between these blocks is secured by the use of a computationally hard cryptographic hash function (Nakamoto, 2008). As the chain of blocks grows, the difficulty involved in recalculating the hash value also grows to make any alteration to past data expensive. This growth in difficulty leads to a deterministic guarantee of data immutability.

The participants of the blockchain-based network have to reach consensus on a single state of this append-only structure. Blockchain-based systems utilize a peer-to-peer distributed system with a clever incentive mechanism (Baliga, 2017) to accomplish this consistency of data in an unconstrained distributed environment. Proof-of-work (PoW) and Proof-of-stake (PoS) are two prominent examples of consensus mechanisms used in blockchain-based systems. In PoW, the participants are expected to perform computationally expensive operations to solve a puzzle. The first participant to solve and propagate the solution to a majority of the network is rewarded. PoW is often criticized for the extensive use of electricity (O'Dwyer & Malone, 2014). This issue of electricity usage is addressed in PoS, where the reward distribution is based on the monetary assets of the participants (Nguyen et al., 2019). Other notable consensus algorithms include Proof-of-Authority, Proof of Elapsed Time, and Delegated Proof-of-Stake; we refer the reader to Mingxiao, Xiaofeng, Zhe, Xiangwei, and Qijun (2017) for an in-depth review of consensus algorithms.

As discussed earlier, based on the type of consensus mechanism deployed and the constraints imposed, we can segment blockchain-based systems in three broad categories: Public, Private, and Consortium. In private and consortium-based blockchain systems, the participation in consensus is limited to users approved by a trusted authority. However, in Public blockchain systems, the participation in consensus is open to any individual with appropriate computing and networking capabilities. This unconstrained access to controlling power for all participants in the network is referred to as decentralization. Bitcoin and other public blockchains establish consensus on the blockchain through a decentralized, pseudonymous protocol. This protocol can be considered a core innovation and possibly the most crucial ingredient to the success of public blockchains (Bonneau et al., 2015).

The possibility of decentralized control over a computing network without oversight has resulted in many novel applications of the blockchain technology in information systems to improve efficiency or increase the security of the operations (Hileman & Rauchs, 2017). The blockchain technology provides a general-purpose approach to managing information in a non-trusted computing environment enabling a plethora of information systems use cases such as auditing of big data (Li, Wu, Jiang, & Srikanthan, 2020), secure information management (Putz, Dietz, Empl, & Pernul, 2021), countering fake news (Chen, Srivastava, Parizi, Aloqaily, & Ridhawi, 2020), cloud computing (Baniata, Anaqreh, & Kertesz, 2021), health data management (Hardin & Kotz, 2021), copyright management (Jing, Liu, & Sugumaran, 2021), IoT management (Chen et al., 2020; Zhao, Chen, Liu, Baker, & Zhang, 2020) and assisting autonomous vehicles in reaching consensus on events (Esposito, Ficco, & Gupta, 2021; Khalid et al., 2021; Oham, Michelin, Jurdak, Kanhere, & Jha, 2021).

2.1. Decentralization and public blockchain

Decentralization is an essential property of public Blockchain systems where participants can read, write data, and contribute to consensus without authorization (Davidson et al., 2016). In this subsection, we review the existing discussion around decentralization in the blockchain.

Consensus on the state of data in a public blockchain is attained by the acceptance of a valid block by the network in a time interval determined by a stochastic process to maintain a predefined expected time interval. To deter malicious participants from accepting fraudulent blocks, the majority of the control must be decentralized. This decentralization of control ensures that the blockchain is secure from malicious participants as long as the majority of the network remains honest. This interplay of security and decentralization makes it fundamental that the system remains decentralized.

A survey paper by He et al. (2017) identifies decentralization, among other features, as a prominent reason to adopt blockchain technology for business applications. This view is supported by numerous studies which demonstrate the application of decentralized Blockchains to the liberalization of financial asset management (Guo & Liang, 2016), the Internet of Things (Panarello, Tapas, Merlino, Longo, & Puliafito, 2018; Zhu, Loke, Trujillo-Rasua, Jiang, & Xiang, 2019), healthcare (Dwivedi, Srivastava, Dhar, & Singh, 2019) and smart cities (Xie et al., 2019). The extent of literature surveyed by these review articles demonstrates the significance of decentralization in blockchain applications.

As decentralization is core to the secure functioning of public blockchains, it may be taken as a fundamental given. This assumed association between decentralization and public blockchains may be a vulnerability that malicious actors attack. Security research on the blockchain has focused on the assumption of an honest majority. A survey paper by Li, Jiang, Chen, Luo, and Wen (2017) identifies the centralization of consensus power as a significant security threat to that network. Centralization of consensus power is intrinsic to attacks on the public blockchain, such as the 51% attack (Bradbury, 2013) and Selfish Mining (Sapirshtein et al., 2016).

In the 51% attack, the attacker is assumed to have gained control of more than half of the consensus power, which can then be used to enter fraudulent transactions in the blockchain. Unlike the 51% attack, in selfish mining, the attacker only needs to control 26% consensus power to cause harm to the network (Sai et al., 2019a). More detail on the security of blockchains is provided in Zhang et al. (2019).



Fig. 2. Architecture of public blockchain.

Studying blockchain as from solely a technical perspective may be misleading due to the inherent socio-technical nature of the blockchain (De Domenico & Baronchelli, 2019). As the study of centralization in public blockchain is still fragmented, current conceptual models, such as security and privacy models, do not provide adequate insights. To overcome this limitation, we devise a novel centralization taxonomy focusing on the different architectural layers of blockchain to categorize centralization concerns. We employ a two-step research approach, first conducting a systematic literature review to construct a taxonomy of centralization, and refine this further through expert interviews.

2.2. Architecture of public blockchains

The first public blockchain, Bitcoin, incorporated the blockchain data structure and consensus mechanism in-depth, but omitted any formalization of the networking structure (Nakamoto, 2008). Since the introduction of Bitcoin, numerous attempts have been made to describe the structure of public blockchains more formally.

Some of these attempts have been aspect-specific with a microscopic focus on one or a few components of the blockchain. For example, Garay, Kiayias, and Leonardos (2015) describe the architecture of blockchain in terms of consensus mechanisms and participants of the network. Another notable description of blockchain architecture is given by Gervais et al. (2016), who focus on security and scalability by describing consensus and a peer-to-peer network.

Since the aim of our review is to analyze public blockchains more holistically to capture the factors causing centralization, we adhere to a more generic description of blockchain used by Zhang et al. (2019) and Zhu et al. (2019). In this generic description, the authors propose a layered architecture of blockchain. As a blockchain is a peer-to-peer distributed network, it is intuitive that blockchain systems will share many similarities with a generic, distributed computing architecture, such as the traditional OSI layered model of a network (Briscoe, 2000).

This layered architecture, illustrated in Fig. 2, describes how the data is stored (Data Layer) and shared (Network Layer) between different participants of the network. Once the data is shared with peers in the network, the network is tasked with agreeing a single view of the data (Consensus Layer). Public blockchains attain consensus in the network by incentivizing non-malicious participants using an incentive mechanism (Incentive Layer). Incentive and consensus operations are performed by the execution of computational scripts (Contract Layer). The computational capabilities of a blockchain are not just limited to these two operations; many different applications can be built on top of the blockchain such as cryptocurrencies and decentralized applications (DAPPS) (Application Layer) (Antonopoulos & Wood, 2018).

In the following subsection, we describe these layers in-depth:

2.2.1. Data layer

The data layer contains the definition of the data structure used by the system, including how transactions are stored, thus encompassing the transactions component proposed by Bonneau et al. (2015). Other data layer components include the cryptographic primitives employed on the blockchain. The network participants must adhere to the data layer specifications to participate in the network, i.e., use the same protocol to communicate. Application layer blockchain clients implement these specifications for the end-user.

2.2.2. Network layer

The network layer specifies the behavior of the nodes (network participants) in a distributed network. This behavior includes the network connection establishment and intercommunication mechanism. The network layer is responsible for the discovery of other nodes on the network and for efficient communication among nodes. The network layer serves as the information dissemination mechanism of the system. This network layer is identical to the network subsystem in the structure proposed by Judmayer, Stifter et al. (2017).

A.R. Sai et al.

2.2.3. Consensus layer

Once the participating nodes are connected in a predefined topology, the next step is to generate blocks to contribute to the growing ledger. As all the participating nodes are tasked with the creation of the next block, it is crucial that the network can agree on a single state of the ledger. The aim of the blockchain network is to deterministically agree on a single state of the data. The consensus layer assures that the network reaches a consensus with a certain degree of assurance.

2.2.4. Incentive layer

This deterministic assurance in prominent consensus algorithms such as Proof-of-work, and Proof-of-Stake, is based on the assumption of an honest majority, i.e., the network has greater than 50% non-malicious participants. Blockchain systems use incentive engineering to ensure that the majority of the network is honest (Sai et al., 2019a). This incentive is often in the form of a block reward which is assigned to the node that successfully adds a new block to the blockchain. The incentive layer describes the mechanism used for issuance of reward and the distribution of reward. This layer acts as an interface between the user-facing layers and the technical implementation layers.

2.2.5. Contract layer

To process transactions in the network, Bitcoin uses a scripting language called *script* (Antonopoulos, 2017). This scripting language is significantly limited in terms of functionality as it lacks Turing completeness (Buterin et al., 2013). One example of this is the lack of loops in *Script*. Despite the lack of such functionality, the scripting language serves as the building block of Bitcoin cryptocurrency, enabling complex financial transaction processing.

The limitations on the scripting language of Bitcoin served as a motivation for Ethereum's developers (Wood et al., 2014). Ethereum implements a Turing complete computing engine on top of a distributed blockchain. Applications on top of the blockchain exploit this programmable nature of blockchain. The contract layer also acts as the interface between information systems and the blockchain (Beck et al., 2017).

2.2.6. Application layer

Public blockchains provide a mechanism that can be used to interact with and run user-defined code on the computing engine provided by the contract layer. JSON HTTP API is an example of one such public API provided by Ethereum (Lee, 2019). These public APIs serve as an interface between different Broker–Dealer services such as Wallets and Exchanges and the blockchain. These services are primarily used by end-users to interact with the blockchain (Chu, 2018).

2.3. Taxonomy development methodology

Classification of logically related objects is a fundamental problem in many disciplines. Taxonomies are considered an important tool to logically classify objects to better understand complex domains (Guerra García, Espinosa Torre, & García Gómez, 2008). The concept of taxonomy was initially proposed by Carolus Linnaeus (Lindley, 1836) to group organisms in Biology. Since then, taxonomies have been used in different knowledge domains such as social science (Bailey, 1994), computer science (Buckley & Exton, 2003), and information systems (Oberländer, Lösser, & Rau, 2019).

Due to the emerging nature of blockchain technologies, the state taxonomies in the field is preliminary. The most prominent of taxonomies in blockchain have been architecture (Xu et al., 2017) and security-specific (Zheng, Xie, Dai, Chen, & Wang, 2018). However, these taxonomies often treat blockchain as a single-dimensional computer science artifact whereas, as discussed earlier, the secure functioning of the blockchain-based assets is the result of the socio-technical nature of information systems. In the following subsection, we describe the state of the art in information system specific taxonomy formation and how our methodology aligns with this existing research.

Information system researchers have recognized the importance of taxonomies in knowledge organization. Specifically, Nickerson, Varshney, and Muntermann (2013) observed that despite the significance of taxonomies in information systems, the taxonomy development process remained largely ad hoc. To address this research gap, Nickerson et al. (2013) proposed a taxonomy development method specific to information systems. The development method proposed by Nickerson et al. (2013) has served as the guidelines followed by many information systems taxonomies (Oberländer et al., 2019). In this article, we follow the seven-step method proposed by Nickerson et al. (2013). We have illustrated the seven-step method in Fig. 3.

The first step of taxonomy construction is the determination of meta-characteristics. According to Nickerson et al. (2013), metacharacteristic is the most comprehensive characteristic that will serve as the basis for the choice of characteristics in the taxonomy. For example, if the researcher wants to classify a computer platform based on performance, the meta-characteristics are the hardware and software characteristics such as processing power, storage, and software optimization. Nickerson et al. (2013) also highlight the evolving nature of the meta-characteristic as many characteristics only become apparent through the taxonomy construction. In our taxonomy, we ground our meta-characteristic in the generic architecture described in Section 2.2.2.

After establishing the meta-characteristic, Nickerson et al. (2013) suggest the determination of ending conditions. As the taxonomy construction process is iterative in nature, it is crucial to establish end conditions. In the Nickerson et al. (2013) model, there are two types of end conditions: objective and subjective. For our taxonomy construction, we establish one objective and one subjective end condition. The objective end condition is the exhaustive examination and classification of all survey objects (aspects of centralization). The subjective end condition for our taxonomy is the determination of comprehensiveness through expert interviews.



Fig. 3. Taxonomy construction methodology (Nickerson et al., 2013).



Fig. 4. Methodology.

Once we have established the meta-characteristic and ending conditions, Nickerson et al. (2013) propose two taxonomy construction approaches. In the first approach, conceptual-to-empirical, the researcher attempts to conceptualize the taxonomy dimensions without an exhaustive analysis of objects. The second approach, empirical-to-conceptual, relies on a review of the objects before the taxonomy constructions. This is often done in the form of a literature review. In our taxonomy construction, we follow the empirical-to-conceptual approach.

Within the empirical-to-conceptual model of taxonomy construction in the information system, there are three distinctive steps. In the first step, we identify a subset of objects. In our taxonomy, we identify new objects through a systematic literature review: Phase 1 and Phase 2 in Fig. 4. The second step is to identify common characteristics and group objects. We perform data extraction and analysis of the objects shortlisted through the systematic literature review to construct these grouped objects. This is done in Phase 3 of our methodology, as illustrated in Fig. 4. In the last step of the empirical-to-conceptual model, we group characteristics into dimensions to create or revise the taxonomy. For our taxonomy construction, Phase 4 attempts to construct a conceptual taxonomy that is refined through iterative cycles and structured via the architectural-layers lens.

In the following section, we describe our research methodology in detail.

3. Methodology

In this section, we describe the research methodology employed for our systematic literature review (SLR) of blockchain through which we sought to provide a more cohesive overview of centralization in public blockchains. We follow the SLR guidelines proposed by Kitchenham (2004) to identify the factors associated with centralization. We then use a classification scheme based on the generic architecture presented in Section 2.1 to map the identified factors and associated measurement techniques. This mapping is loosely based on the approach proposed by Petersen, Feldt, Mujtaba, and Mattsson (2008). The mapping of obtained data to the generic architecture produces an initial taxonomy, which we then refined by conducting ten expert interviews to improve the taxonomy. This process is graphically illustrated in Fig. 4.

3.1. Systematic literature review

The systematic literature review guidelines suggested by Kitchenham (2004) span four phases:

- In the first phase, we define the two primary research questions for the review and produce relevant keywords for the subsequent search.
- In phase two, we systematically extract relevant articles from leading research repositories. We filter the resultant articles through a manual review of titles and abstracts.
- In phase three, the shortlisted articles are then used for data extraction, which is driven by an extraction protocol.
- In phase four, we perform the mapping of the data extracted from phase three to the generic architecture presented in Section 2.1, leading towards an initial taxonomy of centralization in public blockchains.

Fig. 5 illustrates the literature review employed in the study in more detail.

3.1.1. Phase 1: Research questions and query formation

The primary aim of our review is to provide richer insight into the different types of centralization present in public blockchain. We also identify techniques used to measure these aspects of centralization quantifiably. This will inform the development of our initial centralization taxonomy of public blockchains. We define the research questions of our study as follows:

- RQ1: What are the different aspects of centralization in public blockchains?
- RQ2: What techniques are employed to measure these centralization aspects?

Regarding RQ1, if a paper presented a novel centralization-causing factor, it is mapped to the architecture. If our generic architecture cannot accommodate the identified factor, we modify the architecture. This process is repeated for every novel factor identified. If a paper identified a factor already present in our taxonomy, we retain the reference to the article, using number-of-articles to define a proxy for the significance of that particular factor.

For every identified factor, we also recorded any measurement technique used to quantify the factor. If multiple papers employ different measurement techniques for a single factor, we retained all measurement techniques.

These research questions form the basis of article identification and selection, as they define the relevance of a particular article to our review. As we aim to capture factors from different socio-technical aspects of the blockchain, we conducted an exhaustive search on the following leading digital repositories: **Google Scholar**, **ACM Digital Library**, **IEEE Digital Library**, **ISI Web of Science**, **Science Direct**, **Scopus and Springer Link**. These repositories provided us with access to a wealth of articles, including gray literature.

Having identified the search repositories, we formed the search query. We adopted a systematic approach to keyword generation to form the search query:

1. **Initial set of keywords**: We formulated an initial set of keywords for the search consisting of "Blockchain" and "Centralization" with the following synonyms and alternate words:

Blockchain: bitcoin, ethereum, blockchain, cryptocurrencies, cryptocurrency, distributed ledger, DLT, Merkel tree, smart contract platform, tokenized asset.

Centralization: centralization, centralism, consolidation, decentralisation, decentralization, devolution, dominating, domination, managed, monopolization, monopoly, singular, unipolar.

2. Text Corpus Creation: Complementary to the initial set of keywords, we also reviewed existing studies on centralization to extract more relevant keywords. We selected the two most cited relevant studies from Google Scholar (Gencer et al., 2018; Gervais et al., 2014). We performed forward, and backward snowballing on these two articles and generated a list of the most used keywords from this set. We selected the top 5 keywords from this set. This leads to the inclusion of "digital currency" and "oligopoly" to our initial set of keywords.

The resultant queries from query formation step are present in Appendix A.

3.1.2. Phase 2: Article search and selection

Given that decentralization is fundamental to a public blockchain, we expect that the search will return a high number of articles. We implement a filtering process to limit the search to relevant articles. We restrict our search to articles published in English after the introduction of Bitcoin in 2009. We refrain from treating citations as a proxy for quality to filter articles, as it has been questioned in the past (Galster, Weyns, Tofan, Michalik, & Avgeriou, 2013).

After the execution of a search query, Google Scholar returned the highest number of articles with 4380 results. However, due to the restrictions imposed by Google Scholar, we can only retrieve the first 1000 most relevant articles (Razzaq, Wasala, Exton, & Buckley, 2018). After applying the language and publication date constraints, we retrieved 982 articles from Google Scholar. We also



Fig. 5. Overview of systematic literature review.

Table 1			
Quality assignment matrix.			
Attribute	No	Yes	
1. Centralization factor identified	0.0	1.0	
2. Factor measurement technique proposed	0.0	1.0	

retrieved additional 2737 articles from all other sources resulting in a total of 3728 articles. All of these articles were cross-checked to identify duplicate entries. After the removal of duplicate articles, the final set contained 3572 articles.¹

Due to the high number of articles, we first analyzed the title and abstract to establish relevance. This was based on explicit inclusion criteria. The shortlisted, relevant articles were then scanned further to assign a quality score. These shortlisted articles

¹ A list of selected articles is available at www.github.com/ashishrsai/centralization.

were assessed for quality with regards to our research questions. To ensure that the assessment process is reliable, we followed the inclusion criteria for titling, abstraction, and full-text screening. This process obeyed the following inclusion criteria:

- 1. The paper's title mentions centralization, or any of the synonyms mentioned above, or is potentially relevant to the study of centralization.
- 2. The abstract is relevant to the identification or measurement of centralization-causing factors.

During the review of the title, we tried to avoid eliminating articles that might have some relevance to the topic of centralization. This relevance was evaluated by the review of the abstract. We excluded articles that did not pass both criteria.

The first author conducted this analysis. To test for reliability, we performed cross-validation by following Fleiss and Cohen (1973). We specifically use the guidelines proposed by Sim and Wright (2005) for the calculation of sample size. We select 89 articles with a confidence level of **95%** and a margin of error of **10%**. This sampling contained an equal number of accepted and rejected articles by the first author to eliminate the possibility of only sampling accepted or rejected articles. The second author was then tasked with the evaluation of these 89 articles based on the guidelines provided above. Results from the cross-validation suggest that both the reviewers were in almost perfect agreement over the acceptance and rejection of the articles with the Cohen's Kappa.² exceeding 0.8 (Landis & Koch, 1977).

Using this process, we retrieved 212 relevant articles for our study. Subsequently, we performed quality assessment of these articles by conducting full-text review. We assigned a quality score between 0 to 2 based on the relevance of the article to our research question. Table 1 outlines the assignment matrix employed for quality assessment.

We reviewed each article on two attributes - (1) factor identification and (2) measurement techniques used. If an article identifies a novel centralization-causing factor, we assign a score of 1.0 for Attribute 1. Articles that do not identify a novel centralization or refer to already identified factors are assigned a score of 0.0 for Attribute 1.³

We follow a similar quality assignment scheme for Attribute 2, where we assign a score of 1 for the identification of a novel measurement technique. Articles not proposing or using any existing measurement techniques are assigned a score of 0.0 for attribute 2.

To ensure that the quality assignment process is reliable, we again perform a similar reliability test but with a smaller data set of 9 articles. We observe that both the reviewers (first and fourth authors) agree on eight score assignments with one score difference for the ninth article. This disagreement is resolved when the article is reviewed by the third author.

This filtering process resulted in a set of 89 articles. These articles are used in the third phase of our study: Data Extraction.

3.1.3. Phase 3: Data extraction

Having identified relevant studies, the next step is to extract relevant data from them. For this purpose, we design a protocol to analyze the articles towards the development of an initial taxonomy of centralization. In this context, we focused on the factors identified and measurement techniques proposed or used. We reviewed all of the shortlisted articles to create a list of factors and associated measurement techniques. The extracted data from this step serves as a building block for our taxonomy.

3.1.4. Phase 4: Development of initial taxonomy

As we aim to structure the findings of the review in an initial taxonomy, we use the data extracted in Phase 3 and map it to appropriate layers in the generic blockchain architecture. We repeat this process for all identified factors; if a factor cannot reasonably be mapped to the existing layers, we typically refine the architecture by including an additional layer. This iterative refinement results in a blockchain architecture specific to the study of centralization. Results from this mapping analysis are illustrated in Fig. 6. Out of all shortlisted articles, 63 considered the consensus layer as prone to centralization, the highest reported count for any layer in our survey: This is represented in Fig. 6 by the size of the bubble, but we discuss these results in more depth in Section 4.

To further validate the initial taxonomy and refined architecture, we conducted interviews with industry and academic experts.

3.2. Interview with experts

The initial taxonomy, as referred to in Section 3.1, is based on the review of existing literature. To raise confidence that the initial taxonomy proposed by the study provides relevant coverage and is accurate, we further refine and validate it by interviewing experts.

To identify experts in the blockchain field, we relied on the epicenters of the bibliographic map generated by Ramona, Cristina, Raluca, et al. (2019). We approached 112 active researchers based on their prominence determined by their location on the bibliographic map. Out of 112 researchers approached for the study, we received a response from 10 and subsequently interviewed them. We interviewed four academic experts (I_1 to I_4) and six experts from industry (I_5 to I_{10}). Interviews were typically one hour in duration and involved open-ended questions⁴ These open-ended questions were designed to:

² Cohen's kappa is a statistic measure of the agreement between two raters based on the classification of items in mutually exclusive categories.

 $^{^{3}}$ Although we do record the paper, as this helps us identify the significance of that centralization aspect in the literature.

⁴ The interview script is available at www.github.com/ashishrsai/centralization.



Fig. 6. Article titling and abstraction process.

- 1. Extract the view of the expert on centralization and the significance of it in their respective field, i.e., security, economics, information systems, and industrial application.
- 2. If needed, refine the taxonomy and/or the architecture.
- 3. Validate the generic architecture of the blockchain used in this study (Section 2).
- 4. Assess the accuracy of the initial centralization taxonomy.

The transcripts of these interviews are available in anonymized form⁵ These transcripts are color-coded based on the relevance of the conversation to factor identification and measurement.⁶

3.3. Illustrative walk-through of the four phases

Thus far in this Section, we have described the four phases used for taxonomy development and refinement. In this subsection, we present an explanatory walk-through of 2 articles through these phases. For this illustration, we select the following two articles: Gencer et al. (2018) and Peck (2017).

In phase one, we formulate the search query through an initial set of keywords and snowballing on seminal work in centralization. Gencer et al. (2018) is one of the two articles used for snowballing and keyword formulation due to the high citation count. After constructing the query, we move to phase 2: executing the query and shortlisting the appropriate articles.

During title screening in phase 2, after performing the search across the academic databases, Gencer et al. (2018) is included for abstract screening as the title points to the state of decentralization. Likewise, the second illustrative article, Peck (2017), is also shortlisted as the title refers to the difference between different blockchain forms.

In the abstract screening step, the article Gencer et al. (2018) is considered relevant because the abstract makes direct reference to the state of decentralization. The second article, Peck (2017), is also shortlisted as the abstract points to the risk of limiting controlling power to a select few participants.

Both the articles are now evaluated for quality by conducting a full-text review. In the first article, Gencer et al. (2018) describe the fundamentals of centralization on the network layer in blockchain, identifying novel centralization causing factors, and suggesting novel measurement techniques. Following the quality assessment matrix in Table 1, we assign a quality score of 2 to Gencer et al. (2018).

The full-text analysis of Peck (2017) reveals that the article does not identify or measure any centralization causing factor, therefore obtaining a quality score of 0. This article is henceforth excluded from taxonomy formulation.

Having identified the relevant articles, we perform data extraction in Phase 3. In data extraction, we first extract all the factors identified by Gencer et al. (2018): Consensus Power Distribution (Section 4.3.), Geographic Distribution (Section 4.2.2.), Bandwidth Concentration (Section 4.2.3) and Routing Centralization (Section 4.2.4). After identifying the centralization causing factors, we extract the measurement techniques used or suggested in the article: The authors proposed using a percentage-based value for Consensus Power Distribution, a latency-based measurement for identifying the geographic location for participating nodes, clustering for bandwidth concentration, and using AS (autonomous systems) coverage as a metric for routing centralization (Section 4.2.4).

After extracting the centralization causing factors and measurement techniques, we move to phase 4, constructing the initial taxonomy. In our representative example article, Gencer et al. (2018) have identified four centralization-causing factors. In this step, we venture to map these four factors to the generic blockchain architecture described in Section 2. The first centralization causing factor, consensus power distribution, is mapped to the consensus layer as this factor is within the layer's scope. The remaining three centralizations causing factors are all related to the networking aspects of the blockchain network. The geographic distribution

⁵ The transcripts can be obtained from www.github.com/ashishrsai/centralization.

⁶ More details on the coding scheme provided in Appendix C.

Table 2

Taxonomy of centralization in public blockchains.

Layer	Centralization factor	Measurement techniques
Application layer	Wallet concentration	Not found
	Exchange concentration	Centrality & Percentage value
	Reference client concentration	Satoshi index
Operational layer	Storage constraint	Ratio of growth
	Specialized equipment concentration	Not found
Incentive layer	Wealth concentration	Gini coefficient & Percentage value
Consensus layer	Consensus power distribution	Percentage value & Gini coefficient & Theil index & Centralization factor
Network layer	Node discovery protocol control	Not found
	Geographic distribution	Gini coefficient & Latency
	Bandwidth concentration	Clustering of provisioned bandwidth
	Routing centralization	AS-Level coverage
Governance layer	Owner control	Fractional measurement
	Improvement protocol	Centrality metrics

results from the open peer-to-peer network topology, whereas the routing, and bandwidth centralization target the network layer's information dissemination aspect.

After mapping these factors and their measurement techniques to the architecture, we construct an intermediate form of the centralization taxonomy. This intermediate form is iteratively refined as we process more articles through the four-phased approach. The resultant taxonomy is described in-depth in the following Section.

4. Taxonomy of centralization of public blockchain

In this Section, we map the results of the systematic review, and the interviews with experts, to the initial taxonomy of centralization outlined in Table 2.

As discussed in Section 3.1, this generic architecture is refined to reflect the centralization-related aspects of the blockchain better. To this end, we refined the generic architecture by removing the Data and Contract layers as none of the surveyed articles suggested any centralization aspects for either of these layers. As can be seen from Table 2, on average two centralization factors were identified for each resultant layer. As is also presented in the table, there are some factors for which there are no proposed measurement techniques (for example 'Wallet Concentration'). We also note that the existing generic architecture was unable to capture governance-related aspects of the blockchain system. For example, as blockchain systems evolve, it is crucial to have a mechanism to handle improvements such as security patches of the system. We account for the governance-related aspects of centralization by including a Governance Layer.

Another set of centralization causing issues that the generic architecture does not capture are associated with the operation of a node on the network. These issues include the computational requirements for participation, such as proprietary hardware and storage. In accordance with the recommendation of interviewee I_{10} , we include an Operational layer to represent the centralization associated with operating as a node on the blockchain.

Table 3 considers the factors identified in Table 2 from the perspective of 'prevalence-of-occurrence' in the literature and the interviews, where prevalence is considered as a proxy for whether the factor is "established" or not. The literature references in the table identify that particular factor as a potential source of centralization.⁷ The interviewer identifiers are used to indicate explicit recognition of the factor as a contributor to centralization in the associated interview. Interestingly, based on the data presented in this table, most of the factors can be considered well established, with the possible exception of Bandwidth Concentration and Routing Centralization. Even though Node Discovery Protocol Control was only referred to by one academic article, the majority of interviewees perceived it as a relevant factor.

Based on our taxonomy, we define centralization of public Blockchains as *the process by which one or more architectural dimensions* (*aspects*) of the Blockchain are restrictive to the majority of participants by direct or indirect economic, social, or technical constraints. We report a total of 13 aspects spread over six architectural layers. The governance layer aims to capture the social constructs of building and maintaining a public blockchain, specifically reporting on the incentives to build (Owner Control) and maintain a public blockchain (Improvement Protocol). The governance layer feeds into the economic aspects of the Blockchain in forms of incentives, this is captured by the Incentive layer, where we review the wealth inequality (Wealth Concentration). This inequality is in part caused by the technical constraints of participation ranging from Networking aspects such as bandwidth and routing requirements to operational requirements restrict participation in the consensus, which is observable in the consensus layer. We also report on the centralization of end-user applications such as wallets and exchanges. The following subsections discuss the taxonomy in detail.

⁷ A complete list of articles is available in Appendix B.
Table 3

Centralization causing factors found in literature and interviews.

Centralization factor	Refereed articles	Interviews
Wallet concentration	R11, R13, R36, R40, R76, R78, R84, R86, R88	I_4, I_5, I_7, I_8
Exchange concentration	R11, R13, R27, R34, R37, R40, R57, R64, R73, R78, R84, R86, R89	$I_1, I_3, I_4, I_5, I_7, I_8$
Reference client concentration	R4, R6, R8, R26, R36, R50, R67, R83	$I_2, I_5, I_8, I_9, I_{10}$
Storage growth rate	R9, R24, R38, R39, R63, R80	I_2, I_{10}
Specialized equipment concentration	R23, R51–R53, R55, R62, R67	$I_4, I_5, I_7, I_8, I_9, I_{10}$
Wealth concentration	R16, R51, R52, R55, R62, R67	$I_1, I_2, I_3, I_4, I_5, I_6, I_7, I_9$
Consensus power distribution R1-R3, R5, R7, R9, R11-R17, R19-R22, R25, R26, R28-R33, R35, R36, R39		$I_1, I_2, I_3, I_4, I_5, I_6, I_7, I_8, I_9, I_{10}$
	R42–R47, R49, R52–R56, R58, R60, R61, R65–R72, R74–R79, R81, R82, R87	
Node discovery protocol control	R59	$I_1, I_2, I_3, I_5, I_{10}$
Geographic distribution	R5, R30, R40, R47, R50, R76	$I_1, I_2, I_3, I_4, I_5, I_6, I_7$
Bandwidth concentration	R35, R87	I_2, I_{10}
Routing centralization	R3, R20, R35	I ₂
Owner control	R14, R18, R26, R41, R48	$I_1, I_4, I_5, I_7, I_8, I_9$
Improvement protocol	R4–R6, R10, R26, R36, R41, R48, R76, R83, R85	$I_1, I_2, I_3, I_4, I_5, I_7$

Table	4
-------	---

Categories of centralization in governance layer.

Ref	Owner control		Improvement protocol		
	Identification	Measurement	Identification	Measurement	
R4	×	x	√	×	
R5	×	×	\checkmark	×	
R6	×	×	\checkmark	Centrality metrics	
R10	×	×	\checkmark	×	
R14	\checkmark	×	×	×	
R18	\checkmark	Fractional measurement	×	×	
R26	\checkmark	×	\checkmark	×	
R36	×	×	\checkmark	×	
R41	\checkmark	×	\checkmark	×	
R48	\checkmark	×	\checkmark	×	
R76	×	×	\checkmark	×	
R83	×	×	\checkmark	×	
R85	×	×	\checkmark	×	

4.1. Governance

Blockchain, like any other information system, is subject to evolutionary changes that are governed by a governance structure. These evolutionary changes may include security patches, scalability provisions, and improvement proposals. Wang et al. (2017) theorizes the relationship between the value proposition of blockchain and the governance structure in place. They reason that the core value proposition of blockchain is rooted in decentralization. This property of decentralization is considered valuable by investors.

Decentralized governance was also indicted as a vital component of public blockchains by our interview participants. 80% mentioned governance as a significant centralization threat (I_1 , I_2 , I_3 , I_4 , I_5 , I_7 , I_8 , I_9). This is best illustrated by a quote from I_1 , with respect to the implication of centralized governance structure: "*if you are talking about the centralization of governance, that for me is the prime example of a private permissioned Blockchain*".

Despite the significance of decentralization for blockchain, Wang et al. (2017) argue that a high level of decentralization may slow down the strategic decision-making process. Contrary to the proposition in favor of some centralization by Gervais et al. (2014) and Wang et al. (2017) argue against the concentration of decision making power by pointing out instances of unilateral decision making by core developers in the short history of bitcoin; for example, when the core developers unilaterally decided to lower the minimum transaction fee. This criticism of governance centralization is shared by Roubini (2018a) who criticizes the centrality of control over governance as it may concentrate the decision power to a few entities involved in governance of the blockchain. Atzori (2015) expands the analysis of blockchain governance issues towards the emergence of blockchain governance oligarchy. Azouvi et al. (2018) conducts an empirical analysis of two of the most prominent blockchain projects, Bitcoin and Ethereum, by comparing the state of governance to other major open-source projects. They conclude that control governance is usually concentrated in a handful of people in Bitcoin and Ethereum, which is a big centralization factor.

As reported by Wang et al. (2017), the centralization on the governance layer may not be detrimental due to the advantages of rapid strategic decision-making. We expand on the argument in favor of some centralization (Wang et al., 2017) in Section 6, where we discuss how the adverse impact of centralization varies across the different layers of the taxonomy.

Based on the literature review and subsequent interviews, we further divide the issue of governance into *owner control* and *improvement protocol*. These results are presented in Table 4.

A.R. Sai et al.

4.1.1. Owner control

As described by Wang et al. (2017), the developers of the blockchain often retain some control over the implementation on the governance level. This can be in the form of, for example, the native cryptocurrency owned by the developers. Wang et al. (2017) describes this as *Owner Control*.

Measurement Technique : This type of owner control can be measured by examining the total cryptocurrency accumulated by the owners in the early adoption period (Wolfson, 2015). This early adoption period also includes the pre-mined⁸ cryptocurrency (Wang et al., 2017). We report studies such as (Chohan, 2019; Wolfson, 2015) that have implemented a proportional measure to quantify owner control. Owner control can be measured as the fraction of the total allowed cryptocurrency if the supply is capped, as measured by Eq. (1), where $C_{OurnerControl}$ represents the fraction of total cryptocurrency that the owner controls.

$$C_{OwnerControl} = V_{OwnerBalance} / V_{CappedSupply}$$

(1)

(2)

If the supply is uncapped, owner control is measured as the fraction of total currency in circulation, as illustrated in Eq. (2).

 $C_{OwnerControl} = V_{OwnerBalance} / V_{CurrentSupply}$

Most interview participants indicated that the use of fractional measurement for owner control was appropriate. However, I_9 suggested a refinement: "The fractional calculation of the owner control varies with the supply; a simpler approach might be to use a metric such as how much power over the network can be achieved with the money in the owner control. How much hardware can you afford, and what hash power can you get with it. Relating the cryptocurrency to the hashing power would be more informative".

Implication of high owner control: Depending on the consensus mechanism used, the owner control has severe impacts on the network. This adverse impact is particularly worrying in the case of Proof-of-stake based cryptocurrency, where the consensus power is determined by the quantity of native cryptocurrency owned by the participant. Having a large amount of pre-mined or early adoption period accumulated cryptocurrency will give the owner a significant advantage over others, resulting in a more centralized network. This high consensus power pose a security threat as an owner with over 50% consensus power can conduct a double spending attacks. Ethereum is a prime example of such wealth concentration due to pre-mined cryptocurrency.

The Ethereum platform was crowdfunded by investors who were rewarded in the form of ETH⁹ during the creation of the first block in Ethereum. An estimated 60 Million ETH were distributed among the early investors; another 12 Million were distributed among the developers of Ethereum (Etherscan, 2019a). We calculate the value of $C_{OuvereControl}$ by considering the 12 Million pre-mined ETH that developers control and the total current supply of ETH obtained (from Etherscan, 2019b):

$$C_{OwnerControl} = 12,000,000/106,514,407.78 = 0.11$$
(3)

It should be noted that the value of $C_{OwnerControl}$ feeds into the issue of Wealth Concentration, which is a significant cause of economic centralization. A high wealth concentration in a cryptocurrency is against the founding principle and premise of cryptocurrency providing a more even monetary system. This can consequently disincentivize the adoption.

4.1.2. Improvement protocol

As discussed earlier, evolutionary changes require blockchains to have a robust governance structure in place. As decentralized blockchains do not have any authorized entities moderating the changes, the process of moderation is delegated to the participants. Bitcoin improvement protocol (BIP) is a prime example of such an improvement system (Anceaume, Lajoie-Mazenc, Ludinard, & Sericola, 2016). The formal voting protocol, such as that in BIP, is used to establish consensus over proposed changes, often through voting. 60% of interview participants (I_1 , I_2 , I_3 , I_4 , I_5 , I_7) mentioned that the improvement protocol performs an essential function in the network with I_7 suggesting: "Wheever controls the improvements will inevitably shape the future of the network".

The literature review points out the similarities between the Python Enhancement Proposals and BIPs, both of which heavily draw from the *"canonical"* approach to consensus (De Filippi & Loveluck, 2016). In the *"canonical"* based BIP, all the suggested changes have to be made available to the public for open discussion. However, the final decision as to how proposed changes will be implemented is taken by the core developers (Gervais et al., 2014).

Measurement Technique: The centralization in a formal voting protocol is measured by analyzing the moderation control. If specific developers or owners can moderate the voting, the moderation may jeopardize changes these that developers or owners disagree with. Thus the determination of the control level is done by examining the voting protocol in place and the controls imposed on it.

As public blockchains often have an open platform for proposing improvements, such as BIP for Bitcoin, and EIP for Ethereum, Azouvi et al. (2018) suggests reviewing the number of improvement proposals made by each author and the respective states of those proposals (i.e., approved, rejected or under review). The authors also suggest reviewing the comments on each proposal to examine the discussions. Based on the data obtained from the author/number of proposals, complemented by comments per author on the proposal, Azouvi et al. (2018) suggests calculating metrics for centralization measurement. Fig. 7 illustrates this measurement technique graphically.

⁸ Pre-mined cryptocurrency refers to the native cryptocurrency issued with the creation of the first block in the blockchain.

⁹ ETH is the ticker mark for Ether, the cryptocurrency used by Ethereum platform.



Fig. 7. Improvement protocol centralization measurement technique.

These centrality metrics include Mean, Median, interquartile range (IQR), and interquartile mean (IQMean). IQR is a measure of variability that assists in locating where the majority of values lie in the data sample. It is calculated as the difference between 75th and 25th percentiles of the data. However, IQR is sensitive to noisy outliers, which can impact the overall result. This can be overcome by using the IQMean, which allows us to eliminate the outliers from our data set by calculating the median of IQR.

Implication of control over improvement protocol: If a subset of all participants moderate the improvement protocol, it will result in control over improvements or modifications to the network. The debate over block size in Bitcoin an example of an issue arising due to this type of control over the network (Bitcoin, 2019; De Filippi & Loveluck, 2016). Other significant control implications over the improvement protocol include the unilateral decision making in both Bitcoin and Ethereum, where the governance structure implemented a change not widely supported by the community. This includes the notable transaction fee reduction in Bitcoin (Gervais et al., 2014) and Ethereum hard fork due to DAO attack which led to the subsequent creation of Ethereum classic (Wirdum, 2016). More incidents of unilateral decision making include the changes to the Ethereum consensus algorithm in 2018, where developers decided to modify the algorithm to disable newer mining hardware (Kim & Zetlin-Jones, 2019). These incidents not only represent the lack of a systematic governance model in terms of improvement but also present a challenge in terms of newer participation and updates. This type of centralization impacts the presumed open nature of the Blockchain, which is one of the core contributions of Blockchain to the field of financial technologies.

4.1.3. State-of-the-art for centralization on the governance layer

Based on the literature review, we report that there are two distinctive approaches to centralization in governance. In the first approach, pioneered by Wang et al. (2017), governance centralization is essential for rapid strategic decisions. This approach is countered by the empirical analysis of Azouvi et al. (2018). Those authors report that other non-cryptocurrency open source projects have attained a higher level of decentralization in governance, specifically in the form of improvement protocol without the need for centralized control. These two contrasting approaches highlight the need for a more in-depth analysis of the importance of rapid strategic decision making in the context of blockchain various other open-source projects.

4.2. Network

The network layer acts as the information dissemination mechanism for the blockchain instance. As the decentralized network cannot have centralized nodes that act as relay points to transmit messages between the participants, the network is largely a peer-to-peer system. The network layer acts as the information dissemination mechanism for the blockchain instance. This peer-to-peer network serves as an essential security and usability measure as pointed out by I_8 : "In this peer to peer network, there is no single point of failure and participants can join and leave the network without risking interruption or degradation of the network".

Network connectivity of a node is an important aspect of performing the mining operation (Sapirshtein et al., 2016). Higher network connectivity results in a higher likelihood of adding the next block on the longest chain as the miner can propagate the block to a large number of nodes in the network. This interplay between the reward from adding a block to the blockchain and network connectivity has resulted in networking phenomena such as strategizing networking resource concentration in the form of bandwidth (Gencer et al., 2018) and strategizing geographic distribution of nodes in the network (Kim et al., 2018; Roubini, 2018b).

Based on the literature review, we identify another source of centralization on the network layer as the topology formation of the network. This formation includes the node discovery protocol for finding peers in the network (Neudecker & Hartenstein, 2018) and the routing structure of the network (Apostolaki, Zohar, & Vanbever, 2017). The relevant studies identified by our review are presented in Table 5. We describe each of the outlined factors in detail in the following Subsections.

Table 5

Categories of centralization in network layer.

Ref Node discovery			Geographic distribution		Bandwidth		Routing	
	Identification	Measurement	Identification	Measurement	Identification	Measurement	Identification	Measurement
R3	×	×	×	×	×	×	\checkmark	AS coverage
R5	×	×	\checkmark	×	×	×	×	×
R20	×	×	×	×	×	×	\checkmark	×
R30	×	×	\checkmark	×	×	×	×	×
R35	×	×	\checkmark	Latency	\checkmark	Clustering	\checkmark	AS coverage
R40	×	×	\checkmark	×	×	×	×	×
R47	×	×	\checkmark	×	×	×	×	×
R50	×	×	\checkmark	×	×	×	×	×
R59	\checkmark	×	×	×	×	×	×	×
R76	×	×	\checkmark	×	×	×	×	×
R87	×	×	×	×	\checkmark	×	×	×

4.2.1. Node discovery protocol control

In a peer-to-peer topology, participating nodes directly communicate with other participants to transmit data packets. A node discovery protocol is used to discover nodes in the network with which to communicate (Miller et al., 2015). The node discovery protocol often relies on a set of seed DNS nodes that distribute the address of other active nodes on the network. These predefined DNS nodes may be a potential source of security threat, as demonstrated by Jin, Zhang, Liu, and Lei (2017) and Tapsell, Akram, and Markantonakis (2018). If one of the seed nodes becomes inaccessible, it may result in many participants of the network becoming undiscoverable. As the new nodes in the network discover others by querying these predefined seed DNS nodes, the literature identifies seed nodes as a contributor to centralization on the network layer (Neudecker & Hartenstein, 2018).

Measurement Technique : After the review of all relevant articles in our study, we conclude that no measurement technique focuses on the Node Discovery protocol. Studies such as (Jin et al., 2017; Tapsell et al., 2018) investigate the issue of seed DNS nodes from a security perspective, specifically focusing on the single point of failure issue. We reason that further investigation into centralization in node discovery level is warranted due to the significant security threats that it poses.

Implication of control over DNS: Centralized DNS services are linked to security threats in the network (Jin et al., 2017). They also allow the DNS owners to observe the participants of the network. These centralized DNS services can also act as a single point of failure, which is of particular concern in the case of a Denial of Service attack (Dietrich, Long, & Dittrich, 2000). As core developers select these DNS nodes, the issue of node discovery protocol also feeds into that of trust in the core developers (Tapsell et al., 2018). A malicious developer can also change the DNS seed nodes to conduct an eclipse attack. Several Monte Carlo simulations have shown the effectiveness of such eclipse attacks on Bitcoin and Ethereum (Heilman, Kendler, Zohar, & Goldberg, 2015).

4.2.2. Geographic distribution

Bitcoin and similar cryptocurrencies have been able to gain significant attention from governments around the world due to their decentralized uncensored nature. This has prompted many to argue that a significant concentration of the nodes in any geographic area may be a threat to the network (Roubini, 2018b). This type of geographic concentration may lead to centralization on the network layer as the nodes become prone to geopolitical manipulation. 70% of interview participants indicated that geographic concentration is harmful to the network. I_6 suggested that geographic centralization may be disadvantageous for miners who are not centrally located: "I fear that in a geographically-focused network, people within the same geographic location will have an edge over others, they will receive and send transactions first".

The nodes are distributed over the participating countries in the network. In an ideal case, the distribution of nodes should be equal in all participating countries so as to be able to withstand a geopolitical blockade. Findings from our review suggest there is a trend towards geographic concentration of nodes in both Bitcoin and Ethereum (Gencer et al., 2018; Khairuddin & Sas, 2019; Kim et al., 2018; Roubini, 2018b).

Measurement Technique : Our review suggests that the geographic location measurement in blockchain can be done by measuring latency in the peer-to-peer network (Gencer et al., 2018; Kim et al., 2018). This approach draws heavily from Saroiu, Gummadi, and Gribble (2001), where the authors proposed using latency as a measurement tool in Gnutella. Gencer et al. (2018) first proposed measuring the distance between their geographically distributed nodes and other peers in the network by sending a data packet and measuring the round-trip time. Based on the round-trip time, Gencer et al. (2018) calculated upper and lower bounds between two remote peers in the network. If two nodes take a similar time to respond to the data packet sent by their nodes, it is reasoned that these two nodes are likely geographically close. This approach is further refined by Kim et al. (2018), who consider the average of bounds for final latency estimation.

Fig. 8 illustrates this measurement technique graphically using a toy example. In this example, we have two geographic regions A and B. To identify a blockchain network participant's relative geographical locations, we deploy two measurement nodes that send a data packet to the network participant and wait for the response. Upon the receipt of a response, we can calculate the network latency. In this toy example, the measurement node in geographic area A returns a lower latency; thus, we can assume that the blockchain participant is geographically closer to area A than area B. In their analysis (Gencer et al., 2018), the authors conduct this experiment with a large number of measurement nodes spread out geographically.



Fig. 8. Geographic distribution measurement technique.

Implications of geographical centralization: The most prominent issue with geographic centralization is the potential for geopolitical manipulation of the network (Roubini, 2018b). Other issues with geographic clustering include the possibility of faster transmission of packets to nearby nodes promoting faster network propagation. This can lead to more clustering, since participant must propagate the solution to the majority of the network in order to get rewarded in Proof-of-work based blockchains. If the majority is located in a geographical cluster away from the participant, that may translate to a loss of revenue. As suggested by Gencer et al. (2018), a low number of geographic clusters are considered good for the decentralization of the network. This is due to the association of potentially high block rewards due to faster network propagation. As shown in Sapirshtein et al. (2016), network connectivity is directly related to the ability to successfully conduct selfish mining attacks, which can support a double spending attack.

4.2.3. Bandwidth concentration

In a public blockchain's peer-to-peer network, the network bandwidth often acts as a crucial factor in the successful propagation of data packets. In Proof-of-Work based blockchain, every consensus cycle acts as a race to first calculate the solution to the cryptographic puzzle followed by dissemination of the solution to a majority of the network. Dissemination requires a large number of network connections with peers in the network, thus increasing the bandwidth requirements. This arms race to attain higher bandwidth may lead to the centralization of mining equipment to services like a centralized data center with high bandwidth (Gencer et al., 2018).

Measurement Technique : Gencer et al. (2018) proposed measuring the bandwidth of each peer by requesting a large amount of data and estimating the speed by observing the time taken for the transmission. Once they estimate the speed of each accessible peer, they calculate and cluster the provisioned bandwidth in groups.

Implication of bandwidth concentration: A high bandwidth requirement may limit the participation to only the participants with significant bandwidth (Zheng et al., 2018). It may also result in a high concentration of networking devices in centralized spaces such as data centers (Gencer et al., 2018). This potential increment in bandwidth requirement may limit the participation to only those entities with high network capabilities making the consensus participation not viable in a domestic setting. The inability to participate in the network violates the open nature of the public blockchain preventing a widespread adoption of the technology.

4.2.4. Routing centralization

As public blockchain networks run over the existing networking stack, they rely on the networking structure used by IP (Internet Protocol). Centralization present in the networking structure of IP transfers to the blockchain as well. Our review reports that this centralization has been studied in blockchain from the privacy (Feld, Schönfeld, & Werner, 2014) and security (Apostolaki et al., 2017) perspectives. Gencer at al. (2018) reports that concentration on AS-Level¹⁰ as a source of centralization for a public blockchain (Gencer et al., 2018). Interestingly, none of the industrial participants mentioned this concern unprompted, suggesting that it might be more of an academic concern than a real-world one. However, when the concern was mentioned, one industry participant agreed.

Measurement Technique : Our review suggests that there is a common network traversing strategy used to determine the network structure from the AS-Level perspective (Apostolaki et al., 2017; Feld et al., 2014; Gencer et al., 2018). To measure the number of ASes in a peer to peer network, the observer node traverses the network by recursively collecting IP addresses of each peer and querying every reachable address. This process is repeated until no new reachable nodes are available in the IP list. For the determination of AS of each IP, Feld et al. (2014) recommend using Maxmind's free Geo API.¹¹

Implication of control over ASes: Centralization on AS-Level is reported to have privacy implications for blockchain users as it allows more traceability on a network level (Feld et al., 2014). This concentration of IP addresses under a few ASes is directly linked with potential network security issues in Bitcoin (Apostolaki et al., 2017) and Ethereum (Gencer et al., 2018). However,

¹⁰ Autonomous systems (AS) in computer networks refers to the collection of connected IP routing prefixes under the authority of one or more networking entities.

¹¹ https://dev.maxmind.com/geoip/geoip2/geolite2/.

Table 6

Categories of centralization in consensus layer.

Consensus power distribution	Selected studies
Identification	R1, R2, R3, R5, R7, R9, R11, R12, R13, R14, R15, R16, R17, R19, R20, R21, R22, R25, R26, R28, R29, R30, R31, R32, R33, R35, R36, R39, R40, R42, R43, R44, R45, R46, R47, R49, R52, R53, R54, R55, R56, R58, R60, R61, R65, R66, R67, R68, R69, R70, R71, R72, R73, R74, R75, R76, R77, R78, R79, R80, R81, R82, R87
Factor measurement	
Percentage based measure	R1, R7, R12, R14, R21, R26, R29, R31, R33, R35, R36, R43, R46, R47, R49, R49, R53, R55, R56, R60, R61, R67, R71, R72, R73, R77, R78, R80
Gini	R15, R16

these privacy and security threats remain largely academic with no real world incident reports in our sample set of articles. This is further evident through our interviews, where no academic or industrial experts pointed to control over ASes as a centralization threat unprompted.

4.2.5. State-of-the-art for centralization on the network layer

In a peer-to-peer, network-based blockchain system, both the network connectivity and the network capabilities have an impact on the likelihood of profit (Sapirshtein et al., 2016). Our survey reports on four types of network-based centralization: node discovery, geographic distribution, bandwidth, and routing. Among the reported centralization avenues, there are no measurement techniques to quantify node discovery protocol centralization despite the security threats associated with centralization (Jin et al., 2017). We reason that a further investigation into the centralization of node discovery protocol is warranted. We also report that the research into the geographic distribution and bandwidth centralization is primarily focused on Bitcoin and Ethereum. Due to the association of the monetary reward and the network connectivity and capabilities, we reason that a further empirical investigation into the network layer for more cryptocurrencies may assist in better understanding network participation's profitability.

4.3. Consensus

The consensus layer establishes an agreement on a single state of the data in the public blockchain. As described in Section 2.2, in the case of Proof of Work, it is attained by inducing a race to solve a mathematical problem. The first person to solve and propagate receives a monetary reward as an incentive. The likelihood of finding the solution to the mathematical problem depends on the computational power devoted to the solution. Thus a high concentration of computational power is a direct signifier of centralization in the blockchain. As identified by articles in Table 6, the **consensus power distribution** is a key contributor to the centralization of the Proof-of-Work based blockchain. Eight interviewees mentioned this aspect unprompted, suggesting that this is a prevalent concern. In this subsection, we review how the literature defines and measures the consensus power centralization.

4.3.1. Consensus power distribution

In the case of a Proof-of-Work based blockchain, the Consensus power is also known as the hash power of the miner (participating node). The centralization of hash power can pose a significant security threat to blockchain solutions such as Bitcoin and Ethereum. One key contributing factor to centralization is commercial mining pools. The income from mining operations depends on the probability of finding and propagating the solution of the puzzle before everyone else. The probability of successfully calculating the solution depends on the hash power of the computing device used for the calculation. Lower probability leads to a lack of stable income and may prompt users to mine as a group and share the profit. This group mining is also known as pooled mining (Lewenberg, Bachrach, Sompolinsky, Zohar, & Rosenschein, 2015). Based on the analysis of the shortlisted literature, we report that the concept of pooled mining in itself is not considered a threat to the decentralization of the network; however, the literature is in agreement over the harms of a centrally run commercialized mining pool. In these centrally run mining pools, the pool manager decides which transactions to include in a block and subsequently distributes the workload among participants of the pool. This type of structure requires trusting the manager of the pool thus limiting the decentralization in the blockchain (Chesterman, 2018).

Measurement Technique : Studies including (Beikverdi & Song, 2015; Gencer et al., 2018; Judmayer, Zamyatin, Stifter, Voyiatzis and Weippl, 2017; Sai et al., 2019a), have deployed an experimental setup to measure consensus centralization. Judmayer, Zamyatin et al. (2017) refer to this approach as a "block attribution scheme". In this experimental set-up, a participating node is connected to the blockchain that actively sniffs the network to extract mined blocks and coinbase addresses.¹² The coin base address is then used to query public blockchain explorers to determine if it belongs to a known mining pool. Based on the results, a list of the mining pools and the proportion of the blocks mined by each respective public mining pool is constructed. Using this approach, we can calculate the proportion of total computational power that each mining pool controls. Fig. 9 illustrates the block attribution scheme graphically.

This proportion can be represented as a percentage value as suggested by referred articles in Table 6 or by using the Gini values, based on the Lorenz Curve (Bruschi, Rana, Gentile, & Sciuto, 2019; Caccioli, Livan, & Aste, 2016).

¹² The coin base address refers to the address of the node that gets the reward for successfully mining a block.



Fig. 9. Block attribution scheme.

Gini Value Measurement: These are economic measures of inequality (Dorfman, 1979; Gastwirth, 1971) for consensus power concentration (Bruschi et al., 2019; Caccioli et al., 2016).

The *Lorenz curve* is a graphical representation of the distribution of wealth. The curve illustrates the proportion of the income earned by any given percentage of the population. This curve has proven to be of significant importance in economic disparity measurement. To numerically describe this distribution, we can use the Gini Coefficient, which is based on the difference between the Lorenz curve and the line of equality.¹³ We can calculate the Gini Coefficient as follows:

$$Gini = A/(A+B) \tag{4}$$

where A is the area between the line of equality and Lorenz curve, and B is the area under the line of equality. The value of Gini can range between 0 to 1, where 0 represents complete equality, and 1 represents complete inequality.

Implications of consensus power centralization: The impact of centralization in consensus power has been widely studied in security literature (Chen et al., 2017; Gervais et al., 2016; Karame, 2016; Sai et al., 2019a; Sapirshtein et al., 2016; Zhang et al., 2019). A concentration of 26% in proof of work-based blockchain can lead to successful selfish mining attacks. Whereas a consensus power concentration of over 51% can result in a 51% attack.

Smaller cryptocurrecies tend to be more prone to 51% attack as evident by successful attacks on Aurum Coin, Bitcoin Gold, Ethereum Classic, Flo Blockchain, Monacoin, Verge, Vertcoin and ZenCash (Sayeed & Marco-Gisbert, 2019). These 51% attacks have, on average, resulted in a loss of \$2.5 million per cryptocurrency (Sayeed & Marco-Gisbert, 2019). The significance of these attacks is evident by the agreement of all our interviewees on the centralization implications of a 51% attack caused by consensus power concentration.

4.3.2. State-of-the-art for centralization on the consensus layer

Recent successful double-spending attacks due to the consensus power centralization have highlighted the need for a more encompassing understanding of the incentives behind the honest behavior for participation. Based on the results of Sai et al. (2019a), it seems that a better understanding of the economics behind the consensus system is needed to ensure secure operation. This is especially important for smaller cryptocurrencies as the barrier to conduct a 51% attack is lower when compared to major cryptocurrencies. We suggest a further in-depth investigation of the economic incentives behind conducting a double-spending attack against smaller cryptocurrencies be conducted.

4.4. Incentive layer

Bitcoin and similar decentralized cryptocurrencies are inherently dependent on the economics associated with rewards (Sai et al., 2019a). Sai et al. (2019a) reports that the exchange rate of Bitcoin is related to the overall consensus power of the network. If the exchange rate falls below a given threshold of profitability, the participants of the network may withdraw from active mining, which may result in a fall in overall hashing power of the network. A low value of hashing power of the network makes it easier for attackers to attain a higher consensus proportion; thus it may increase the threat of selfish mining and 51% attack. This interplay between the monetary aspect of public cryptocurrencies and security makes it essential to inspect centralization on the economy

¹³ Line of equality refers to the equal distribution of hashing power among miners.

Table 7					
Categories	of	centralization	in	incentive	layer.

Ref	Wealth concentration	
	Identification	Measurement
R16	√	×
R51	\checkmark	Gini
R52	\checkmark	Percentage value
R55	\checkmark	×
R62	\checkmark	Percentage value
R67	\checkmark	Gini

driven incentive aspect of the network. A high concentration of wealth to a select few may be an aspect of centralization that can prove to be harmful to the network. Attacks such as the Whale Transaction Attack (Liao & Katz, 2017) have exploited wealth concentration. In a whale transaction attack, the attacker attempts to induce disagreement¹⁴ between the participants by providing a high transaction fee in an already published block.

The issue of wealth concentration was raised by 60% of our interview participants unprompted. P_7 , for example, noted how they focused on wealth concentration: "In general, I follow the money. If the trail of funds leads to one natural person or group of natural persons (regardless of number of addresses), then the process is relatively centralized along the spectrum of centralized-decentralized blockchain".

Table 7 outlines the result of our review, identifying relevant articles and shortlisted techniques for measurement. In this subsection, we review the centralization based on Wealth Concentration in depth. This type of centralization may be of significance for a blockchain solution that employs a wealth-oriented consensus mechanism such as Proof-of-Stake (Kiayias, Russell, David, & Oliynykov, 2017).

4.4.1. Wealth concentration

High accumulation of native cryptocurrency may give a unique advantage to an adversary. The high wealth concentration can also be used to increase the overall cost of transactions (Liao & Katz, 2017), as demonstrated in the iFish attack on the Ethereum network (Cryptoslate, 2018). In the iFish attack, the attacker induced a large number of transactions with a high transaction fee in a short period. This influx of high transaction fees resulted in a considerable increase in the transaction fee. Another form of network abuse arising from high wealth concentration involves transaction fee manipulation by artificially increasing the overall fee required for a successful transaction.

Based on the results from our review, we point that this wealth concentration also has economic impacts on the network. As reported by Kondor, Pósfai, Csabai, and Vattay (2014), already wealthy nodes in the bitcoin's transaction graph tend to increase their wealth at a higher speed than smaller nodes. They call this phenomena the "*rich get richer*" scheme.

Measurement Technique: Wealth concentration measurement is at the center of disparity studies in economics (Gini, 1921). One of the most commonly used measures is the Gini Coefficient calculated from the Lorenz Curve. The wealth concentration is measured in the form of inequality based on the population and what proportion of population controls how much wealth. Translating this directly to the blockchain could mean calculating Gini over a cryptocurrency and all existing addresses on the blockchain. But we argue that this may not be the most efficient way as techniques such as Hierarchical Deterministic Wallets (Gutoski & Stebila, 2015) promote the generation of new addresses for every transaction. To overcome this limitation, Srinivasan (2017) proposes establishing a lower bound value on the cryptocurrency contained in the address for inclusion in the measurement, i.e., a wallet with 0 cryptocurrencies may be excluded from the study, as it most likely resembles an inactive address. Another reported measurement technique is to use a percentage measure. However, a simple percentage measure fails to capture the distribution. Machine learning-based transaction clustering approaches have also been employed to extract behavioral patterns (Hu et al., 2021); these heuristics may also be useful for the calculation of wealth distribution.

Implications of Wealth Concentration: Wealth concentration is linked with a number of potential attacks, such as the possibility of a 51% attack in the case of a wealthy attacker during a fall in exchange rate (Sai et al., 2019a). Whale attack, as discussed above, is another example of a wealth oriented security threat to the network. However, both of these potential attacks are without any real-world incident reports.

One example of wealth concentration in a real-world attack is the transaction fee price manipulation caused by the iFish attack (Cryptoslate, 2018). During the iFish attack, the attacker was able to artificially inflate the transaction fee of Ethereum by 35%. Another example of a wealth oriented attack is the bZx hack, where a smart contract designed for lending Ether was exploited by sending high-value transactions and manipulating the platform (Zmudzinski, 2020).

A public blockchain with high wealth concentration contradicts the foundational notion of a more even and open monetary system. This has a direct implication on the adoption of the technology.

¹⁴ The disagreement is in the form of blockchain fork.

Table 8	
Categories of centralization in operational layer.	

Ref	Storage constrai	nt	Specialized equipmen	cialized equipment concentration		
	Identification	Measurement	Identification	Measurement		
R9	\checkmark	×	×	×		
R23	×	×	\checkmark	×		
R24	\checkmark	Rate of growth	×	×		
R38	\checkmark	×	×	×		
R39	\checkmark	×	×	×		
R51	×	×	\checkmark	×		
R52	×	×	\checkmark	×		
R53	×	×	\checkmark	×		
R55	×	×	\checkmark	×		
R62	×	×	\checkmark	×		
R63	\checkmark	×	×	×		
R67	×	×	\checkmark	×		
R80	\checkmark	×	×	×		

4.4.2. State-of-the-art for centralization on the incentive layer

High wealth concentration in public blockchain poses security threats, specifically in the form of manipulating economics associated with the blockchain system, such as transaction fees and exchange rates. However, measuring the current state of wealth concentration in the public blockchain is a nontrivial problem due to widespread adoption of technologies that aim to increase users' anonymity. Based on our review, we note that the use of clustering techniques and setting a lower bound on the amount of cryptocurrency stored at an address can help establish the state of wealth distribution. We suggest that further investigation into deanonymizing the blockchain can improve the accuracy of the techniques used to calculate the Gini value and that this is an important agenda for research going forward.

4.5. Operational layer

The uncertainty of reward imposes a constraint on participation for rational investors. This reasoning is primarily based on the cost of mining (Sai, Le Gear and Buckley, 2019). A miner can earn rewards in the form of mining incentives and accumulated transaction fees from the mined block but to profitably mine on a Proof-of-Work blockchain, the difference between rewards earned and the expenses of the mining operation should be positive. This is the 'operations' we are referring to in this 'operational' layer. The expenses of mining operations include capital costs such as the acquisition of adequate hardware and other recurrent costs such as the cost of electricity.

After conducting the systematic review, we report two types of centralization associated with operational aspects of the public blockchain. The first is the move from commercially available mining equipment to proprietary application-specific integrated circuit machines. This increased capital, operational cost has proven to be a significant barrier to entry for new miners in Bitcoin (Borge et al., 2017). We categorize this type of specialized hardware centralization as *Specialized Equipment Concentration*.

Another factor that contributes to the cost of mining is the storage requirements for operating on the network. As all full nodes in the network are required to store and process all the transactions, the data stored increases (Dai, Zhang, Wang, & Jin, 2018). This imposes a significant barrier as traditional computing devices may not be able to participate in the network given high storage requirements. This may limit the participation in consensus to only the participants who can afford greater computational resources imposing a constraint on participation. A significant storage requirement may deter users with conventional computing devices from participating in the consensus altogether, resulting in a more centralized network converged on participants with high computational capabilities. This high storage requirement has been discussed as a centralization causing factor (Guo, Gao, Mei, Zhao, & Yang, 2019; Reddit, 2019; StopAndDecrypt, 2018).

In this layer of centralization, interviewee I_{10} had an interesting perspective suggesting a restructuring of contract layer to widen our definition of the layer to include other operational concerns.

In this subsection, we report the centralization caused by the operational cost involved in participating in the consensus of the blockchain. We also manifest the result of our systematic literature review in Table 8.

4.5.1. Size of the blockchain

The traditional computing devices are often limited in-memory capabilities and can only hold a constrained amount of data. Attaining a higher storage capacity may prove to be costly if the growth rate of the storage requirement is significantly high (Guo et al., 2019). This growth in requirement may act as a deterring factor for non-organizational users as the requirement of the investment may be significant (Raman & Varshney, 2017), thus prompting centralization of mining effort.

The issue of storage requirement was articulated by 20% of our interview participants. I_{10} said: "Nothing really stops blockchains from becoming so large that we will run out of capacity. Personally, I have just experienced the first challenge because my Linux partition ran out of capacity; however, if I bought additional hard-disks, I will still be able to run a full node, but it is getting more expensive to run full nodes".



Fig. 10. Storage growth rate measurement technique.

Measurement Technique : To capture the storage-oriented centralization, Raman and Varshney (2017) suggests using the growth rate as a metric. This growth rate is determined based on historical data about the total size of the blockchain. The growth rate can be calculated periodically, ideally after every difficulty recalibration.¹⁵ Fig. 10 illustrates the growth rate measurement technique graphically.

 I_{10} stated their expectations for storage growth rate: "considering that Moore's Law applies to hard drives, it will be interesting to measure the growth rate in comparison with Moore's law".

Implication of high storage requirements: Every blockchain instance may have a different storage requirements, based on its implementation. For example, Bitcoin does not pose significant storage issues as the overall requirement is still low. In contrast, Ethereum has an important storage requirement where the growth rate may limit participation. A growing storage requirement for Ethereum may result in fewer people being able to participate in the network as the participating nodes on Ethereum are expected to store code of smart contracts. A low number of participating nodes increases the likelihood of a successful DDoS attack as it reduces the attack surface.

4.5.2. Specialized equipment concentration

Proof-of-work based blockchains have seen a surge in the overall computational power of the network (Sai et al., 2019a). This surge has made it harder to get higher proportional control over consensus and, consequently, over the rewards associated with incentive. This higher computational requirement has induced an arms race in miners to acquire more efficient and specialized hardware (Ekblaw, Barabas, Harvey-Buschel, & Lippman, 2016). This type of specialized hardware is often not open source and gives the developers an advantage over others (Ekblaw et al., 2016).

60% of our interview participants acknowledged specialized equipment concentration as an issue for a public blockchain. I_7 suggested that this concentration may undermine the whole proposition of public blockchains: ".. but blockchain does not live in a vacuum, so really it was/is the externalities (ASICs and other special hardware for example) that threw the biggest spanner in the experiment".

Measurement Technique: Despite the significance of specialized equipment in Proof-of-work based mining operations, there is no existing metric to measure the centralization of hardware. Based on our literature review, we reason that this may be due to the non-public nature of this specialized hardware. As discussed earlier, most of these hardware implementations are not open source and often not available for public use.

Implication of Specialized Equipment Concentration: As reported by several studies listed in Table 8, the specialized equipment concentration may have given commercial entities an advantage over normal users. If this results in those commercial entries becoming focal, they may utilize the efficient computing equipment to attain higher consensus power and only release it to the public when it becomes less profitable to operate that computing equipment. This approach to hoarding efficient computing equipment is illustrated as the superhashing power dilemma by Bruschi et al. (2019). As a result of our review, we suggest that further investigation is warranted into the measurement of specialized equipment and its impact on centralization.

Apart from the above reported DDoS attack due to the low number of nodes, the specialized equipment requirement severely contains the participation. This higher barrier of entry and lack of profitability with old hardware makes it impractical to contribute to the network without significant investment. This lack of involvement has been shown to increase the likelihood of a successful selfish mining and double-spending attack (Sai et al., 2019a).

4.5.3. State-of-the-art for centralization on the operational layer

Our survey reports on two operational aspects of blockchain: the size of the blockchain and the specialized equipment concentration. The append-only nature of blockchain leads to an ever-increasing size of the ledger that needs to be stored in full

¹⁵ In Proof-of-Work based blockchains, the difficulty of finding the solution of the computational puzzle is updated after a predefined number of blocks to maintain a static block creation time.

Categories	of	centralization	in	application	layer.

Ref	Wallet concentra	ation	Exchange concentration		Reference client	
	Identification	Measurement	Identification	Measurement	Identification	Measurement
R4	×	×	×	×	\checkmark	×
R6	×	×	×	×	\checkmark	Satoshi index
R8	×	×	×	×	×	×
R11	\checkmark	×	\checkmark	×	×	×
R13	\checkmark	×	\checkmark	×	×	×
R26	×	×	×	×	\checkmark	×
R27	×	×	\checkmark	×	×	×
R34	×	×	\checkmark	×	×	×
R36	\checkmark	×	×	×	\checkmark	×
R37	×	×	\checkmark	Centrality	×	×
R40	\checkmark	×	\checkmark	×	×	×
R50	×	×	×	×	\checkmark	×
R57	×	×	\checkmark	×	×	×
R64	×	×	\checkmark	×	×	×
R67	×	×	×	×	\checkmark	×
R73	×	×	\checkmark	Centrality	×	×
R76	\checkmark	×	×	×	×	×
R78	\checkmark	×	\checkmark	×	×	×
R83	×	×	×	×	\checkmark	×
R84	\checkmark	×	\checkmark	×	×	×
R86	\checkmark	×	\checkmark	×	×	×
R88	\checkmark	×	×	×	×	×
R89	×	×	\checkmark	×	×	×

nodes. This growth in the storage requirement is considered an adoption barrier. The current research only suggests metrics to measure it; however, there is a lack of techniques to counter the issue of growth in the storage requirement. We recommend the development of strategies to reduce the storage requirement in order to decrease this type of centralization. In terms of specialized equipment concentration, we note that this is yet to be quantifiably measured and we identify this as a high-potential avenue of future work.

4.6. Application layer

Users often rely on third-party applications to facilitate user interaction with the blockchain (Sai et al., 2019b). These third-party applications include reference implementations, wallets, and exchanges (Gervais et al., 2014). As a result of our review, we report on centralization on these three application layer entities. We also suggest that a monopoly in the user end applications for a blockchain instance is a contributor to the centralization of the blockchain. This issue of centralization on third-party applications was also pointed out by I_8 : "If you remember the catastrophe that centralized implementations such as Mt. Gox, Bitfinex have brought to the blockchain world, you can clearly see the desperate need for decentralization in user-facing applications".

Results from our literature review are outlined in Table 9.

This subsection is a manifestation of the identified centralization prone application layer entities.

4.6.1. Reference client development concentration

As described in Section 2.2, the data layer definition is implemented by a reference client, which acts as the gateway to the blockchain system. As any client that implements the protocol can become a part of the network, it is desirable from the decentralization point of view to have as many developers working on the reference implementation. Each client is expected to fulfill the protocol specification suggested by the core protocol. The development of the core protocol is decentralized by developing an open-source reference implementation. If a select few developers primarily drive the development of the core client, it contributes to centralization (Azouvi et al., 2018; Gervais et al., 2014). The decentralized protocol development factor captures this type of centralization. We note that this centralization is different from the improvement protocol centralization as the focus here on the development of a reference client and not improvements to the protocol.

Despite the reported adverse impact of this type of centralization on the blockchain, in Section 6, we present an argument in favor of some centralization in client development as the developer concentration may be a result of highly skilled developers making useful contributions.

Measurement Technique : (Gervais et al., 2014) suggests examining the number of unique developers contributing to the open-source project with the number of commits on the main core client codebase. This approach is then extended by Azouvi et al. (2018), where they propose using the Satoshi index, which represents the minimum percentage of all contributors required to reach 51% of data contribution.

Implication of reference client development concentration: If only a select few developers work on the reference implementation, they may gain unfair influence over the network. This concentration of power in the hand of select few feeds into the governance issues discussed earlier. As discussed by Azouvi et al. (2018) and Gervais et al. (2014), this type of concentration is harmful to the decentralization of the network as a few developers may influence the implementation of change to the codebase. One of the major implications of influential actors in the public blockchain ecosystem is the defiance of open and equal monetary system assurance provided by the blockchain. As this open and equal system is one of the primary contributions of the public blockchain, the existence of influential entities severely limits systems capabilities to perform in an open and equal manner.

4.6.2. Exchange

Incentives for honest behavior are at the core of the decentralized, trustless transaction ledger. These incentives are often offered in the native cryptocurrency such as BTC and Ether. The real-world value of these cryptocurrencies has been debated (Sai et al., 2019a) with the recommendation that they be determined by the exchange rate to traditional fiat currencies. The exchange of cryptocurrency to traditional fiat currency is aided by application layer entities known as exchanges. These exchanges act as the means of consensus formation around the exchange value. This process is also known as *Price Discovery*. Due to the vital importance of the exchanges, the exchange applications must not be monopolized.

Measurement Technique : To measure the state of centralization in exchanges, Marvin (2017) propose measuring the centrality of exchanges by examining the flow of cryptocurrencies between addresses on the blockchain (Marvin, 2017). Addresses with high centrality in transactions may point to exchanges. This is observed by graphing the transaction flow and identifying nodes with a high degree of centrality. This was followed by the calculation of a Gini Coefficient that reports on the trend of centralization due to exchanges.

Other studies, such as Hileman and Rauchs (2017), have employed a percentage-based value measure, where they measure the proportion of all bitcoin transactions processed by exchanges.

Implication of centralized exchanges: A large number of successful attacks on Bitcoin and Ethereum have focused on exploiting vulnerabilities in exchanges (Chia et al., 2018). These centralized systems act as a single point of failure in case they also serve as a central repository of keys. A prominent example of this is the closure of Mt. Gox due to numerous security flaws leading to loss of Bitcoins owned by its users (Abrams, Goldstein, & Tabuchi, 2014).

Attacks on centralized exchanges not only impact the users of the exchange but the broader cryptocurrency community as it can instill doubts over the security of the ecosystem. These security attacks contribute to the barring trust and adoption by the wider community.

The use of these centralized exchanges may also be reflected through the uneven wealth distribution in the blockchain. Using the block attribution scheme discussed in Section 4.3.1, we report that major centralized exchanges such as Binance can store over \$ 1 Billion in a single wallet location.¹⁶

4.6.3. Wallet concentration

Wallet applications are another form of centralized service on the application layer, as these applications are often developed and maintained by centralized organizations (Sai et al., 2019b).

Measurement Technique: Based on the review of the relevant literature, we report that there are no suggestions regarding the measurement of wallet concentration. We reason that this may be due to the nature of how wallets operate in a closed commercial environment. However, as most of these wallets use an exchange service to transmit funds such as Coinbase (Hileman & Rauchs, 2017), it may be reasoned that exchange centralization may provide a rough proxy for wallet based centralization as well.

Implication of centralized wallet: Applications such as wallets have been identified as a single point of failure and are considered a security threat (Sai et al., 2019b). A high concentration of wealth in centrally managed wallets may give the host an advantage feeding into the issue of wealth concentration. This concentration may also result in a dependence on centralized organization, consequently reducing the decentralization.

Similar to exchanges, a centralized wallet poses a potential barrier of entry in the ecosystem. Due to the technical ability required to host their wallets, most end users tend to prefer hosted wallets, which provides attackers with a small attack surface. This can aid attackers in conducting more targeted yet profitable attacks on the centralized wallet hosting service.

4.6.4. State-of-the-art for centralization on the application layer

Based on the literature review, we report that the dependence on centralized third party services for user-end applications such as exchanges, wallets, and reference client leads is a centralization causing factor. To measure the reference client centralization, Azouvi et al. (2018) suggest using the Satoshi index. However, the empirical data present in the surveyed reports is primarily focused on Bitcoin and Ethereum. We suggest that further investigation into smaller cryptocurrencies may assist in better understanding the current state of centralization across cryptocurrencies. Another form of user-end application is wallet applications. Based on our survey, we report that there are no measurement techniques to quantify this type of centralization. The presence of centralized wallets can also lead to more wealth concentration and this is specifically true for blockchain addresses used by centralized exchanges, as reported earlier. As such we propose this as an important area for future research.

¹⁶ The identified Binance wallet address on Bitcoin leader is 1NDyJtNTjmwk5xPNhjgAMu4HDHigtobu1s.

Table 10

State of centralization in Bitcoin and Ethereum.

Centralization factor	Bitcoin	Ethereum
Wallet concentration	No measurement	No measurement
Exchange concentration	7 exchanges served more than 97.24% of all trades	7 exchanges served more than 28.27% of all trades
Reference client concentration	Single developer authored 25.11% of all files	Single developer authored 40% of all files
Storage growth rate	0.5 GB per week	0.68 GB per week
Specialized equipment concentration	No measurement	No measurement
Wealth concentration	Gini = 0.56	Gini = 0.64
Consensus power distribution	Top 4 mining pools with 50.36% consensus power	Top 4 mining pools with 63.04% consensus power
Node discovery protocol control	No measurement	No measurement
Geographic distribution	Network latency 26.7% less than Ethereum (Gencer et al., 2018)	26.7% higher than Bitcoin (Gencer et al., 2018)
Bandwidth concentration	1.9 to 2.7 times greater than Ethereum (Gencer et al., 2018)	1.9 to 2.7 times less than Bitcoin (Gencer et al., 2018)
Routing centralization	30% network with 10 ASes (Apostolaki et al., 2017)	28% network with 1 AS (Gencer et al., 2018)
Owner control	$C_{OwnerControl} = 0.033$	$C_{OwnerControl} = 0.11$
Improvement protocol	Mean = 11.41, Median = 2.0, IQR = 6.5, IQMean = 2.95	Mean = 9.16, Median = 2.0, IQR = 5.0, IQMean = 2.56

5. State of centralization in Bitcoin and Ethereum

The following subsection provides an overview of empirical evidence specific to the two most prominently used blockchain-based cryptocurrencies: Bitcoin and Ethereum. We present the view of the literature on the centralization of these two cryptocurrencies. Where feasible, we also report the present state of centralization by conducting measurements following the taxonomy guidelines. To structure this investigation, we use the initial taxonomy. The results from this investigation are manifested in Table 10.

5.1. Governance layer

Owner Control: Satoshi Nakamoto is largely credited for the authorship and development of bitcoin (Nakamoto, 2008). This as-yet unknown individual or organization is said to have performed active mining in the early days of Bitcoin, accumulating a considerable amount of BTC (Bohr & Bashir, 2014). As Bitcoin implements provisions for anonymity, the amount of BTC held by Satoshi is not publicly known. Based on the data obtained from the Bitcoin blocks mined in 2009 (Blockchain luxembourg s.a, 2019; Sergio, 2013) performed clustering of similar wallets to identify large entities gathering BTCs. They also report that the largest gain of around 700,000 BTC belonged to a single entity performing mining in 2009. This gives us a value of $C_{OwnerControl} = 0.033$ for Bitcoin.

This value is significantly less than that of Ethereum, where the value of $C_{OwnerControl}$ is 0.11. Bai, Zhang, Xu, Chen, and Wang (2020) argues that Ethereum is very unfair since "the rich are already very rich". This high wealth disparity may allow select participants to conduct attacks based on economic manipulation of the Ethereum ecosystem, such as the Whale transaction attack and transaction fee manipulation. A reason for the high value has been presented in Section 4.

Improvement Protocol: According to Bitcoin (2019), all changes must be approved by all the developers of the core client. However, Gervais et al. (2014) reported on one violation of this rule. In this violation, a subset of developers unilaterally decided to lower the minimum transaction fee to 0.0001 BTC. This illustration strengthens (Gervais et al., 2014) questioning of the transparency, in the process of handling improvement proposals.

Azouvi et al. (2018) measures the centralization by calculating the centrality metrics for both Bitcoin and Ethereum reported in Table 10. They report that the collected commits and comments data set contained many outliers pointing out that the top 25% of developers contribute significantly more than others. They also point out that in their data set for Ethereum, a vast majority of EIP contributions are from a single user, Vitalik Buterin, the founder of Ethereum. They report a similar trend in Bitcoin, where only a handful of people are contributing to the improvement protocol allowing these select groups of people to dictate the changes that are implemented in the protocol. In the past, the block size debate surrounding the scalability has often been cited as a prime example of this type of governance control. Based on our current analysis, we report that the state of centralization in improvement protocol remains largely unchanged since the calculation of Azouvi et al. (2018) with Vitalik Buterin still dominating the EIP contributions in Ethereum and Gavin Andresen remaining responsible for the vast majority of accepted BIP in Bitcoin¹⁷

5.2. Network layer

Node Discovery Protocol Control: As reported in Section 4, our literature review suggest that no study has focused on the measurement of centralization of seed DNS nodes. We postulate that further investigation is required to measure this type of centralization adequately.

¹⁷ The results from our improvement protocol analysis on 02/01/2020 are available at www.github.com/ashishrsai/centralization.



Fig. 11. Consensus centralization in Bitcoin and Ethereum.

Geographic Distribution: Gencer et al. (2018) conducted an extensive review of both Bitcoin and Etherum to measure centralization in the network layer. They reported that the Bitcoin network is more geographically centralized than Ethereum. The average peer-to-peer network latency of Ethereum is 26.7% higher than Bitcoin, suggesting that Ethereum nodes are located at a greater geographic distance. They reason that this is due to the data center focused approach to mining for Bitcoin, whereas Ethereum can be mined by using consumer hardware. This association between geographical distribution and operational centralization neatly illustrates the interdependency between different aspects of centralization, even those based in different layers.

Bandwidth Concentration: Gencer et al. (2018) states that nodes in Bitcoin tend to have about 1.9 to 2.7 times more network bandwidth than Ethereum nodes. They also report that based on the bandwidth, it can be assumed that Bitcoin nodes are located in data center clusters, whereas Ethereum exhibits a more spread out distribution of bandwidth.

Routing Centralization: Feld et al. (2014) reports that 30% of the bitcoin network was only made up of 10 ASes, which presents a level of security threat. This work was expanded by Apostolaki et al. (2017), where they report that 13 ASes covered about 30% of the network but only consisted of 36 IP prefixes. These 36 IP prefixes cover about 50% of mining power. However, the only investigation that has reported on AS-Level centralization in Ethereum, Gencer et al. (2018) reports that 28% of Ethereum nodes belonged to a single AS.

Replicating the experimentation of Gencer et al. (2018) to get the current measurement for network layer centralization requires access to an extensive geographically spread beacon network. Due to resource constraints, we leave such experimentation for other, larger research groups.

5.3. Consensus layer

Centralization of consensus power of bitcoin has been studied thoroughly in the literature (Beikverdi & Song, 2015; Gervais et al., 2014, 2016; Karame & Androulaki, 2016; Sai et al., 2019a). Beikverdi and Song (2015) uses a percentage based centralization value to derive a new metric called Centralization Factor. They report that at the beginning of 2011, 30% of all hashing power was controlled by eight mining pools. This concentration sees a significant increase in 2014 when, according to Gervais et al. (2014), the top mining pool alone controls close to 40% of all hashing power of the network.

Gencer et al. (2018) expands these analyses by also examining Ethereum's network. During the observation period, Gencer et al. (2018) reports that Bitcoin had a less centralized consensus mechanism than Ethereum. On average, the top four mining pools in Bitcoin controlled 53% of the hashing power, whereas in Ethereum the top three mining pools controlled 61% hashing power.

We followed the block attribution scheme for both Bitcoin and Ethereum as discussed in Section 4.3.1 for a week's period.¹⁸ Our results are in line with the observations of Gencer et al. (2018), with the top four mining pools in Ethereun constituting 63.04% of the hashing power, whereas, in Bitcoin, the top four mining pools controlled 50.36% of the hashing power. The results from our analysis are illustrated in the Fig. 11.

5.4. Incentive layer

According to Malik (2016), as of 2016, 11,000 unique Bitcoin addresses, out of a total of 12 million, contained 75.2% of all Bitcoin in circulation. This disparity shows a significant concentration of wealth to a select few. Chohan (2019) also supports the claim of significant inequality in the Bitcoin network. The author claims that the level of inequality reflects that of traditional economies and voids the proposed purpose of Bitcoin: decentralization. Gupta and Gupta (2017) conducted an in-depth investigation of the inequality of Bitcoin. They report that Bitcoin had a Gini value of 0.995 in the year 2013. This result is then refined by Srinivasan

¹⁸ From DEC-25-2020 12:00:00 PM +UTC until JAN-01-2021 12:00:00 PM +UTC.



Fig. 12. Trend of Gini value in Bitcoin and Ethereum.

(2017), where they set a lower bound on the Bitcoin account to account for Hierarchical Deterministic wallets as described in Section 3. They report that in 2018, Bitcoin had a Gini value of 0.65, where they set the minimum threshold to 185 BTC per account. This Gini value suggests that wealth in bitcoin is highly centralized when compared to real economies where, according to the World Bank (World Bank, 0000), the highest reported Gini value is 0.63.

According to Srinivasan (2017), Ethereum demonstrates a similar trend of significant centralization with a Gini value of 0.76 with a minimum threshold of 2477 ETH per account. This suggested trend is in line with the report by Huobi Blockchain (0000), where they claim Ethereum to be more centralized in terms of wealth distribution.

In line with the investigation of Gupta and Gupta (2017) and Srinivasan (2017), we parsed daily transactions for both Bitcoin and Ethereum starting from the genesis block. Our results are in accordance with previous reports of significant wealth inequality in Bitcoin and Ethereum. However, we also report that there is a downwards trend in Bitcoin where the Gini has been steadily decreasing with the current Gini value of 0.56. The current Gini value for Ethereum is 0.64, which is lower than the Figure suggested by Srinivasan (2017), implying a downwards trend. We have visualized these results in Fig. 12.

5.5. Operational layer

In Pustišek, Umek, and Kos (2019), the authors report that the Bitcoin full node requires 204 GB storage space. This storage requirement is slightly lower than the 385 GB required by Ethereum for a full node (Afanasev, Krylova, Shorokhov, Fedosov, & Sidorenko, 2018). Pustišek et al. (2019) also reports that the storage growth rate is about 0.1–0.5 GB per day. Our review was unable to identify any longitudinal studies that observe the growth in storage requirements over a long time. To account for this, we collected the storage growth rate data for both Ethereum and Bitcoin for a period of a month by hosting full nodes. We report that Bitcoin's storage growth rate is on average 0.50 GB per week, whereas Ethereun tends to grow at a faster rate with the weekly growth rate of 0.68 GB.¹⁹

As reported in Section 4, numerous studies identify specialized equipment concentration as a cause of centralization. Despite the significant attention to this issue, our review suggests that there are no proposed measurement techniques.

5.6. Application layer

Reference Client Concentration: According to Azouvi et al. (2018), a single author wrote about 30% of all files in the bitcoin reference implementation.²⁰ This is significantly higher in Ethereum, where an individual author wrote 55% of all files. They also analyze the comments on the GitHub pages of Bitcoin and Ethereum reference clients. They report that only eight people contributed to half of all comments representing 0.3% of all commenters. This concentration in comments is also observable in Ethereum, where 0.6% commenters contributed to 50% of comments.

We analyzed the core clients for Bitcoin and Ethereum ($Geth^{21}$) to observe the current state of centralization in the development. We report that Ethereum has a higher contribution of single developers than Bitcoin. In Ethereum, a single author has contributed to over 40% of all commits, whereas in Bitcoin, a single author wrote 25.11% of all commits. These observations are in line with Azouvi et al.'s (2018) results. We have reported the top 5 contributors to Ethereum and Github core clients in Fig. 13.

Exchange Concentration: Intermediary services such as Exchanges that also act as central key stores for Bitcoin have been suggested as a centralization causing factor by Böhme et al. (2015). A prominent example of the harm caused by exchange

²⁰ Azouvi et al. (2018) propose using the Satoshi Index to measure centralization in client development. However, the specific values of Satoshi Index for Bitcoin and Ethereum are not available.

¹⁹ Both the full nodes were hosted from DEC-01-2020 to JAN-01-2021, the daily growth reports are available at www.github.com/ashishrsai/centralization.

²¹ https://github.com/ethereum/go-ethereum.



Fig. 13. Github contributions of top 5 authors for Bitcoin and Ethereum.

concentration is the collapse of Mt. Gox in 2014 (Abrams et al., 2014). In 2014, Mt. Gox was the leading exchange for Bitcoin, and its closure resulted in a total loss of \$450 Million. Böhme et al. (2015) reports that the concentration of exchanges was still high in 2015 when the seven largest exchanges served more than 95% of all bitcoin trades.

An empirical analysis conducted by Böhme et al. (2015) reported that out of 40 Bitcoin exchanges examined, 18 had closed, wiping out customers' account balance as they stored the private keys of customers. They argue that these exchanges operate as the de facto centralized authorities in the Bitcoin network.

As for Ethereum, we report that there are no studies that explicitly report on the behavior of exchanges for Ethereum. However, as suggested by Kim and Lee (2018), most of the Bitcoin exchanges also exchange multiple other cryptocurrencies, including Ether.

To account for the lack of empirical data for Ethereum, we measure the centrality of exchanges by observing the flow of cryptocurrencies between addresses on the blockchain as discussed in Section 4.6.2. We measure centrality for both Bitcoin and Ethereum for a period of a week.²² We report that the top 7 exchanges on Bitcoin processed over 97.24% of all trades. This ratio was significantly lower for Ethereum, with the top 7 exchanges contributing to 28.27% of all transactions.

As discussed earlier, based on our systematic review, we conclude that there is no suggestion regarding a measurement technique to capture wallet based centralization.

So, in terms of Bitcoin, the main centralization threats are at the Network, Consensus, and Application layers. Specifically, the centralization aspects of the Network layer: geographic distribution, bandwidth, and routing are vulnerabilities for bitcoin in that they allow the specific threats of geopolitical manipulation of the network, high resource requirement for participation, and possibility of network attacks. These threats for bitcoin are augmented by the high concentration of consensus power to centralized mining pools and application layer operations such as exchanges and wallets.

Ethereum also shares the issues of centralization on the application layer as they lead to reliance on centralized entities such as exchanges and wallets for participation in the network. Other significant centralization threats for Ethereum include the Governance, Consensus, and Incentive layers. Especially the centralization aspects of the Governance and Incentive layers may induce vulnerabilities for Ethereum in that they allow unilateral decision making on the governance layer and high wealth concentration on the incentive layer.

6. Discussion

In this first in-depth investigation of the centralization of public blockchain solutions, we conducted a systematic review of existing literature to produce an initial taxonomy of centralization. We then refined this initial taxonomy through expert interviews. We provide an overview of centralization in different aspects of the blockchain. We examine different means of measuring centralization, also pointing out the absence of measurement techniques in these research studies. This initial taxonomy provides a framework for a more systematic discussion around the centralization of major blockchain systems. The following section discusses the findings of our survey.

6.1. Non binary nature of centralization

We observe that decentralization in the public Blockchain literature is a loosely-defined term that can take many shapes and forms. We also observe that most of the non-decentralization-specific articles reviewed treat decentralization as a binary construct. That is: a blockchain instance is either centralized or decentralized. However, based on our taxonomy, we define centralization of

²² From DEC-25-2020 12:00:00 PM +UTC until JAN-01-2021 12:00:00 PM +UTC.

public Blockchains as the process by which one or more architectural dimensions (aspects) of the Blockchain are restrictive to the majority of participants by direct or indirect economic, social, or technical constraints and so argue that centralization is not suited to binary classification.

This latter observation aligns with expert interviews, where 60% of participants preferred a spectrum of values for centralization rather than the conventional binary notion. However, the interviewees also acknowledged that the complexity of a more granular definition might dilute the meaning to non-experts in the blockchain domain. For example, I_5 said: "I am an engineer, so I prefer precision and a multidimensional model, but I know when you are presenting to business people, a single score might be what they are looking for".

This survey presents a novel, initial taxonomy to address this dilution concern and allows for structured discussion on centralization. The following text discusses the key findings of the taxonomy.

Consensus power concentration was the most recognized form of blockchain centralization by both the literature and experts interviewed. We reason that this wide recognition is due to the dependence of significant security threats such as the Double Spending (Karame et al., 2012) and Selfish mining (Sapirshtein et al., 2016) attacks on the consensus power concentration. The practical implication of this centralization is the heavy the impact of mining pools when operating a profitable mining operation. The dominance of mining pools is observable in both Ethereum and Bitcoin. In Bitcoin the top 4 mining pools control over 53% of the hashing power, whereas in Ethereum the top 3 mining pools control over 61% of the hashing power (See Table 10).

A high concentration of consensus power can induce an arm's race to attain the most efficient hardware (Sai et al., 2019a). Our survey reports that this race often results in specialized proprietary hardware. The practical implication of this type of hardware concentration is an indirect limitation to participation as only efficient, and often proprietary hardware can result in a profitable operation. To remedy this situation, studies such as Cho et al. Hyungmin (2018) (Cho, 2018), have proposed using a consensus algorithm that is memory heavy, for which specialized hardware design is inefficient.

Surprisingly on a similar operational constraint, the Storage growth rate was less widely recognized to contribute to centralization. However, I_{10} raised an interesting issue on the ever-increasing append-only nature of Blockchain that may result in consistent growth in storage requirements. As reported in Table 10, the current growth rate for Bitcoin is around 0.1 to 0.5 GB per day. The practical implication of this increased storage requirement is the inability of conventional computing devices to serve as nodes in the blockchain (Guo et al., 2019). Guo et al. (2019) propose a storage optimization scheme based on the redundant residual number system that can reduce the storage requirement. We suggest that a further investigation into storage optimization in public Blockchain is warranted.

Another unexpected finding of our survey was that 50% of the interviewees accepted node discovery protocol control as a threat to decentralization, despite only one research article reporting on the issue. We reason that this may be due to the practical implications of setting up a new node such as the potential delay in network connection for new nodes due to high traffic through DNS nodes. This type of delay is often not accounted for in network simulation tools such as NS3, employed by studies such as (Gervais et al., 2016; Sai et al., 2019a).

Contrary to the previous example, routing and bandwidth centralization in the network was not widely recognized by the interviewees. One potential explanation could be the experimental nature of the measurement associated with the routing and bandwidth centralization. Despite these being recognized as issues, both the bandwidth and routing do not cause operational issues to most participants at present.

Another network-oriented centralization concern widely recognized by both the literature and interviewees is the geographic distribution of the nodes. Our findings suggest that the Ethereum network is more geographically spread out than Bitcoin. We reason that this is due to the possibility of using conventional hardware such as GPUs to participate in Ethereum. Despite the recognition, our literature review did not identify potential strategies to address this centralization. We suggest that strategies to limit geographic concentration should be investigated.

The lack of mitigation techniques is also persistent in the application layer aspects. The wallet and exchange centralization have been reported on by the literature and also recognized as centralization issues by expert interviews. As reasoned earlier, the centralized store of cryptocurrencies may give an advantage to the exchange or wallet operator. This advantage is often in the form of wealth concentration and can be observed in the centralization of Bitcoin exchange platforms, where only seven exchanges were reported to serve more than 95% of all trades.

Interviewees and literature also agree on the implication of wealth concentration on the decentralization. Surprisingly, despite the apparent issue of a *"Rich getting Richer"* effect in Proof-of-Stake cryptocurrencies (Fanti et al., 2019), most of the reported literature focused on the wealth concentration in Proof-of-Work. We suggest that the issue of wealth concentration be investigated in the context of Proof-of-Stake cryptocurrencies.

Another factor that may result in a "*Rich getting Richer*" effect is the distribution of wealth at the very start of the Blockchain captured by owner control in our taxonomy. The issue of owner control is also associated with how the Blockchain is governed. Governance centralization in Blockchain is widely recognized by both the literature and interviewees. Interestingly, Wang et al. (2017) argue for some centralization in the governance to facilitate quick response to security threats. We expand on this line of reasoning in the following subsection.

6.2. Aspect based measurement of implications of centralization

As pointed out earlier, not all aspects of our taxonomy are an equal contributor to the overall centralization of the blockchain. This was also substantiated by six interviewees agreeing that a combined value of centralization for the overall blockchain would

not be meaningful. For example, storage constraint oriented centralization may be an issue in Ethereum due to the requirement to store smart contracts. In contrast, this may not be a significant issue for Bitcoin as only transactions drive the storage requirements. We expand on this category-based significance reasoning that not all centralization is necessarily equally bad for the network:

The governance layer based centralization argument presented by Gervais et al. (2014) assumes that concentrating decision making power to a select few is bad for blockchain. However, we question this argument, as true decentralization is an impossibility in real world scenarios (Kwon et al., 2019; Szabo, 1970). The concentration in decision making had also proven to be useful in instances of network attacks when a prompt response was mandated (Wang et al., 2017). Delegation of controlling power during the cases of security bugs or attacks may have proven to be detrimental to the network. Despite the lack of decentralization in governance, it may be to the overall benefit of the network. We present this as a potential future research avenue to explore the most suitable governance structure for decentralized systems.

We also argue that the results obtained by Azouvi et al. (2018) regarding the centralization in source code development for core client implementation may not necessarily be bad. It may just be the case that only a handful of developers have an in-depth understanding of the source code to make useful contributions to the system. This reasoning of limited expertise feeds into the argument against the decentralization of the improvement protocol. As pointed out by Azouvi et al. (2018), the vast majority of the Ethereum Improvement Protocol recommendations originated from a single developer, Vitalik Buterin. We reason that this may be due to the quality of suggestions proposed by Vitalik.

These arguments in favor of some centralization are an example of the complex nature of decentralization in distributed systems. We propose that the significance of each aspect of centralization be determined based on the empirical evidence specific to each blockchain instance.

7. Conclusion

In this paper, we conduct a systematic literature review to provide a summary of the research done on the centralization aspect of blockchain. We structure our findings in a novel initial taxonomy of centralization. This taxonomy is then refined and validated through expert interviews.

Given the significant growth in the application of blockchain technology in information systems (Chen et al., 2020; Khalid et al., 2021; Li et al., 2020; Putz et al., 2021), it becomes imperative to understand centralization's socio-technical nature in the blockchain. This refined understanding of centralization helps us better understand the security and performance implications associated with adopting a blockchain-based decentralization approach in existing information systems. Another important aspect associated with the adoption of blockchain-based decentralization is the management of the decentralized system. This taxonomy report on the issues associated with the governance of a decentralized system and its potential implications.

7.1. Contribution

Decentralized blockchain solutions provide a means of monetary asset transfer without a trusted third party; this is attained through the delegation of the validation power to all participants of the system rather than the administrator. This delegation of control is often referred to as the original contribution of blockchain systems (Bonneau et al., 2015). Based on previous studies, Cong et al. (2019), Gencer et al. (2018), Gervais et al. (2014) and Sai et al. (2019a), we reason that the preconceived notion that blockchains are inherently decentralized may not hold in the present situation and that raises the potential of severe issues for blockchain instances. Due to the lack of an objective measure of centralization, it becomes impractical to discuss improvement in terms of centralization.

Centralization is a challenging variable to research, in part because of the multiple definitions and measures of centralization applicable in blockchain and, to date, the implicit nature of several of those aspects and the lack of an encompassing framework. We report on these myriads of definitions, conceptualizations, and dimensions used to describe this concept by segmenting them based on a generic architecture proposed by Zhang et al. (2019). Our study contributes to the existing body of knowledge by systematically surveying and synthesizing the blockchain literature, reporting on the adverse impact of centralization such as security threats, as well as identifying research gaps such as the lack of Ethereum specific research on centralization.

With this systematic review, we provide the reader with an overview of various forms of centralization in Blockchain resulting in an initial taxonomy. This taxonomy also contains numerous existing measurement techniques used to measure centralization. It may help researchers evaluate the centralization of a blockchain instance, but will also allow researchers add more aspects of centralization as they become known, providing them with a vocabulary of centralization that will allow them address the issues that arise.

We have also reported on the platform-specific findings for the two most prominently used blockchain-based cryptocurrencies: Bitcoin and Ethereum. We report that both Bitcoin and Ethereum have similar centralization issues with regards to reference client implementation, decentralized protocol development, and exchanges. However, in terms of wealth concentration, Ethereum is more centralized than Bitcoin, primarily due to high owner control. This trend continues with consensus power concentration, where Ethereum is reported to be more centralized than Bitcoin. Ethereum nodes, however, are geographically more spread out than Bitcoin, resulting in a low geographic concentration when compared to Bitcoin.

We also discuss that centralization on all aspects is not necessarily adverse for the blockchain by expanding the argument in favor of some centralization by Wang et al. (2017). We suggest that the unpropitious impact of centralization be measured on each aspect based on empirical evidence. This aspect-specific investigation may assist the move from the binary notion of decentralization to a multidimensional scale encompassing adequate measurement and control where necessary.

7.2. Threats to validity

As decentralization is fundamental to a public blockchain, the term is frequently used in the title and abstract of articles relating to public blockchains. To not omit any relevant articles, we kept the search queries generic by including any article that includes the term *"Blockchain"* and *"Decentralization"* along with suggested alternate words in Section 3. We acknowledge that despite the broad terms used, we may have missed relevant articles not present, or with different phrasing, on these leading search repositories. These missed articles may include *"gray literature"*, which is of significant importance in the blockchain research domain (Casino, Dasaklis, & Patsakis, 2019). To overcome this limitation, we included Google Scholar in our search process. However, as reported earlier, the Google Scholar search was limited to the top 1000 entries, even though the relevant articles dropped off significantly after the top four hundred returned articles.

The literature review may also be limited due to the strict inclusion and exclusion criteria for the title and abstract filtering. We reason that these strict criteria are warranted due to a large number of articles retrieved by the search queries (3574 non-duplicate entries). To overcome this limitation, we employed a two-step filtration by reviewing both the title and abstract. We also performed cross-validation of the filtration process by the independent review of the articles by two authors. This cross-validation process resulted in Cohen's Kappa value of 0.84, which is considered an almost perfect agreement. We repeated a similar cross-validation process for the full-text filtration.

The review process aimed to extract factors from all shortlisted articles despite their core focus. As the study of centralization in public blockchain is still in the early stage, we included articles where the core focus was not centralization. This inclusion may have limited the quality of shortlisted articles, as observed by the exclusion of 148 articles after full-text filtration. To overcome this limitation, we performed a quality review of all 212 shortlisted articles and shortlisted a final set of 89 articles.

To further evaluate the literature-review findings, we interviewed ten experts. The recruitment process was based on the prominence of authors in the bibliographic map generated by Ramona et al. (2019). As with any other qualitative research method, interviews have several limitations, as pointed out by Opdenakker (2006). In addressing them, we adhere to the validity dimensions put forth by Maxwell (1992) for qualitative studies. The first validity threat is the descriptive validity of the data obtained through interviews. To limit this, we transcribed the audio-captured interview in verbatim form. However, in the interviews that relied on contemporaneous notes, it is possible that the interviewer may have missed some observations. The second threat to validity is the interpretive validity of the interviews. To address this, we used open-ended questions and restricted the questions strictly to the research questions presented in Section 3.2. We also coded the interviews based on the terms used by the interviewees rather than an interpretation. The transcripts and notes were individually checked by researchers from the author list. The interviewees were also given back the interview transcripts and notes for validation.

7.3. Future work

Having provided a comprehensive overview of centralization in public blockchain, a case study focused on individual cryptocurrencies, and blockchain implementations would complement our study. This case study could include an in-depth centralization review of, for example, Bitcoin, Ethereum, and Libra (Pilkington, 2019).

The taxonomy developed by our study can also be expanded to provide an objective measure of centralization for blockchain instances, as a whole, to facilitate comparison. This objective measure may prove to be useful for the evaluation of centralization from a novice user, or governance perceptive. Four of our ten interviewees stated that they would prefer a single score to measure centralization objectively, and thought it would assist end-users and nonspecialist researchers.

We also hope to develop different flavors of this initial taxonomy that are specific to implementation details. For instance, the presented taxonomy is generic and does not consider consensus specific issues such as Stake bleeding (Gaži, Kiayias, & Russell, 2018). It also omits the consideration of source code dependencies in Smart Contracts. In future, we intend to statistically examine the source code of smart contracts to observe if a handful of libraries dominate the smart contracts in Ethereum.

The work presented here only examines the already identified factors that may lead to centralization and does not analyze the existence of other novel forms of centralization. As a part of future work, we will consider a thorough review of one of the reference blockchain implementations to identify factors that may also contribute to centralization directly or indirectly.

We also aim to review existing literature to identify potential solutions to the centralization avenues suggested by our review. These solutions may facilitate integrating centralization considerations during the development of public blockchains.

CRediT authorship contribution statement

Ashish Rajendra Sai: Conceptualization, Methodology, Data curation, Writing - original draft. Jim Buckley: Writing - review & editing, Supervision, Project administration, Validation. Brian Fitzgerald: Writing - review & editing, Supervision. Andrew Le Gear: Visualization, Investigation, Supervision, Validation, Writing - review & editing.

Acknowledgments

This work was supported, in part, by Science Foundation Ireland grant 13/RC/2094 and co-funded under the European Regional Development Fund through the Southern Eastern Regional Operational Programme to Lero - the Irish Software Research Centre (www.lero.ie).

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.ipm.2021.102584.

References

Abrams, Rachel, Goldstein, Matthew, & Tabuchi, Hiroko (2014). Erosion of faith was death knell for mt. Gox. New York Times.

- Afanasev, Maxim Ya, Krylova, Anastasiya A., Shorokhov, Sergey A., Fedosov, Yuri V., & Sidorenko, Anastasiia S. (2018). A design of cyber-physical production system prototype based on an ethereum private network. In 2018 22nd conference of open innovations association (FRUCT) (pp. 3–11). IEEE.
- Akram, Shaik V., Malik, Praveen K., Singh, Rajesh, Anita, Gehlot, & Tanwar, Sudeep (2020). Adoption of blockchain technology in various realms: Opportunities and challenges. Security and Privacy, 3(5), Article e109.
- Alzahrani, Naif, & Bulusu, Nirupama (2018). Towards true decentralization: A blockchain consensus protocol based on game theory and randomness. In International conference on decision and game theory for security (pp. 465–485). Springer.
- Anceaume, Emmanuelle, Lajoie-Mazenc, Thibaut, Ludinard, Romaric, & Sericola, Bruno (2016). Safety analysis of bitcoin improvement proposals. In 2016 IEEE 15th international symposium on network computing and applications (NCA) (pp. 318–325). IEEE.
- Androulaki, Elli, Barger, Artem, Bortnikov, Vita, Cachin, Christian, Christidis, Konstantinos, De Caro, Angelo, et al. (2018). Hyperledger fabric: a distributed operating system for permissioned blockchains. In Proceedings of the thirteenth eurosys conference (p. 30). ACM.
- Antonopoulos, Andreas M. (2017). Mastering bitcoin: Programming the open blockchain. " O'Reilly Media, Inc.".
- Antonopoulos, Andreas M., & Wood, Gavin (2018). Mastering ethereum: building smart contracts and dapps. O'reilly Media.
- Apostolaki, Maria, Zohar, Aviv, & Vanbever, Laurent (2017). Hijacking bitcoin: Routing attacks on cryptocurrencies. In 2017 IEEE symposium on security and privacy (SP) (pp. 375–392). IEEE.
- Atzori, Marcella (2015). Blockchain technology and decentralized governance: Is the state still necessary? Available at SSRN 2709713.
- Azouvi, Sarah, Maller, Mary, & Meiklejohn, Sarah (2018). Egalitarian society or benevolent dictatorship: The state of cryptocurrency governance. In International conference on financial cryptography and data security (pp. 127–143). Springer.
- Bai, Qianlan, Zhang, Chao, Xu, Yuedong, Chen, Xiaowei, & Wang, Xin (2020). Evolution of ethereum: A temporal graph perspective. arXiv preprint arXiv: 2001.05251.
- Bailey, Kenneth D. (1994). Typologies and taxonomies: An introduction to classification techniques, Number 102. Sage.
- Baliga, Arati (2017). Understanding blockchain consensus models. Persistent, 2017(4), 1-14.
- Baniata, Hamza, Anaqreh, Ahmad, & Kertesz, Attila (2021). Pf-BTS: A privacy-aware fog-enhanced blockchain-assisted task scheduling. Information Processing & Management, 58(1), Article 102393.
- Beck, Roman, Avital, Michel, Rossi, Matti, & Thatcher, Jason Bennett (2017). Blockchain technology in business and information systems research. Springer.
- Beck, Roman, Müller-Bloch, Christoph, & King, John Leslie (2018). Governance in the blockchain economy: A framework and research agenda. Journal of the Association for Information Systems, 19(10), 1020–1034.
- Beikverdi, Alireza, & Song, JooSeok (2015). Trend of centralization in bitcoin's distributed network. In 2015 IEEE/ACIS 16th international conference on software engineering, artificial intelligence, networking and parallel/distributed computing (SNPD) (pp. 1–6). IEEE.
- Berdik, David, Otoum, Safa, Schmidt, Nikolas, Porter, Dylan, & Jararweh, Yaser (2020). A survey on blockchain for information systems management and security. Information Processing & Management, 58(1), Article 102397.
- Bitcoin (2019). Bitcoin improvement proposals github repository. URL=https://github.com/bitcoin/bips.
- Blockchain luxembourg s. a (2019). Blocks mined. Bitcoin Block Explorer and Currency Statistics, URL=https://www.blockchain.com/btc/blocks.
- Böhme, Rainer, Christin, Nicolas, Edelman, Benjamin, & Moore, Tyler (2015). Bitcoin: Economics, technology, and governance. Journal of Economic Perspectives, 29(2), 213–238.
- Bohr, Jeremiah, & Bashir, Masooda (2014). Who uses bitcoin? an exploration of the bitcoin community. In 2014 twelfth annual international conference on privacy, security and trust (pp. 94–101). IEEE.
- Bonneau, Joseph, Miller, Andrew, Clark, Jeremy, Narayanan, Arvind, Kroll, Joshua A., & Felten, Edward W. (2015). Sok: Research perspectives and challenges for bitcoin and cryptocurrencies. In 2015 IEEE symposium on security and privacy (pp. 104–121). IEEE.
- Borge, Maria, Kokoris-Kogias, Eleftherios, Jovanovic, Philipp, Gasser, Linus, Gailly, Nicolas, & Ford, Bryan (2017). Proof-of-personhood: Redemocratizing permissionless cryptocurrencies. In 2017 IEEE European symposium on security and privacy workshops (EuroS&PW) (pp. 23–26). IEEE.
- Bradbury, Danny (2013). The problem with bitcoin. Computer Fraud & Security, 2013(11), 5-8.
- Briscoe, Neil (2000). Understanding the OSI 7-layer model. PC Network Advisor, 120(2).
- Bruschi, Francesco, Rana, Vincenzo, Gentile, Lorenzo, & Sciuto, Donatella (2019). Mine with it or sell it: the superhashing power dilemma. ACM SIGMETRICS Performance Evaluation Review, 46(3), 127–130.
- Buckley, Jim, & Exton, Christopher (2003). Bloom's taxonomy: A framework for assessing programmers' knowledge of software systems. In 11th IEEE international workshop on program comprehension, 2003. (pp. 165–174). IEEE.
- Buterin, Vitalik, et al. (2013). Ethereum white paper. GitHub Repository, 22-23.
- Caccioli, Fabio, Livan, Giacomo, & Aste, Tomaso (2016). Scalability and egalitarianism in peer-to-peer networks. In *Banking beyond banks and money* (pp. 197–212). Springer.
- Caffyn, Grace (2015). What is the bitcoin block size debate and why does it matter. URL: http://www.coindesk.com/what-is-the-bitcoin-block-size-debate-and-why-does-it-matter/(visited on 27/11/2015).
- Casino, Fran, Dasaklis, Thomas K., & Patsakis, Constantinos (2019). A systematic literature review of blockchain-based applications: current status, classification and open issues. *Telematics and Informatics*, 36, 55–81.
- Chen, Qian, Srivastava, Gautam, Parizi, Reza M., Aloqaily, Moayad, & Ridhawi, Ismaeel Al (2020). An incentive-aware blockchain-based solution for internet of fake media things. Information Processing & Management, 57(6), Article 102370.
- Chen, Lin, Xu, Lei, Shah, Nolan, Gao, Zhimin, Lu, Yang, & Shi, Weidong (2017). On security analysis of proof-of-elapsed-time (poet). In International symposium on stabilization, safety, and security of distributed systems (pp. 282–297). Springer.
- Chesterman, Xavier (2018). The P2pool mining pool (PhD thesis), Ghent University.
- Chia, Vincent, Hartel, Pieter, Hum, Qingze, Ma, Sebastian, Piliouras, Georgios, Reijsbergen, Daniel, et al. (2018). Rethinking blockchain security: Position paper. arXiv preprint arXiv:1806.04358.
- Cho, Hyungmin (2018). ASIC-Resistance of multi-hash proof-of-work mechanisms for blockchain consensus protocols. *IEEE Access*, 6, 66210–66222. Chohan, Usman W. (2019). Cryptocurrencies and inequality. Notes on the 21st Century (CBRI).
- Chu, Dennis (2018). Broker-dealers for virtual currency: Regulating cryptocurrency wallets and exchanges. Columbia Law Review, 118(8), 2323-2360. Cong, Lin William, He, Zhiguo, & Li, Jiasun (2019). Decentralized mining in centralized pools: Technical report, National Bureau of Economic Research.
- Cong, Lin Winnan, He, Zinguo, & Li, Jiasun (2019). Decentratizea mining in centratizea pools: Technical report, National Bureau of Economic Research.
- Conti, Mauro, Kumar, E. Sandeep, Lal, Chhagan, & Ruj, Sushmita (2018). A survey on security and privacy issues of bitcoin. *IEEE Communications Surveys & Tutorials*, 20(4), 3416–3452.
- Cryptoslate (2018). Ethereum network under assault: Gas price manipulation may indicate covert EOS attack [interview].

- Dai, Mingjun, Zhang, Shengli, Wang, Hui, & Jin, Shi (2018). A low storage room requirement framework for distributed ledger in blockchain. *IEEE Access*, 6, 22970–22975.
- Davidson, Sinclair, De Filippi, Primavera, & Potts, Jason (2016). Economics of blockchain. Available at SSRN 2744751.
- De Domenico, Manlio, & Baronchelli, Andrea (2019). The fragility of decentralised trustless socio-technical systems. EPJ Data Science, 8(1), 2.
- De Filippi, Primavera, & Loveluck, Benjamin (2016). The invisible politics of bitcoin: governance crisis of a decentralized infrastructure. Internet Policy Review, 5(4).
- Dietrich, Sven, Long, Neil, & Dittrich, David (2000). Analyzing distributed denial of service tools: The shaft case.. In LISA (pp. 329-339).
- Dorfman, Robert (1979), A formula for the gini coefficient. The Review of Economics and Statistics, 146-149,
- Dwivedi, Ashutosh Dhar, Srivastava, Gautam, Dhar, Shalini, & Singh, Rajani (2019). A decentralized privacy-preserving healthcare blockchain for IoT. Sensors, 19(2), 326.
- Ekblaw, Ariel, Barabas, Chelsea, Harvey-Buschel, Jonathan, & Lippman, Andrew (2016). Bitcoin and the myth of decentralization: Socio-technical proposals for restoring network integrity. In 2016 IEEE 1st international workshops on foundations and applications of self* systems (FAS* W) (pp. 18–23). IEEE.
- Esposito, Christian, Ficco, Massimo, & Gupta, Brij Bhooshan (2021). Blockchain-based authentication and authorization for smart city applications. Information Processing & Management, 58(2), Article 102468.
- Etherscan (2019a). Ethereum transaction information for the gensis block. URL=https://etherscan.io/txs?block=0.
- Etherscan (2019b). Total ether supply and market capitalization. URL=https://etherscan.io/stat/supply.
- Fanti, Giulia, Kogan, Leonid, Oh, Sewoong, Ruan, Kathleen, Viswanath, Pramod, & Wang, Gerui (2019). Compounding of wealth in proof-of-stake cryptocurrencies. In International conference on financial cryptography and data security (pp. 42–61). Springer.
- Feld, Sebastian, Schönfeld, Mirco, & Werner, Martin (2014). Analyzing the deployment of bitcoin's P2p network under an AS-level perspective. Procedia Computer Science. 32, 1121–1126.
- Fleiss, Joseph L., & Cohen, Jacob (1973). The equivalence of weighted kappa and the intraclass correlation coefficient as measures of reliability. Educational and Psychological Measurement, 33(3), 613–619.
- Galster, Matthias, Weyns, Danny, Tofan, Dan, Michalik, Bartosz, & Avgeriou, Paris (2013). Variability in software systems—a systematic literature review. IEEE Transactions on Software Engineering, 40(3), 282–306.
- Garay, Juan, Kiayias, Aggelos, & Leonardos, Nikos (2015). The bitcoin backbone protocol: Analysis and applications. In Annual international conference on the theory and applications of cryptographic techniques (pp. 281-310). Springer.
- Gastwirth, Joseph L. (1971). A general definition of the lorenz curve. Econometrica, 1037-1039.
- Gaži, Peter, Kiayias, Aggelos, & Russell, Alexander (2018). Stake-bleeding attacks on proof-of-stake blockchains. In 2018 crypto valley conference on blockchain technology (CVCBT) (pp. 85–92). IEEE.
- Gencer, Adem Efe, Basu, Soumya, Eyal, Ittay, Van Renesse, Robbert, & Sirer, Emin Gün (2018). Decentralization in bitcoin and ethereum networks. arXiv preprint arXiv:1801.03998.
- Gervais, Arthur, Karame, Ghassan O., Capkun, Vedran, & Capkun, Srdjan (2014). Is bitcoin a decentralized currency? IEEE Security & Privacy, 12(3), 54-60.
- Gervais, Arthur, Karame, Ghassan O., Wüst, Karl, Glykantzis, Vasileios, Ritzdorf, Hubert, & Capkun, Srdjan (2016). On the security and performance of proof of work blockchains. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security* (pp. 3–16).
- Gini, Corrado (1921). Measurement of inequality of incomes. The Economic Journal, 31(121), 124-126.
- Great Britain. Government Office for Science (2016). Distributed ledger technology: Beyond block chain. Government Office for Science.
- Guegan, Dominique (2017). Public blockchain versus private blockhain. Documents de travail du Centre d'Economie de la Sorbonne 2017.20 ISSN : 1955-611X.
- Guerra García, José Manuel, Espinosa Torre, Free, & García Gómez, José Carlos (2008). Trends in taxonomy today: an overview about the main topics in taxonomy. Zoológica Baetica, 19, 15–49.
- Guo, Zhaohui, Gao, Zhen, Mei, Haojuan, Zhao, Ming, & Yang, Jinsheng (2019). Design and optimization for storage mechanism of the public blockchain based on redundant residual number system. *IEEE Access*, 7, 98546–98554.
- Guo, Ye, & Liang, Chen (2016). Blockchain application and outlook in the banking industry. Financial Innovation, 2(1), 24.
- Gupta, Manas, & Gupta, Parth (2017). Gini coefficient based wealth distribution in the bitcoin network: A case study. In International conference on computing, analytics and networks (pp. 192–202). Springer.
- Gutoski, Gus, & Stebila, Douglas (2015). Hierarchical deterministic bitcoin wallets that tolerate key leakage. In International conference on financial cryptography and data security (pp. 497–504). Springer.
- Halpin, Harry, & Piekarska, Marta (2017). Introduction to security and privacy on the blockchain. In 2017 IEEE European symposium on security and privacy workshops (EuroS&PW) (pp. 1–3). IEEE.
- Hardin, Taylor, & Kotz, David (2021). Amanuensis: Information provenance for health-data systems. *Information Processing & Management*, 58(2), Article 102460. He, Pu, Yu, Ge, Zhang, Y. F., & Bao, Y. B. (2017). Survey on blockchain technology and its application prospect. *Computer Science*, 44(4), 1–7.
- Heilman, Ethan, Kendler, Alison, Zohar, Aviv, & Goldberg, Sharon (2015). Eclipse attacks on bitcoin's peer-to-peer network. In 24th {USENIX} security symposium ({USENIX} security 15) (pp. 129–144).
- Hileman, Garrick, & Rauchs, Michel (2017). Global cryptocurrency benchmarking study. Cambridge Centre for Alternative Finance, 33.
- Hu, Teng, Liu, Xiaolei, Chen, Ting, Zhang, Xiaosong, Huang, Xiaoming, Niu, Weina, et al. (2021). Transaction-based classification and detection approach for ethereum smart contract. Information Processing & Management, 58(2), Article 102462.
- Huobi Blockchain Big Data Weekly Insights. Vol. 8.
- Iansiti, Marco, & Lakhani, Karim R. (2017). The truth about blockchain. Harvard Business Review, 95(1), 118-127.
- Jin, Tong, Zhang, Xiang, Liu, Yirui, & Lei, Kai (2017). Blockndn: A bitcoin blockchain decentralized system over named data networking. In 2017 ninth international conference on ubiquitous and future networks (ICUFN) (pp. 75–80). IEEE.
- Jing, Nan, Liu, Qi, & Sugumaran, Vijayan (2021). A blockchain-based code copyright management system. Information Processing & Management, 58(3), Article 102518.
- Judmayer, Aljosha, Stifter, Nicholas, Krombholz, Katharina, & Weippl, Edgar (2017). Blocks and chains: Introduction to bitcoin, cryptocurrencies, and their consensus mechanisms. Synthesis Lectures on Information Security, Privacy, & Trust, 9(1), 1–123.
- Judmayer, Aljosha, Zamyatin, Alexei, Stifter, Nicholas, Voyiatzis, Artemios G., & Weippl, Edgar (2017). Merged mining: Curse or cure? In Data privacy management, cryptocurrencies and blockchain technology (pp. 316–333). Springer.
- Karame, Ghassan (2016). On the security and scalability of bitcoin's blockchain. In Proceedings of the 2016 ACM SIGSAC conference on computer and communications security (pp. 1861–1862). ACM.
- Karame, Ghassan O., & Androulaki, Elli (2016). Bitcoin and blockchain security. Artech House.
- Karame, Ghassan O., Androulaki, Elli, & Capkun, Srdjan (2012). Double-spending fast payments in bitcoin. In Proceedings of the 2012 ACM conference on computer and communications security (pp. 906–917).
- Khairuddin, Irni Eliana, & Sas, Corina (2019). An Exploration of Bitcoin mining practices: Miners' trust challenges and motivations. In Proceedings of the 2019 CHI conference on human factors in computing systems (1-13).
- Khalid, Adia, Iftikhar, Muhammad Sohaib, Almogren, Ahmad, Khalid, Rabiya, Afzal, Muhammad Khalil, & Javaid, Nadeem (2021). A blockchain based incentive provisioning scheme for traffic event validation and information storage in VANETs. *Information Processing & Management*, 58(2), Article 102464.

- Kiayias, Aggelos, Russell, Alexander, David, Bernardo, & Oliynykov, Roman (2017). Ouroboros: A provably secure proof-of-stake blockchain protocol. In Annual international cryptology conference (pp. 357-388). Springer.
- Kim, Chang Yeon, & Lee, Kyungho (2018). Risk management to cryptocurrency exchange and investors guidelines to prevent potential threats. In 2018 international conference on platform technology and service (PlatCon) (pp. 1–6). IEEE.
- Kim, Seoung Kyun, Ma, Zane, Murali, Siddharth, Mason, Joshua, Miller, Andrew, & Bailey, Michael (2018). Measuring ethereum network peers. In Proceedings of the internet measurement conference 2018 (pp. 91–104).
- Kim, Tae Wan, & Zetlin-Jones, Ariel (2019). The ethics of blockchain networks] the ethics of contentious hard forks in blockchain networks with fixed features. Frontiers in Blockchain, 2, 9.

Kitchenham, Barbara (2004). Procedures for performing systematic reviews.

- Kondor, Dániel, Pósfai, Márton, Csabai, István, & Vattay, Gábor (2014). Do the rich get richer? An empirical analysis of the bitcoin transaction network. PloS One, 9(2).
- Kwon, Yujin, Liu, Jian, Kim, Minjeong, Song, Dawn, & Kim, Yongdae (2019). Impossibility of full decentralization in permissionless blockchains. arXiv preprint arXiv:1905.05158.

Landis, J. Richard, & Koch, Gary G. (1977). The measurement of observer agreement for categorical data. biometrics, 159-174.

- Lee, Wei-Meng (2019). Using the web3. js APIs. In Beginning ethereum smart contracts programming (pp. 169-198). Springer.
- Lewenberg, Yoad, Bachrach, Yoram, Sompolinsky, Yonatan, Zohar, Aviv, & Rosenschein, Jeffrey S (2015). Bitcoin mining pools: A cooperative game theoretic analysis. In Proceedings of the 2015 international conference on autonomous agents and multiagent systems (pp. 919–927). International Foundation for Autonomous Agents and Multiagent Systems.
- Li, Xiaoqi, Jiang, Peng, Chen, Ting, Luo, Xiapu, & Wen, Qiaoyan (2017). A survey on the security of blockchain systems. Future Generation Computer Systems.
- Li, Jiaxing, Wu, Jigang, Jiang, Guiyuan, & Srikanthan, Thambipillai (2020). Blockchain-based public auditing for big data in cloud storage. Information Processing & Management, 57(6), Article 102382.
- Liao, Kevin, & Katz, Jonathan (2017). Incentivizing blockchain forks via whale transactions. In International conference on financial cryptography and data security (pp. 264–279). Springer.
- Lindley, John (1836). A Natural System of Botany, or, A systematic view of the organization, natural affinities, and geographical distribution, of the whole vegetable kingdom: together with the uses of the most important species in medicine, the arts, and rural or domestic economy. Longman, Rees, Orme, Brown, Green, and Longman.
- Malik, Vladimir (2016). The history and the future of bitcoin. Praha: Bankovní Institut Vysoká škola Praha.
- Marvin, Ian (2017). Decentralised? A Study of Concentration in the Bitcoin Network (PhD thesis), University of Cape Town.
- Mattila, Juri (2016). The blockchain phenomenon-the disruptive potential of distributed consensus architectures: Technical report, ETLA working papers.
- Maxwell, Joseph (1992). Understanding and validity in qualitative research. Harvard Educational Review, 62(3), 279-301.
- Meijer, David, & Ubacht, Jolien (2018). The governance of blockchain systems from an institutional perspective, a matter of trust or control? In Proceedings of the 19th annual international conference on digital government research: Governance in the data age (pp. 1–9).
- Miller, Andrew, Litton, James, Pachulski, Andrew, Gupta, Neal, Levin, Dave, Spring, Neil, et al. (2015). Discovering bitcoin's public topology and influential nodes. et al.
- Mingxiao, Du, Xiaofeng, Ma, Zhe, Zhang, Xiangwei, Wang, & Qijun, Chen (2017). A review on consensus algorithm of blockchain. In 2017 IEEE international conference on systems, man, and cybernetics (SMC) (pp. 2567–2572). IEEE.
- Nakamoto, Satoshi (2008). Bitcoin: A peer-to-peer electronic cash system.
- Neudecker, Till, & Hartenstein, Hannes (2018). Network layer aspects of permissionless blockchains. *IEEE Communications Surveys & Tutorials*, 21(1), 838–857. Nguyen, Cong T, Hoang, Dinh Thai, Nguyen, Diep N, Niyato, Dusit, Nguyen, Huynh Tuong, & Dutkiewicz, Eryk (2019). Proof-of-stake consensus mechanisms
- for future blockchain networks: fundamentals, applications and opportunities. IEEE Access, 7, 85727-85745.
- Nickerson, Robert C., Varshney, Upkar, & Muntermann, Jan (2013). A method for taxonomy development and its application in information systems. European Journal of Information Systems, 22(3), 336–359.
- Oberländer, Anna Maria, Lösser, Benedict, & Rau, Daniel (2019). Taxonomy research in information systems: A systematic assessment.

O'Dwyer, Karl J., & Malone, David (2014). Bitcoin mining and its energy footprint. IET.

Oham, Chuka, Michelin, Regio A., Jurdak, Raja, Kanhere, Salil S., & Jha, Sanjay (2021). B-FERL: Blockchain based framework for securing smart vehicles. Information Processing & Management, 58(1), Article 102426.

- Opdenakker, Raymond (2006). Advantages and disadvantages of four interview techniques in qualitative research. In Forum qualitative sozial forschung/forum: qualitative sozial research, Vol. 7.
- Panarello, Alfonso, Tapas, Nachiket, Merlino, Giovanni, Longo, Francesco, & Puliafito, Antonio (2018). Blockchain and iot integration: A systematic survey. Sensors, 18(8), 2575.
- Peck, Morgen E. (2017). Blockchain world-do you need a blockchain? This chart will tell you if the technology can solve your problem. *IEEE Spectrum*, 54(10), 38–60.
- Petersen, Kai, Feldt, Robert, Mujtaba, Shahid, & Mattsson, Michael (2008). Systematic mapping studies in software engineering. In 12th international conference on evaluation and assessment in software engineering (EASE) 12 (pp. 1–10).
- Pilkington, Marc (2019). The libra project: A transnational monetary dystopia-analysis of the disruption generated by the facebook-led stable coin. Available at SSRN 3434079.
- Pustišek, Matevž, Umek, Anton, & Kos, Andrej (2019). Approaching the communication constraints of ethereum-based decentralized applications. Sensors, 19(11), 2647.
- Putz, Benedikt, Dietz, Marietheres, Empl, Philip, & Pernul, Günther (2021). Ethertwin: Blockchain-based secure digital twin information management. Information Processing & Management, 58(1), Article 102425.

Raman, Ravi Kiran, & Varshney, Lav R. (2017). Dynamic distributed storage for scaling blockchains. arXiv preprint arXiv:1711.07617.

Ramona, Orăștean, Cristina, Mărginean Silvia, Raluca, Sava, et al. (2019). Bitcoin in the scientific literature-a bibliometric study. Studies in Business and Economics, 14(3), 160–174.

- Razzaq, Abdul, Wasala, Asanka, Exton, Chris, & Buckley, Jim (2018). The state of empirical evaluation in static feature location. ACM Transactions on Software Engineering and Methodology (TOSEM), 28(1), 1–58.
- Reddit (2019). R/monero blockchain size issue in future? reddit, URL=https://www.reddit.com/r/Monero/comments/9gymaf/blockchain_size_issue_in_future/. Roubini, Nouriel (2018a). The big blockchain lie. *Project Syndicate. Blog Post*, 15.
- Roubini, Nouriel (2018b). Blockchain isn't about democracy and decentralisation it's about greed | nouriel roubini.
- Sai, Ashish Rajendra, Buckley, Jim, & Le Gear, Andrew (2019a). Assessing the security implication of bitcoin exchange rates. Computers and Security.
- Sai, Ashish Rajendra, Buckley, Jim, & Le Gear, Andrew (2019b). Privacy and security analysis of cryptocurrency mobile applications. In 2019 fifth conference on mobile and secure services (mobisecserv) (pp. 1–6). IEEE.
- Sai, Ashish Rajendra, Le Gear, Andrew, & Buckley, Jim (2019). Centralization threat metric.
- Sapirshtein, Ayelet, Sompolinsky, Yonatan, & Zohar, Aviv (2016). Optimal selfish mining strategies in bitcoin. In International conference on financial cryptography and data security (pp. 515–532). Springer.

Saroiu, Stefan, Gummadi, P. Krishna, & Gribble, Steven D. (2001). Measurement study of peer-to-peer file sharing systems. In *Multimedia computing and networking* 2002, Vol. 4673 (pp. 156–170). International Society for Optics and Photonics.

Sayeed, Sarwar, & Marco-Gisbert, Hector (2019). Assessing blockchain consensus and security mechanisms against the 51% attack. Applied Sciences, 9(9), 1788. Sergio (2013). The well deserved fortune of satoshi nakamoto, bitcoin creator, visionary and genius.

Sim, Julius, & Wright, Chris C. (2005). The kappa statistic in reliability studies: use, interpretation, and sample size requirements. *Physical Therapy*, 85(3), 257–268.

Srinivasan, Balaji S. (2017). Quantifying decentralization. Medium.

StopAndDecrypt (2018). The ethereum-blockchain size has exceeded 1tb, and yes, it's an issue.

Szabo, Nick (1970). The dawn of trustworthy computing.

Tapsell, James, Akram, Raja Naeem, & Markantonakis, Konstantinos (2018). An evaluation of the security of the bitcoin peer-to-peer network. In 2018 IEEE international conference on internet of things (IThings) and IEEE green computing and communications (GreenCom) and IEEE cyber, physical and social computing (CPSCom) and IEEE smart data (SmartData) (pp. 1057–1062). IEEE.

Walport, Mark (2016). Distributed ledger technology: beyond blockchain. UK Government Office for Science: Technical report, Tech. Rep.

Wang, Wenbo, Hoang, Dinh Thai, Xiong, Zehui, Niyato, Dusit, Wang, Ping, Hu, Peizhao, et al. (2018). A survey on consensus mechanisms and mining management in blockchain networks. (pp. 1–33). arXiv preprint arXiv:1805.02707.

Wang, Sha, Vergne, Jean-Philippe J. P., & Hsieh, Ying-Ying (2017). The internal and external governance of blockchain-based organizations: Evidence from cryptocurrencies. In *Bitcoin and beyond* (pp. 48–68). Routledge.

Wirdum, Aaron van (2016). Rejecting today's hard fork, the ethereum classic project continues on the original chain: Here's why. Bitcoin Magazine, 20. Wolfson, Shael N. (2015). Bitcoin: the early market, Journal of Business & Economics Research (JBER), 13(4), 201–214.

Wood, Gavin, et al. (2019). Ethereum: A secure decentralised generalised transaction ledger. *Ethereum Project Yellow Paper*, 151(2014), 1–32.

World Bank, GINI index (World Bank estimate). URL=https://data.worldbank.org/indicator/si.pov.gini.

Wüst, Karl, & Gervais, Arthur (2018). Do you need a blockchain? In 2018 crypto valley conference on blockchain technology (CVCBT) (pp. 45-54). IEEE.

Xie, Junfeng, Tang, Helen, Huang, Tao, Yu, F Richard, Xie, Renchao, Liu, Jiang, et al. (2019). A survey of blockchain technology applied to smart cities: Research issues and challenges. *IEEE Communications Surveys & Tutorials*, 21(3), 2794–2830.

Xu, Xiaoqiong, Sun, Gang, Luo, Long, Cao, Huilong, Yu, Hongfang, & Vasilakos, Athanasios V. (2021). Latency performance modeling and analysis for hyperledger fabric blockchain network. *Information Processing & Management*, 58(1), Article 102436.

Xu, Xiwei, Weber, Ingo, Staples, Mark, Zhu, Liming, Bosch, Jan, Bass, Len, et al. (2017). A taxonomy of blockchain-based systems for architecture design. In 2017 IEEE international conference on software architecture (ICSA) (pp. 243–252). IEEE.

Yli-Huumo, Jesse, Ko, Deokyoon, Choi, Sujin, Park, Sooyong, & Smolander, Kari (2016). Where is current research on blockchain technology?—A systematic review. *PLoS One*, 11, Article e0163477.

Zhang, Rui, Xue, Rui, & Liu, Ling (2019). Security and privacy on blockchain. ACM Computing Surveys, 52(3), 1-34.

Zhao, Quanyu, Chen, Siyi, Liu, Zheli, Baker, Thar, & Zhang, Yuan (2020). Blockchain-based privacy-preserving remote data integrity checking scheme for IoT information systems. *Information Processing & Management*, 57(6), Article 102355.

Zheng, Zibin, Xie, Shaoan, Dai, Hongning, Chen, Xiangping, & Wang, Huaimin (2017). An overview of blockchain technology: Architecture, consensus, and future trends. In 2017 IEEE international congress on big data (BigData Congress) (pp. 557–564). IEEE.

Zheng, Zibin, Xie, Shaoan, Dai, Hong-Ning, Chen, Xiangping, & Wang, Huaimin (2018). Blockchain challenges and opportunities: A survey. International Journal of Web and Grid Services, 14(4), 352–375.

Zhu, Qingyi, Loke, Seng W, Trujillo-Rasua, Rolando, Jiang, Frank, & Xiang, Yong (2019). Applications of distributed ledger technologies to the internet of things: A survey. ACM Computing Surveys, 52(6), 1–34.

Zmudzinski, Adrian (2020). Decentralized lending protocol bzx hacked twice in a matter of days.

EXHIBIT F

Carmela Troncoso, Marios Isaakidis, George Danezis, and Harry Halpin

Systematizing Decentralization and Privacy: Lessons from 15 Years of Research and Deployments

Abstract: Decentralized systems are a subset of distributed systems where multiple authorities control different components and no authority is fully trusted by all. This implies that any component in a decentralized system is potentially adversarial. We revise fifteen years of research on decentralization and privacy, and provide an overview of key systems, as well as key insights for designers of future systems. We show that decentralized designs can enhance privacy, integrity, and availability but also require careful trade-offs in terms of system complexity, properties provided, and degree of decentralization. These trade-offs need to be understood and navigated by designers. We argue that a combination of insights from cryptography, distributed systems, and mechanism design, aligned with the development of adequate incentives, are necessary to build scalable and successful privacy-preserving decentralized systems.

Keywords: Decentralization, Privacy, Peer-to-peer, Systematization of Knowledge.

DOI 10.1515/popets-2017-0052 Received 2017-02-28; revised 2017-06-01; accepted 2017-06-02.

1 Introduction: the Long Road from 2001 to 2016

The successful adoption of decentralized systems such as BitTorrent [24], Tor [57], and Bitcoin [112], and the revelations of mass surveillance against centralized cloud services [74], has contributed to the wide belief that decentralized architectures are beneficial to privacy. Yet, there does not exist a foundational treatment or even an established common definition of decentralization. In this paper we aim at defining decentralization and systematizing the ways in which a system can be decentralized, and, by presenting the key design decisions in decentralized systems, bring forth past lessons that can inform a new generation of decentralized privacyenhancing technologies.

This is not the first time there has been a surge of interest in decentralization. As Cory Doctorow noted at the 2016 Decentralized Web Summit: "It's like being back at the O'Reilly P2P conference in 1999," which signaled a peak of interest around decentralized architectures at the turn of the millennium [118]. The 'hype' around decentralization was followed in the early 2000s by research and deployment activity around decentralized systems.

To some extent, decentralization was originally a response to the threat of censorship. Perhaps the first rallying cry for decentralization was the Eternity Service [8]. Anderson created this system in response to the success of the Church of Scientology at closing down the anon.penet.fi remailer [77] "as a means of putting electronic documents beyond the censor's grasp." This motivation of censorship resistance is clear in more modern systems: Tor using a decentralized network of anonymous relays and a DHT-based hidden services naming infrastructure; Bitcoin emerging as a censorshipresistant way to transfer funds to organizations like Wikileaks after the centralized e-Gold [62] online currency had been shut down by the Department of Justice; or BitTorrent succeeding as a peer-to-peer (P2P) file sharing service using Mainline DHT [164] rather than having a central indexing service like Napster that could be subject to requests to keep track of file copying [6]. In each of these cases, decentralization arose as a response to the shutdown of a centralized authority, aiming to remove that single natural point of failure.

Despite the millennial fervour for decentralization, the 2000s witnessed the rise of massively distributed, *but not decentralized*, data centers and systems as the dominant technical paradigm embodied by the Cloud computing capabilities offered by Google, Facebook, Mi-

Carmela Troncoso: IMDEA Software Institute, E-mail: carmela.troncoso@imdea.org

Marios Isaakidis: University College London, E-mail: m.isaakidis@cs.ucl.ac.uk

George Danezis: University College London, E-mail: g.danezis@ucl.ac.uk

Harry Halpin: INRIA, E-mail: harry.halpin@inria.fr

crosoft, and others. Eventually, users were diverted away from software running locally on their machines, which essentially is a form of decentralization, towards cloud applications that enabled an unprecedented aggregation of user data by the providers. Snowden's revelations in 2013 on mass surveillance programs leveraging the centralized nature of these services gave credence to longstanding privacy concerns brought about by the rise and popularity of centralized services.

The desire to preserve privacy, liberty, and the autonomous control of infrastructure and services have led to a call to "re-decentralize" the Internet [128, 179]. As a result, in the 2010s we are observing an upsurge of alternatives to centralized infrastructures and services, although most alternatives to Cloud-based applications are still under development.

It is important for system designers to neither be nostalgic about past systems nor fatalistic about future ones. Today's networking and computing environments are vastly different from those in 2000: Smart-phones have placed a powerful computer in people's pockets; users are usually connected to the Internet over fast connections without time or bandwidth caps; clients, such as web browsers, are now mature end-used platforms with P2P communications enabled and cryptographic capabilities; and mobile code, in the form of Javascript, is ubiquitous.

Even though the design space for modern decentralized systems is less restricted than in the past, fundamental challenges remain. Our key objective is to support future work on decentralized privacy systems by systematizing the past 15 years of research, between O'Reilly's publication of "Peer-to-Peer: Harnessing the Power of Disruptive Technologies" [118] in 2001, and 2016. We aim at highlighting key findings in classic designs, and also the important problems faced by designers of past systems, so as to inform the choices made by engineers pursuing decentralization today.

2 Epistemology

Scope. There is a wide use of the term 'decentralized'. In this paper, we restrict ourselves to discussing systems that support privacy properties using decentralized architectures. We draw a distinction between *decentralized* and *distributed* architectures, as follows:

Distributed system: A system with multiple components that have their behavior co-ordinated via message passing. These components are usually spatially separated and communicate using a network, and may be managed by a single root of trust or authority. Distribution is beneficial to support robustness against single component failure, scalability beyond what a single component could handle, high-availability and lowlatency under distributed loads, and ecological diversity to prevent systemic failures. Developments led by Google, ranging from BigTable [35] to MapReduce [49] are good examples of distributed systems.

Decentralized system: A distributed system in which multiple authorities control different components and no single authority is fully trusted by all others.

Following Baran [13], systems are conceived of as networks of interconnected components, where all the components of a system form a graph, where the nodes of the graph are the components and the edges the connections between them (see Fig. 1). Due to this analogy with graphs, the terms "decentralized network" and "decentralized system" tend to be used interchangeably. However, decentralized systems are not just network topologies, but systems that exist to fulfill some function or set of functions, otherwise called 'operations.' These operations are accomplished by passing messages between a sender and a receiver node, with other nodes serving as proxies to relay the message [91] (right graph in Fig. 1). On the contrary, in centralized systems messages and operations are orchestrated by a central trusted authority (depicted as an orange circle in the left graph in Fig. 1).

Centralized systems may be distributed, typically for efficiency or scaling, but not for privacy, and so the underlying components are fundamentally trusted. Only external entities are considered adversarial. Widely deployed systems such as Bitcoin, BitTorrent, and Tor are on the other hand decentralized. Contrary to generic distributed systems, in participating parties may choose their relationships of trust autonomously, including the case where there one may not trust any other components. This has profound implications in terms of security and privacy: no single entity that can act as a trusted computing base (TCB) [135] to enforce a global security or privacy policy. Any internal component of the system may be adversarial, in addition to external parties, requiring defences in depth.

In terms of security and privacy we adopt the following broad definitions, that we make more detailed at the corresponding section when the context requires clarification or preciseness.

Security: We consider the security aspects of a system to be those that encompass traditional information se-



Fig. 1. From centralized to decentralized systems

curity properties. This include of course confidentiality, integrity, and authentication; but also less traditional ones such as availability, accountability, authorization, non-repudiation or non-equivocation.

Privacy: We consider the privacy aspects of a system to be those related to the protection of users' related data (identities, actions, etc.). This protection is usually formalized in terms of privacy properties (anonymity, pseudonymity, unlinkability, unobservability) for which we follow the definitions by Pfitzmann and Hansen [121]. These definitions are extended in the privacy-oriented discussion in Section 3.3.

Methods & Model. To systematize knowledge in decentralized privacy-preserving systems we performed a systematic literature review of all papers published in the top 4 computer security conferences (IEEE S&P, ACM CCS, Usenix Security, NDSS) as well as the specialized conferences (PETS, WPES and IEEE P2P) that are proposing or analyzing decentralized systems with privacy properties, from the years 2000 to 2016.

Our first analysis resulted in 165 papers (28 from IEEE S&P, 56 from ACM CCS, 18 from Usenix Security, 11 from NDSS, 11 from PETS, 10 from WPES, and 31 from IEEE P2P). Finally the paper contains only 90 references from these venues (13 from IEEE S&P, 32 from ACM CCS, 10 from Usenix Security, 11 from NDSS, 9 from PETS, 6 from WPES, and 9 from IEEE P2P), 19 are well-known deployed systems that do not have an associated peer-reviewed publication, and the rest come from an additional pool of 30 conferences and work-shops (among them FOCI, WEIS, NSDI, SIGCOMM, SIGSAC, or CRYPTO). The selection was done on the basis of highlighting design decisions that reflect a key lesson worth of future reference.

Due to the vast amount of identified designs, by necessity we do not describe each system in detail, but instead show how each system exemplifies a property or design choice. We do, though, expand upon Tor, Bit-Torrent, and Bitcoin as they are are heavily deployed and have substantial academic analysis. As illustrated in Figure 1, we study the pool of selected designs with the intention to determine:

- 1. How is the system decentralized? (Section 3.1)
- 2. What advantages do we get from decentralizing? (Section 3.2)
- 3. How does decentralization support privacy? (Section 3.3)
- 4. What are the disadvantages of decentralizing? (Section 3.4)
- 5. What implicit centralized assumptions remain? (Section 3.5)
- 6. What can we learn from existing designs? (Section 3.6)

Insights.

- The key difference between distributed systems and decentralized systems is one of authority and trust between components. Differences in architecture and use of security and privacy controls stem from it.
- Decentralized systems embody a complex set of relationships of trust between parties managing different aspects of the system. Untrusted insiders are common, and security controls must be deployed taking into account adversaries within the system.
- In distributed, but not decentralized, systems the existence of a single authority that provisions and manages all components that are trusted enables the use of simple security, many times based on dedicated trusted components that act as roots of trust.
- In decentralized systems no single authority can provision a root of trust or trusted computing base, making security mechanisms reliant on those (such as central access control or traditional public key infrastructures) inapplicable.

3 Decentralization and Privacy

This section runs over the key questions we pose in the previous sections with regards to the current state of affairs in decentralized systems. Table 1 (page 320) provides a summary of the different design decisions and the properties achieved as a result.

3.1 How Is Decentralization Achieved?

We review key architectural decisions: how to orchestrate the infrastructure of the network, how to route messages, and how to distribute trust between nodes.

3.1.1 Infrastructure

A first key choice concerns the distribution of tasks needed for maintaining a service within the system. The provisioning of infrastructure impacts the design in terms of trust and message routing.

User-based Infrastructure. Some decentralized system consist solely of nodes that are users and there is no additional infrastructure. They rely solely on users to collectively contribute resources (bandwidth, storage) in order to provide a service. The advantage of this design is that by nature it does not require a third-party centralized authority. This user-based design can support services such as hosting of encrypted data, e.g. in Freenet [41] and Cachet [117]. A disadvantage is that user-based infrastructure may lead to poor performance due to evolving into sparsely connected topologies, and to "churn" caused by peers constantly joining and leaving the network.

User-independent Infrastructure. Here, the functions of the decentralized system are realized by nodes that are not users. A set of third-parties that are not necessarily trusted may provide all or part of the functionality to users. This design pattern underlies classic open federated protocols such as SMTP [123] and XMPP [9] based on a client-server model. The advantages of user-independent infrastructure include increased availability of the service, a reduced attack surface, and immunity to user churn. Servers do not necessarily threaten user privacy. The Eternity Service [8], as realized in systems like Tahoe-LAFS [139], combined encryption with the use of several servers controlled by different non-collaborating authorities for the private storage and replication of files. Other examples of systems that rely on user-independent infrastructure include DP5 [27] and Riposte [42] in terms of Private Information Retrieval [39] or anonymous communication systems like mix networks [36] or DC-nets [37].

Hybrid Systems. Functions may be shared between users and nodes run by third-parties. An example is Tor, where relays are mainly run by volunteers but Directory Authorities are operated by a closed 'known' group of servers. In terms of privacy and security, new elements such as distributed ledgers decentralize traditionally centralized cryptographic protocols in these hybrid systems. For example, computations can be locally and securely recorded to the blockchain with the support of multi-party computation protocols [189], even without a trusted third party [10, 189], or using a small number of stable entities to ensure reliability and lowlatency, as in the Sharemind MPC system [26].

3.1.2 Network Topology

When considering a decentralized system, there are two distinct topologies. The first, *network topology* describes the connections between nodes used to route traffic; and the second, *authority topology* describes the power relations between the nodes. Thus, the network routing structure does not necessarily have to mirror how authority is decentralized in a system, although it often does. That can greatly affect the security and privacy properties of the system [53]. It must be noted that components of traditional network routing is done in a hierarchical manner, including spanning tree protocols such as in BGP [130] in the current Internet as well as 'next generation' designs like SCION [186].

Mesh. Mesh topologies are unstructured. Nodes can route messages to every other node they are connected with. One advantage is that mesh networks function in settings with no stable connections to other nodes to guarantee service in the presence of massive churn and changing connectivity, such as in mobile adhoc networking and file sharing in early versions of Gnutella [73]. A particularly popular communication means in mesh topologies [112] are *qossip protocols*. In gossiping, as opposed to flooding, a random subset of the nodes in the network are chosen to receive the messages. These nodes then continue to broadcast the message via another independently selected random subset of the network to relay messages. The reliability of message delivery under load is questionable and information propagation experiences delays. Historically mesh networking does not preserve user privacy of their users, but recent secure messaging systems such as Briar [28] use this topology to remain functional during Internet blackouts.

Distributed Hash Tables (DHT). DHTs are network topologies where each node maintains a small routing table of its neighbours, and messages are passed greedily to known nodes that are 'closer' to the intended recipient. Although efficient and decentralized, DHTs do not by themselves provide strong security, privacy and anonymity properties. While decentralized, DHTs are not secure and privacy-preserving by default: Tran et al. [153] show that low latency anonymity systems based on DHTs such as Salsa [113] are vulnerable to having large amounts of traffic captured by adversaries controlling a fraction of the relays. DHT nodes may, however, be grouped into byzantine quorums to defeat adversaries that control a minority of nodes [180].

Super-nodes. Super-nodes are nodes that are endowed with more, and contribute more, resources to the system. This may be in terms of computation power, storage, or network connectivity, stability and up time. In terms of routing, such super-nodes may be used to mediate operations requiring higher network throughput. They can be arranged in structured topologies, designed to leverage them; or they may emerge naturally in unstructured topologies, as a result of some nodes committing more resources. Most P2P systems such as BitTorrent eventually rely on super-nodes [50]. These supernodes have serious implications on availability and integrity, as they may become targets for attack, and privacy, as they mediate, and are in a privileged position to observe, a larger fraction of activities.

Stratified. Some of the more complex decentralized systems use a stratified design where nodes have specialized roles in terms of routing, or other functions. A paradigmatic example is the Tor network. Tor users autonomously form circuits from an open-ended set of Tor relays, in layers of entry guards, middle nodes and exit nodes. A high-integrity global list of these relays is maintained through consensus by a closed group of specialized Directory Authorities. Simultaneously, Tor hidden services are resolved through a Hidden Service Directory maintained by a simple DHT topology. We note that, on some level, Tor has also evolved to use supernodes on its topology and the distribution of traffic sent through Tor relays is far from uniform [84]. Cascades, are a particular case of Stratified topologies in anonymous communications, in which paths are pre-defined. The advantages and disadvantages of such choice as opposed to free routes has been discussed in [52].

3.1.3 Authority

We now consider the relation among nodes in terms of authority and describe mechanisms to mitigate the potentially effects of power disparity that could potentially harm the security and privacy of users.

Ad-hoc: Nodes Interact Directly. In ad-hoc there is no relationship of authority among nodes. Nodes directly interact with each other without the participation of other nodes, and they do so for the benefit of the involved parties only. In terms of routing, ad-hoc requires a mesh topology where nodes do not carry traffic for other nodes. However, note that mesh topologies do not always have a ad-hoc (lack of) authority relations, such as routing based on gossip. An example of this type of system would be point-to-point communication in Briar [28]. For purposes of privacy, direct interaction bypasses possibly compromised nodes, but not network adversaries. As for confidentiality, communications can be encrypted between the two nodes, and can be extended to group communication using group key agreement protocols [138].

P2P: Nodes Assist Other Nodes. P2P designs have no central authority. Unlike ad-hoc interaction, nodes provide services and resources to other nodes, such as routing messages or storing blocks of data. Nodes have equal authority and so each node may equally compel any other node, although services and resources are usually provided according to their capacity. In other words, P2P systems self-organize and all nodes are responsible for carrying out operations for all other nodes, rather than having any pre-configured special position of authority. Since nodes are not motivated by authority to help each other, mechanisms should instead be in place to provide 'incentives' for collaborative behaviour.

There are clear advantages for the security and privacy properties in P2P systems. Information about peers is not centralized and interaction typically remains local to a few nodes, so it is difficult for an adversary to obtain a global view of the system. Yet, relying on peers for functionality poses an additional threat to privacy, since requests may be served by adversarial nodes. These nodes can passively collect information on other nodes or they may actively disrupt the integrity of operations by forging messages or replay attacks that are hard to detect. Furthermore, since P2P systems are usually open, without any admissions control, adversaries may purposely inject a large number of Sybil nodes, to increase their chances of a successful attack [59]. P2P systems are not a silver bullet for decentralization: there is no clear and definite solution to Sybil attacks in P2P networks, although such an attack can be mitigated using reputation [43] or trust [83].

Social-based: Nodes Assist Friends. These designs take advantage of pre-existing decentralized relationships, such as "friendship". In terms of applicability of security mechanisms this approach maintains most advantages of a P2P system. It is less vulnerable to Sybil attacks as adversarial nodes can be excluded from participating in the network or may be easier to detect [47], as it is harder to infiltrate a social network than a network. The downside is that, without cover traffic, a global passive adversary can discover the underlying social graph by monitoring network communications and violate privacy properties such as unobservability and unlinkability. This in turn may lead to user deanonymization [114], and techniques such as perturbation of the underlying graph may not be robust enough to prevent this [107].

A number of systems implement social-based communication to resist Sybil attacks. For instance Drac [44] and Pisces [108] use social-networks to support routing of messages. X-Vine [105] is a mechanism that, applied to distributed hash tables, helps resisting denial of service via Sybil attacks at the cost of higher latency. Tribler [124] uses social-based trust relations to improve performance that exploits similarity to improve performance, content discovery, and downloading in file sharing; or Nasir et al.'s socially-aware DHT [116], which reduce latency and improve the reliability of the communication.

Federated: Providers Assist Users. In federated designs, users are associated to *provider* nodes, which they trust and that act as authorities. Each provider is responsible only for its own users but collaborates with other providers in order to provide a service. No single provider has authority over other providers, and thus there is a "federation" of providers. Federated authorities typically use user-independent infrastructure and act as a super-node in terms of routing. This combination of design choices leads typically to high availability as long as the provider is accessible and not compromised, but the provider is a central point of attack to violate security properties and the provider itself can violate the privacy of nodes. The primary weakness of federated systems is the assumption that federated service providers largely act honestly. Some techniques can relax strong trust assumptions in the provider. End-toend encryption can maintain confidentiality [145] using providers. Computation can be obscured using secret sharing [133] or differential privacy-based solutions [3].

Accountability: Transparency Assists Users. Transparency can be used to make an authority accountable in order to establish trust. It promotes integrity of operations by monitoring the correct behavior of nodes, e.g. a transparent log of a provider's operations in a federated system audited by users or other providers acting in lieu of their associated users. The nature of this auditor's authority is very different from the aforementioned previous types of authority relations and critically relies on the non-collusion of the auditor and the audited authority, e.g., Bitcoin consensus over its blockchain using proof-of-work. Other alternatives, such as Certificate Transparency [92], rely on a set of services and auditors to keep track of X.509 certificates and quickly detect potentially rogue or hacked certificate authorities. Similarly, electronic election protocols [75] achieve robustness through proofs of correct shuffling of votes, e.g., Helios [1]. Yet naïve designs of audit logs may violate the privacy of decentralized nodes by learning too much information.

While decentralized accountability can have clear advantages regarding integrity, there are difficulties in maintaining privacy in any distributed log. This disadvantage can nevertheless be reduced as shown by Zerocash [18], which uses zero-knowledge proofs in order to maintain unlinkability in auditing relationships; or CONIKS [101], that shows that auditing the consistency of a name-key binding through time enables verification of user public keys by the end users collectively and by other providers, while concealing the identities and the number of users at each provider using Verified Random Functions.

Insights.

- Decentralization encompasses a large space of designs from decentralized ad-hoc mesh to federated super-node networks, not just peer-to-peer. These offer a variety of privacy and systems (e.g., availability, or reliability) properties. Developer instincts may often be incorrect in terms of their trade off.
- Despite being separate parts of the design, the network topology in decentralized systems often mirrors the authorities' trust relationships. However, a strict mapping between authority, infrastructure and networking topology is not necessary, and may come at the cost of harming privacy or availability.
- Centralization in terms of federated and super-nodes leads to better availability and system performance.
 However, it introduces single points of failure that impact availability and privacy. P2P models are by design more resilient to unstable routing and compromises, but entail higher engineering complexity.
- All networking topologies suffer under node churn, and pure P2P topologies must effectively address this effectively to be applicable at all.
- Decentralization does not imply the absence of any infrastructure. However, the infrastructure itself needs to be decentralized by being provided by a plurality of authorities. Such infrastructure may enhance performance by offering super-nodes or dedicated high-availability operations.

- De-facto super nodes may emerge naturally in decentralized designs, as a result of different node capabilities, and efficiency in centralizing certain operations. If this occurs outside the context of careful design, those super nodes become a single point of failure, and may lead to de facto re-centralization.
- Lack of relationships of authority imply that nodes must be willing to provide services to each other on a different basis. Designers of decentralized systems must carefully engineer such incentives, to ensure that natural (non adversarial) selfishness does not lead to dysfunction. Monetary incentives, reputation, and reciprocity can be the basis of such incentives – but off the shelf such mechanisms are often central points of failure.

3.2 The Advantages of Decentralization

In this section we discuss a number of perceived intrinsic architectural advantages to decentralized designs that make them appealing compared to their centralized counterparts.

3.2.1 Flexible Trust Models

An intrinsic advantage of decentralized architectures relates to the existence of multiple independent authorities. These create a distributed trusted computing base that ensures that a subset of rogue nodes, at least up to a certain threshold, cannot compromise the overall security properties of the whole system.

Distributed Trust. Decentralized systems leverage multiple independent authorities into a security assumption: for example, all forms of threshold cryptography [141] assure that if some fraction of participants are honest, some security property can be guaranteed. This principle can also be applied to secure multi-party computation, distributed key generation, public randomness and threshold-based decryption, and signing. One such privacy system is Vanish [72] that guarantees deletion after a pre-set expiry date. It illustrates how a multi-authority system implements properties otherwise impossible, or implausible, to when implemented by a single entity. However, the system was in practice defeated by a Sybil attack that the security properties of its DHT did not take into account [172]. Reliance on multiple authorities to regain a degree of privacy has

also been proposed for commercial cloud storage in case some providers are dishonest [146].

No Natural Central Authority. In some settings there exists no central authority and thus a decentralized architecture is a natural choice. This setting has been traditionally studied in the contexts of decentralized access control, as in TAOS [171] and SDSI [64], and 'trust management', such as Keynote [25]. In such systems a set of distributed principals make claims about users and each other, and those claims need to be assembled and used to resolve access control decisions. Bauer et al. [15] show that the task of resolving access control decisions in a decentralized setting is faster than doing so centrally.

Leveraging Existing Trust Networks. In some cases a decentralized infrastructure embeds or expresses a preexisting set of trust relationships that a system may reuse to support security properties. Systems may use the underlying social trust structure to build overlay privacy-friendly social network services, as surveyed by Paul et al. [120]. As an example, the Frientegrity system [68] provides a social network platform using untrusted providers seeing only encrypted data, where users can exchange information with 'friends' protected by cryptographic access control. This use of encryption to defend against the providers themselves is not the case for systems like Diaspora [19], an open-source project that takes a different approach: users connect to a provider they trust – that gains full visibility of their activity – and delegate the access control on the content they share with their social circles to that provider.

3.2.2 Distributed Allocation of Resources Assists with Ease of Deployment

A central premise of P2P networks is that nodes contribute spare resources, and doing away with a central authority that is forced to bear the full costs (such as Google's server costs). This reduces costs and helps ease deployment by spreading these demands amongst multiple parties. Costs are lowered as spare capacity in the existing infrastructure is used, e.g., underutilized resources given by users such as the early SETI@home project [7] and the use of users' storage in Freenet [41].

In terms of availability, decentralized architectures exhibit fewer correlated failures by virtue of being distributed. As an example the Cachet system [117] uses a pool of untrusted peers as a storage back end of a decentralized Online Social Network.

3.2.3 Resilience Against Formidable Adversaries

Location Diversity. Decentralization provides properties that are inherently difficult to centralize, such as the network location diversity needed for Tor bridges [56] to bypass censorship both on the network and legal levels. A number of designs take advantage of this, like Publius [163], in order to resist censorship, although censorship resistance itself is a separate field with many centralized, as well as decentralized, solutions.

Survivability. Decentralized architectures can be designed to survive *catastrophic* attempts to take them down or inflict crippling damage, in a way that centralized systems cannot resist [176]. This property has been used to build highly robust botnets using a peer-to-peer architecture [134]. Although these bot-nets are decentralized on the technical level, they of course maintain central but covert command and control (C&C). Those botnets have demonstrably been harder to take down using conventional techniques, but are also vulnerable to new threats that result from their decentralization, such as poisoning and enumeration of nodes. A further discussion of wider 'Darknet' survivability is provided by Zhou et al. [97].

Separation of Development from Operations. Decentralized architectures clearly separate the authorities that provide public code - and that have no access to operational data and secrets – and those that run the code. Users and nodes, deploying software, can audit any such open source code for integrity, and chose whether to deploy it. The core development team maintains the code, that is publicly visible and auditable, but upgrading is up to independent relay operators. This model is followed by both Tor and Bitcoin. As a result, attempts to coerce the Tor development team can only have an indirect and possibly highly visible effect – rendering such attempts less effective. Similarly in Ethereum, the exploitation of a vulnerability in the DAO smart-contract, led to the core developers proposing a "hard fork", and this fork was voluntarily adopted by the majority of the Ethereum mining node operators.

Publicly Verifiable Integrity. Due to the availability of multiple independent authorities, decentralized systems can implement accountability mechanisms to publicly verify integrity. Adversaries are disincentivised to compromise nodes, by ensuring attacks have an observable effect so that cheating can ideally be discovered before it has a negative effect. Verifiable logs can be used to help enable privacy as ensuring that actions are transparent enables users to know what happened with their data, as when Pulls et al. [125] use decentralization to support transparent audits of personal data accesses. Auditability is also a key feature of secure electronic election systems such as the Helios system [1]. Such systems rely on the existence of multiple authorities in a number of ways in e-voting: threshold cryptography is used for parameter and ballot generation, with privacy enforced via threshold decryption.

Insights.

- Real-world relationships of trust and authority are personal, complex and localized, and rarely hierarchical or all-or-nothing. Decentralized systems offer flexible trust models that can leverage those relationships to support security and privacy properties.
- When it comes to high-availability and survivability against powerful adversaries – particularly with legal authority – decentralized designs are not just best, but sometimes the only available option. Designs that allow operations to continue despite some authorities being adversarial or not available, are necessary to support these properties.
- Decentralization's fundamental advantage in terms of security stems from an attacker having to compromise a set of independent authorities in order to disrupt or weaken the security properties of a system. Decentralized systems that do not offer this property may be more fragile than centralized equivalents.
- Decentralized designs decouple development from operations and have a multistakeholder governance model, where node operators influence the entire system based on the software configuration they choose to deploy.
- Decentralized systems can leverage public accountability to detect and exclude compromised or misbehaving authorities. Such accountability architectures may be used instead of more complex or expensive prevention techniques, but need to ensure that auditing will be effective and eventually acted upon.
- Leveraging spare resources of nodes allows decentralized system to scale, and ease deployment. However, this by itself opens the door to high-churn and cannot be a substitute for robust incentives to participate as the system scales or nodes are asked to take on real costs.

3.3 How Does Decentralization Support Privacy?

In this section we survey the privacy properties obtained through mechanisms that are inherent to decentralized architectures. We limit ourselves to the analysis of technical properties that may be obtained in decentralized systems. We acknowledge that decentralized systems may offer both greater user privacy and autonomous control of the infrastructure. As such they are a possible technological solution to the legally-binding, but often technologically unenforced, demands from data protection laws [67, 136], that often are addressed involving a central authority, the data controller [54]. How decentralized systems relate to the law and business models is out of the scope of this paper.

Confidentiality from Third Parties. Some designs employ a decentralized architecture on the grounds that the lack of centralized components, which have full access to user data and can surveil their actions, would be beneficial to confidentiality and unobservability. Such systems may use threshold encryption [141] in order to trade off information confidentiality and information availability, such as the PASIS [176] architecture. This scheme splits the data in n "shares" and distributes it among peers in such a way that recovering m shares allows one to recover the data, but having less pieces provides no information. Similar solutions are provided by POTSHARDS [148] or Plutus [87].

Confidentiality from Peers. In P2P architectures, nodes must interact with other nodes, but they want their communications or actions to remain confidential. For example, nodes need to perform a joint computation, but do not trust each other nor a third party with their data. In this case, decentralization enables them to exchange encrypted data and obtain the sought after result without relying on any particular entity to preserve their privacy. The P4P framework [60] is such a system, in which further zero-knowledge proofs are integrated to protect computations against malicious users. More recent, blockchain-backed systems, such as Enigma [189] rely more heavily on transparency to achieve this goal. In terms of message-passing, systems that pass endto-end encrypted messages across untrusted federated servers achieve peer confidentiality.

Anonymity. Due to the distribution of resources in decentralized networks, it is expensive for one entity to observe all actions in the network and track all activities from a user. Many [70, 78, 100, 105, 113],

leverage this approach to provide anonymous communication, although the precise properties provided in terms of anonymity differ. Some decentralized systems fail to provide full anonymity but instead provide pseudonymity which is weaker [121], e.g. it allows multiple anonymous actions to be linked, providing weaker privacy, but enabling functionality such as detecting returning users and reducing the complexity of the system. For example, in Bitcoin every transaction is linked to a pseudonym and stored in the blockchain. This allows to trace money flows and avoid doublespending; but on the downside if a pseudonym is ever deanonymized (e.g. [21]), all actions from the person would be revealed. A number of decentralized systems, ranging from mix-nets [36, 45], to DC-nets [37], to Tor [57], provide some degree of anonymity.

Deniability. Deniability enables a subject to safely and believably deny having originated an action, so as to shield her from responsibility associated to performing such action. The fact that actions cannot be linked back to a user (i.e. "unlinkability" [121]), equips users with freedom to perform actions without fear of retaliation. For instance, in Freenet [41] requests are hard to link to their originator, thus users can freely search for information without revealing their preferences.

Plausible deniability is crucial in facilitating anonymous and censorship-resistant publishing, and may be implemented using cryptographic techniques allowing of 'repudiation'. This was the motivation behind the original Eternity service [8] and well-known designs such as Publius [163]or Tangler [162].

Covertness. Some systems protect even the act of participation of nodes in the decentralized network from outside observers ("unobservability"[121] if the items of interest is the existence of users). In addition to more well-known work like Tor pluggable transports [122], the Membership Concealing Overlay Network (MCON) [157] leverages this to provide strong forms of covertness. All nodes in MCON only have links with trusted friends, and a complex overlay network is jointly created that allows all nodes to communicate indirectly with all nodes. As any node only connects to other locally trusted peers, the system defends against attempts to enumerate all users by malicious nodes.

Insights.

- The key bet of decentralized systems in terms of privacy is that a local adversary may not observe all communications, data, or actions. However, global adversaries are increasingly realistic. Thus decentralized systems that rely solely on dispersion of information to provide confidentiality are fragile.

- Decentralization can harm privacy: Distributing trust and resource contribution to multiple authorities may provide adversarial nodes with extended visibility of user data and network traffic. Thus, naive decentralization designs may in fact create more, not fewer, attack points to breach privacy.
- Decentralization alone cannot balance the needs for privacy, integrity and availability. It is only combined with the use of advanced cryptography that decentralized architectures obtain those properties. In particular, the reliance on others to perform actions, may naturally expose personal information to other nodes without the use of cryptography. However, naive encryption alone may not be sufficient to support the integrity of operations that are more complex than end-to-end messaging.
- Decentralized networks can provide privacy properties like anonymity and even covertness. Yet, most real-world decentralized systems do not use the advanced cryptography and traffic analysis resistance necessary for that purpose as it increases design, implementation, operations and coordination costs.

3.4 The Disadvantages of Decentralization

Sadly, there is no free lunch in decentralization. While decentralizing has many advantages, there is no guarantee that the properties and features of centralized systems are maintained in the process. This section summarizes problems emerging when decentralizing designs. A further critique of decentralized systems, focusing on personal data, is provided by Narayanan [115].

3.4.1 Increased Attack Surface

Decentralizing systems across different nodes inherently augments the number of points (attack vectors) that an adversary could use to launch an attack or to observe the users' traffic.

Internal Adversaries. In centralized systems, system components can be monitored and evaluated by a trusted entity to detect malicious insiders. In a decentralized system it is easier to insert a malicious node undetected. A number of such attacks have been documented against decentralized systems: the predecessor attack [174, 175] uncovers communication partners in many anonymous communication schemes [37, 57, 129, 150], or the Sybil attack which can be used to bias reputation scores [59] or corrupt the information exchanged in collaborative decentralized systems [82]. Furthermore, when messages are relayed through other nodes, e.g., to gain anonymity, their content is exposed to potential adversaries, as in Crowds [129] for Web transactions or in Yacy [177] for searching information.

Traffic Analysis. Decentralization inherently implies that information will traverse a network. Even in the presence of encryption, metadata is available to external adversaries. For instance, in anonymous communications networks it has been repeatedly shown that both passive local [103] or (partially) global [84, 111], as well as active adversaries [167], can reduce or break anonymity by looking at traffic patterns.

Inconsistent Views. Decentralization typically implies that nodes have a partial, thus non-consistent, view of the network which can have an impact on integrity. These non-consistent views allow adversaries to "cheat" without being detected. For instance, in Bitcoin adversaries can perform double spending by forcing non-consistency through fast operations [89], or eclipse attacks [76] in which the adversary gains control over all connections of a target node thus isolating her from the rest of the network. Furthermore, the lack of global information results in users not necessarily making the optimal choices with respect to optimizing their privacy, as studied both in the context of anonymous communications [55] and location privacy [71].

3.4.2 Cumbersome Management

An obvious problem of decentralization is that no entity has a global vision of the system, and there is no central authority to direct nodes in making optimal decisions with regard to software updates, routing, or solving consensus. This makes the availability of a decentralized network more difficult to maintain, a factor significant enough to contribute in the failure of a system, as pointed out by the Mojo Nation developers [168]. It is very common that nodes in a decentralized system have hugely varying capabilities (bandwidth, computation power, etc.) [69, 160], making super-nodes attractive targets [102]. Finally, decentralized systems need to overcome the shortcomings of underlying technologies (such as NAT [98]), that favor the client-server paradigm over peer-to-peer networking. **Defense Difficulties**. The lack of central management hinders the establishment of effective protection mechanisms. For instance, the non-consistent view of the network not only enables attacks, but also hampers the use of collaborative approaches to detect incorrect information [88]. Similarly, it becomes extremely difficult to prevent Sybil attacks, and defenses must either leverage local information, for example defenses based on social networks [47, 181], or collaborative approaches that combine information from several nodes [119].

Routing Difficulties. A straightforward consequence of the lack of centralized control is an increased complexity in routing. Nodes do not have an overview of the network and its capabilities [149] and consequently cannot globally optimize routing decisions [183], falling back to inefficient flooding or gossiping methods in mesh topologies. This is made harder by highly diverse nodes [69], the existence of churn [11] and the reliance on possibly malicious nodes [166]. Solutions to these problems include using complex routing algorithms to enable secure and private discovery of nodes [100, 104, 108], or avoiding the use of a centralized directory via nextgeneration DHTs. The lack of centralized routing information in decentralized topologies also impacts performance as it hinders the selection of optimal routes or load balancing. We find two approaches to alleviate this problem: using local estimations to improve performance [4, 5, 152], or providing means for users to make better decisions about routing individually [144]. The latter is known to be prone to attacks [78, 110].

3.4.3 Lack of Reputation

Decentralization is also an obstacle to the implementation of accountability and reputation mechanisms. The negative effect is amplified when privacy and anonymity mechanisms are in place, as it becomes even more difficult to identify misbehaving nodes such as Sybils [79]. An effect of this lack of reputation is that nodes have no incentive to behave correctly and can misbehave to obtain advantages within the system (e.g., better performance). This problem has been identified in many settings such as P2P file sharing [184], multicast communication [182], or reputation [79]. In particular, the presence of churn, which make nodes short-lived and difficult to track over time, makes the establishment of reputation to guarantee veracity a very challenging problem [127], even more if privacy has to be preserved [137]. **Poor Incentives.** Without reputation, reciprocity and retaliation it is hard to establish incentive schemes for nodes to not be selfish, in particular in a privacy preserving manner. A solution to this problem is increasing transparency of actions, e.g. by having witnesses to report on malicious nodes in a privacy-preserving manner [187]. However, the most popular approach is the use of (anonymous) payments that incentivize good and collaborative behavior that benefits all users in the network [17, 38, 90]. In contrast, one example of negative reinforcement is the tit-for-tat strategy to encourage users to share blocks to incentivize sharing, as in BitTorrent.

Insights.

- Decentralized designs may prevent conventional attacks but also introduce new ones. Unless they are carefully designed, they may expose personal information to more, rather than fewer parties; and the need to perform joint computation across many authorities introduces threats to integrity.
- Decentralized systems are particularly susceptible to traffic analysis, compared with centralized designs, since their distributed operations are mediated through networks and adversarial nodes that may use meta-data to compromise privacy.
- Decentralized systems by nature require complex management of routing, naming and consistent state
 due to the lack of a central coordinator. Conventional defences against network attacks, like denial of service, require centralization and cannot be straightforwardly applied.
- Sybil attacks are the great unsolved problem of decentralized systems that allow open and dynamic participation. Solutions based on social networks rely on fragile social assumptions; admission control through identification or payment re-introduce centralization. Proof-of-work defences increase the cost of participation.

3.5 What Is Still Centralized in Decentralized Designs?

Even when systems claim to be decentralized, usually there are "hidden" centralized assumptions and parts of the design that need to be centralized to operate correctly. These are often implicit. In any decentralized system routing packets across the network is a challenge for both operational and privacy reasons. Typically routing can be divided in two main task. The first task is how to find candidate nodes to relay traffic, and second task is how to select among these nodes. While as detailed in Sect. 3.1.2, there are many decentralized algorithms to choose the route, actually finding candidate nodes is difficult, as highlighted in Sect. 3.4.

Centralized Directories. A common solution for the first problem is to assume that there exists a centralized directory that knows all network members. The most prominent example is the Domain Name System (DNS) that resolves easy-to-remember domain names to associated IP addresses in order to allow finding hosts in the largest known decentralized system: the Internet. Though distributed, this centralized service has serious security implications, e.g. for privacy [109] or availability [158], and thus several alternatives are being proposed [161] and deployed [58]. Another example are Tor Directory authorities [57] that provide Tor clients with the full list of onion routers. These directories solve the discovery problem but have become a bottleneck for the scalability of the system [100]. How to decentralize these authorities in an efficient, privacy-preserving manner is an active area of research. Solutions are based on having multiple copies of the publicly verifiable directory kept consistent via consensus protocol and distributed via gossiping, although it risks covertness; or to use friendof-a-friend discovery and routing [100, 106].

Path Selection. Once routing alternatives are known the question remains: Which route to choose? Thus typically, a centralized server is considered that can "rank" routing options to allow for path optimization with respect to adversaries [2, 12, 61, 86], performance [143, 144, 159], or with respect to users' reputation [165]. Such a centralized ranking approach has been shown to be vulnerable to attacks [14, 22]. Typically DHTs are the possible solution, although only a few have the necessary security and privacy properties for use in decentralized systems [46].

Distributed Computations. A number of decentralized systems are designed with the assumption that there is a central entity that performs computations on the data collected by the nodes in the system. Paradigmatic examples of this behavior are decentralized sensor networks [34, 65, 188] where the challenge is to send decentralized measurements to a "master" node, but there exist other applications such as distributed network monitoring for intrusion detection [126], anonymous surveys [80], or private statistics [63] in which, even though nodes perform decentralized computations, interaction with a central authority is needed to produce the final result.

3.5.2 Trust Establishment

A challenge when decentralizing networks is to ensure that nodes can be trusted to perform the actions they are assigned or can authenticate themselves as the intended receiver of a message. Often, to avoid dealing with this problem, a common implicit centralized assumption is that a set of trusted servers is assumed to exist, such as in Dissent [173] or the Directory Authorities in Tor.

Decentralized trust establishment is still an open problem, though some of the excitement around mining in Bitcoin is precisely due to their attempt to avoid this problem and so build a 'trustless' decentralized system.

Authentication. In general certificate infrastructures are not decentralized, e.g., PKI. Therefore, some decentralized systems rely on centralized certification authorities to authenticate nodes that can be used for secure routing [33, 147], user authentication [29], or to enrol users in the system in the context of anonymous credentials [16, 30, 31], a privacy-preserving alternative for authentication without requiring user identification. Such centralized authorities are simpler for deployability or usability, but become a single point of failure as pointed out by Lesueur et al. in [95]. They also introduce an imbalance of power unnatural for decentralized environments since they allow a single entity to revoke peers' authentication credentials. Many decentralized designs do not address authentication (e.g. [117, 142], see [120] for more details), although work from TAOS [170] and SDSI [132] onwards has been working in this direction [20]. Authentication is useful to prevent Sybil attacks, and work on decentralized and privacy-preserving authentication via threshold cryptography is one promising solution [99], as is the use of zero-knowledge systems for anonymous credentials [16].

Authorization. Assuming the existence of a centralized entity is also common when it comes to storing and enforcing authorization policies, as highlighted
by numerous efforts to decentralize policy management and enforcement from SDSI [132] to more recent systems [94, 96, 169]. OAuth was designed to be federated in terms of authorization, but in practice only a few large providers use this standard [140]. So if an adversary compromises a user's single authentication method such as a password, it can compromise them across multiple decentralized systems. Work descending from SDSI [132] to limited-time authorization via pseudonyms and blind signatures present one way forward to decentralize authorization [99].

Abuse Prevention. As mentioned in Sect. 3.4 accountability is a challenge in decentralized systems. Hence, existing abuse-prevention schemes end up relying on centralized parties, often determining global reputation scores. Solutions based on blacklistable credentials (anonymous credentials for which authorization can be selectively revoked) use a centralized authority for enrollment [154, 155], or to store blacklists [85, 156]. Similarly, identity escrow [23] or revocable anonymous communication solutions [40], that allow for re-identification of misbehaving users require a centralized party that stores those identities. In practice, spam prevention in federated email systems also uses centralized lists of known spammers. Typically, these are built from preexisting trusted social networks, and only recently have reputation systems such as AnonRep (based on homomorphic encryption and verified shuffles) allowed reputation to be done in a privacy-preserving and decentralized manner [185].

Payment Systems. In many applications of decentralized services it could be desirable to count on a payment system to reward peers for their contributions. While many alternatives have been presented in the literature specifically aimed at peer to peer systems, e.g. [17, 32, 178], they inherently rely on a centralized authority that opens accounts (the bank) and sometimes even on other authorities that can act as "arbiters" in case of dispute [17], or on authorities that record transactions to help taxation on the operations run in the system, even if the transactions are anonymized [151]. Decentralized crypto-currencies can help ameliorate this problem.

Trusted Developer Community. All decentralized systems work by virtue of having the nodes communicate via the same protocol. Thus, the actual software can be a centralized point of failure if the protocol is flawed. If the protocol is standardized or otherwise uniformly specified, the implementation of the protocol itself may be a failure. Furthermore, the developers themselves could be compromised. his danger is augmented by the software monoculture prevalent in deployed systems, that results in a bug in a popular platform capable of compromising a large set of authorities. One solution is to apply the technique of forcing public transparency and auditing of the integrity of the development process. Open-source development, done in public repositories, is increasingly required. Integrity is ensured via deterministic builds [131] so that everybody can verify the genuine binary, and the authority to run new versions of the software remains in the hands of the operators. This approach is already followed by Tor and increasingly by Bitcoin, where the choice to deploy particular open-source code is up to miners.

Insights.

- Many decentralized systems implicitly rely on centralized components to hold network information for efficient routing or for establishing trust and defending against Sybil attacks.
- Essential user-facing infrastructure, from authentication to authorization is centralized even in decentralized systems. Developing alternatives seems to be an open problem, with no clear established design. For payments, Bitcoin has recently provided a decentralized solution, but it suffers from a number of scalability, privacy, and financial volatility problems.
- The developer community of a system is usually an implicit centralized authority, making social attacks on the developer community itself one of the largest dangers to any decentralized system.

3.6 Systematization of Existing Designs

Table 1 presents a systematic analysis of decentralized designs, clustered based on their principal goal. The columns infrastructure, network topology, authority relations, privacy properties, follow closely the definitions of the previous subsections. We applied some level of simplification to complex systems with multiple components or multiple use-cases. The systematization focuses on parts of the system relevant for its main use-case as used in prototype or deployment.

Insights.

 Many systems that provide good coverage of privacy properties and decentralization (usually via DHTs) have not been widely deployed



Anonymi ability services trailed		/ /	,,	`	· ·	· · / /			· · /		,,	` ` `	` `		`	`			`	>			/ /	` ` `	\$		<pre></pre>		 	`		~ ~	`				> >
peer		>	>	>		>			>			>		>	>	>			>	>			>	>		>	>	>	>	>			>			>	
3rd P		>	>	>	>	>			>			>	>	>	>	>			>	>			>	>	>	>	>	>	>	>			>			>	
Authority		P2P	P2P	P2P	P2P	Social		P2P	P2P	P2P	Federated	Federated	Social	P2P	Federated	Federated		Ad-hoc±	Ad-hoc±	P2P	Ad-hoc±		Federated	Federated	Ad-hoc	Federated	Federated	Federated	Social	P2P		Federated	Social		$Federated{\pm}$	$Federated\pm$	$Federated{\pm}$
Network Topology		Stratified	Super-Node	DHT	Mesh	Mesh		Super-Node	DHT	Super-Node	Mesh	Super-Node	DHT	DHT	Super-Node	Stratified		Super-Node	Super-Node	Mesh	Super-Node		Stratified	Stratified	Mesh	Stratified	Stratified	Stratified	Mesh	DHT		Stratified	DHT		Stratified	Super-Node	Stratified
Infrastructure		Hvbrid	User-independent	User-based	User-based	User-Based	0	User-based	User-based	User-based	User-independent	User-independent	User-based	User-based	User-independent	User-independent		User-based	User-based	User-based	User-based		User-independent	User-independent	User-based	User-independent	User-independent	User-independent	User-based	User-based		User-based	User-based		User-independent	User-based	User-independent
System	User Anonymity	Tor± [57]	Mixnets‡ [36, 45]	I2P‡ [81]	Crowds§ [129]	MCON§ [157]	File Sharing/Censorship Resistance	BitTorrent± [24]	Freenet‡ [41]	Gnutella‡ [73]	Publius† [163]	Eternity§ [8]	Tribbler‡ [124]	Vanish† [72]	Tangler§ [162]	Tahoe-LAFS‡ [139]	Cryptocurrencies	Bitcoin [‡] [112]	Zerocash‡ [18]	MojoNation‡ [168]	Ethereum‡ [66]	Secure Messaging	SMTP+PGP‡ [123]	ХМРР+ОТR‡ [9]	Briar† [28]	DP5† [27]	Riposte§ [42]	Dissent/Buddies§ [173]	Drac§ [44]	ShadowWalker \S [104]	Social Applications	Diaspora‡ [51]	X-Vine§ [105]	Auditable Systems	CONIKS [†] [101]	Enigma† [189]	Certificate Transparency [‡] [92]

- Widely deployed systems either are userindependent federated systems or user-based DHTbased systems, both without advanced privacy properties.
- Hybrid and stratified systems such as Tor provide provide advanced privacy properties at the cost of centralized assumptions.
- The space of ad-hoc, mesh, and covert designs is under-explored.

4 Future Research Lines

4.1 Address Decentralization's Shortcomings

To build the next generation of decentralized systems, good will, slogans, and demands are not enough. What is needed is a clear research plan. A number of designs we review consider decentralization as a goal and virtue in itself and do too little to address the inherent challenge of maintaining privacy properties and deployment with high availability. In particular we studied in Section 3.4 a number of those challenges: an increased attack surface, with corrupt insiders; susceptibility to peers violating privacy and vulnerability to traffic analvsis, integrity and consistency attacks; expensive and fragile routing; potential degradation in performance; loss of central choke points to enforce security controls; peer diversity and lack of incentives. These are serious and real threats, and not acknowledging them and confronting them head on leads to weak systems that cannot credibly compete with centralized solutions. This is demonstrated by the failure of Ethereum to promptly address the DAO vulnerability [48]. Indeed, decentralization in the style of early BitTorrent simply ends up being an inefficient way to do redundancy and availability without a centralized authority — and with no credible privacy properties. Likewise, Bitcoin and Ethereum provides this style of decentralization with the addition of integrity but their simplistic accountability designs harms privacy. Therefore, more research is required looking at systems such as Tor and Bitcoin as platforms rather than purely as channels, including understanding their interfaces, performance, quality of service guarantees and the privacy properties as a whole system in order to deliver better privacy properties.

Availability without centralization is a key promise of decentralized systems, but often fails when the system grows. The most important engineering challenge

of those reviewed is that decentralized systems often do not scale and are inefficient in comparison to centralized systems. In practice, in a world with limited resources and investment, inefficient decentralization leads to a failure of decentralization. This problematic dynamic is built into decentralized designs: maintaining highintegrity requires a majority to honestly participate in decisions. Although one could point to Bitcoin as a success, the larger Bitcoin network of miners grows the less it scales, as all miners need to detect and verify new blocks and transactions. Even worse, Ethereum smart contracts are executed on each node in the network. In both Bitcoin and Ethereum, as the number of nodes grows, the system gets slower. Due to this unfortunate design flaw, Bitcoin and Ethereum will face serious issues when scaling without major design changes that accountability as such does not address. We can be assured the current generation of attempts to "re-decentralize" the Internet will fail without more research on how to scale efficiently.

Finally, there has to be a deeper acceptance that even honest users and peers in decentralized systems will have to be incentivised to participate and behave cooperatively. This is particularly true when stronger privacy protections are implemented and reputation based on repeated and iterated interactions cannot be leveraged. In those cases standard platforms must be developed to prevent Sybil attacks and establish privacy preserving reputation to curtail abuse; accounting and payment mechanisms need to be devised to ensure that those that do work are rewarded to sustain their operations. Systems that do not provide incentives for participation in the infrastructure will fall foul of the tragedy of the commons and will remain mere proofs of concepts.

Even with motivated users, human fallibility must be addressed realistically. Decentralization advocates desire of users to return to a 'lost golden age' of selfhosting services, as in the 're-decentralize' project [128]. However, the popularity of services like Facebook and Gmail shows that most do not have the time or skills to host decentralized nodes unless a powerful incentive exists such as file-sharing. Worse, users may not be qualified at protecting their own systems, when even most skilled professional administrators cannot. Building successful decentralized systems that do not betray the security and privacy of their users is hard, and entails much more than tacking a blockchain or P2P network to a pre-existing problem, but also has to take into account platform security and ease of user operations. Systems that claim to be decentralized today simply often use the adjective in an informal manner, resulting in decentralized "snake oil", as is the case for some blockchain-based start-ups. Unlike formal security definitions, information-theoretic definitions of anonymity, and differential privacy, there are no coherent quantitative metrics to characterize decentralization. Aside from having a common definition of the privacy and security properties, decentralization engineering also requires the development of design strategies that measure both decentralization and its effect on the properties systematized earlier. More often than not, properties are neglected, rarely mentioned or evaluated, including the impact of decentralization and availability. Section 3.1, for instance, illustrates the variety of options in this design space.

Beyond the impact of decentralization on availability, a key missing piece is a systematic means for evaluating the privacy and security properties provided by a given decentralization system. As we evidence, decentralization can support privacy in many ways (Section 3.3), as well as supporting other properties too (Section 3.2). We observe that systems are often designed with one particular privacy goal in mind, which is frequently redefined to suit the design, and system designers tend to resort to ad-hoc evaluation. A particular case in which a lack of systematic evaluation has great impact in terms of understanding the protection provided by decentralized system is the case of compound systems (i.e, systems that combine different schemes to try to improve overall protection); or the case where systems are deployed in environments with different characteristics than those assumed in their design. In decentralized systems, it is not granted that the protection of the whole is greater or equal than the sum of the parts. In fact, the inverse may hold: combining different decentralized systems with different assumptions may violate the properties each system guarantees by itself. For example, while a user may assume using BitTorrent over Tor provides anonymity for file-sharing, in fact the reverse holds: Tor provides no anonymity to UDP-based systems like BitTorrent, and users can even be deanonymized by virtue of running BitTorrent [93]. In other words, systems to not exist in a vacuum. Their analysis and evaluation needs to account for interactions with their environment or other systems.

A similar trend is observed in terms of measuring the severity of disadvantages introduced by decentralization. Though, as we show in Section 3.4, many weaknesses arise from decentralizing, few works evaluate their implications, or do so in a design specific way that is difficult to extrapolate to other systems. As a result it is extremely difficult to compare systems and find promising new directions. This slows the development of robust decentralized systems by obscuring good design decisions. For example, in many systems there is a trade-off between privacy and availability.

Further work is also required to radically simplify the deployment and management of "real-world" decentralized applications, either on larger platforms or as stand-alone distributed systems. Deployability and usable application life-cycle support is at the heart of the current centralized cloud-based 'dev-ops' revolution, and has made centralized app stores and Web applications as popular as they are. Yet, there are no equivalent tools or technologies to facilitate the deployment, management, and monitoring of decentralized systems, let alone their continuous updates, application life-cycle management, and telemetry. This gap negatively affects developer's productivity and makes the engineering and maintenance of decentralized systems very expensive. Building toolchains that support easy management without introducing any central control – is largely an open research problem. Successful projects such as Tor and Bitcoin have developed best practices and running code in that space such as open-source development and reproducible builds [131] to address security concerns that may be generalized.

Key Research Questions for Decentralization.

- Are there generalized techniques to provide privacy and integrity properties for decentralized systems without damaging availability?
- Can we develop systematic techniques to evaluate decentralized systems both in isolation and when they are deployed in different environments?
- How can human users be incentivised to work in a decentralized manner?
- How do real-world deployment of decentralization lead to scalability challenges that change the desired properties and defeat decentralization?
- Can we develop a mathematical metric to define degrees of decentralization?

In the next section we will provide provisional answers to these questions to guide future research. These answers will be based on the observations built in previous sections.

5 Conclusions: Towards Full Decentralization

Availability, Privacy, and Integrity. Our analysis points to some fundamental trade-off between availability, privacy, and integrity in decentralized systems: A good design for one is an unsafe design pattern for another. Systems use a wide variety of infrastructure, network topology, and authority relation choices (as systematized in Table 1). Three widely deployed decentralized systems demonstrate a different set of design goals. Bitcoin comes with high-integrity at the cost of a public ledger with little privacy. Tor routers provide high-privacy at the cost of no available or correct collective statistics to ensure the integrity of the entire system. BitTorrent provides high availability in downloading files, but fails to provide privacy to its users against powerful adversaries.

We believe it is not pre-ordained that there is a trade-off between privacy, availability, and integrity in decentralized systems by virtue of using advanced cryptographic techniques. Unlike Bitcoin, Zerocash[18] combines both privacy and integrity using zero-knowledge proofs. Likewise, many academic systems, such as Drac[44], tackle traffic analysis to defend privacy in a P2P network. Simply put, advanced techniques for providing everything from dummy traffic for anonymity to succinct zero-knowledge proofs are not yet part of the toolbox for many decentralized system engineers.

Interdisciplinarity. Reviewing the literature reveals that to build good secure privacy-preserving decentralized systems, one needs:

- Expertise in building *distributed systems*, as decentralized systems are by definition distributed.
- Knowledge of modern *cryptography*, as complex cryptographic protocols are necessary to achieve simultaneously privacy, integrity and availability.
- An understanding of mechanism design, game theory and sociology to motivate cooperation amongst possibly selfish actors.

The focus on social incentive structures is usually left out, and thus most decentralized systems do not gain real-world wide deployment. In general, the involvement of nodes in decentralized systems varies and this is usually mirrored in the power allowed to authorities, as well as in inter-node relationships that reflect social behavior. Some designs assume centralized components, for better availability and performance. Others push for sheer decentralization, in pursuit of resilience to censorship and network outages. Are these design choices often social or political rather than technical? Most designs, though, fall somewhere in the middle and generally impose cryptographic techniques and rely on real-world dynamics in order to defend against adversarial nodes. Certainly, the way decentralization is achieved affects the privacy of the users and thus their behavior. It falls upon decentralized system designers to achieve satisfactory performance and deployability, while taking into account not just the technical but the necessary social structure of the system.

Real-world Scalability. From our study of the literature, we have shown that a number of key functions of decentralized systems often fall-back to centralized models in practice for scalability, even when unnecessary. First, network directories, key management, and naming often remain centralized. Thus, the there is a need to design of collective high-integrity and re-usable infrastructures to support directories, node discovery, and key exchange. These mechanisms need to scale up and remain decentralized, while not being open to corruption or inconsistencies.

Second, reputation and abuse control often require either centralized entities, or building on pre-existing social networks in user-based infrastructure. Even advanced privacy-preserving techniques, such as anonymous blacklisting, assume that centralized services will issue and bind identities, and e-cash protocols rely on a bank to issue coins and prevent double spending. More work is required in establishing reputation in decentralized systems and preventing abuse without resorting to central points of control.

Third, it is important to make credible assumptions about the platform security and computing environment of end-users or other devices. It is too facile to heavily rely on end-user systems keeping secret keys and data, and ignore that they are often compromised. Achieving perfect end-point security is an ambitious goal in and of itself – and so needed but beyond the strict remit of building secure decentralized systems. Decentralized architectures that display or limit the effect of compromises, and which may 'heal' and recover privacy properties following hacks, should be preferred to those that fail catastrophically or silently under those conditions.

Defining Decentralization. In general, decentralized systems are networks. Yet as shown by the difference between network topologies for routing and the relationships of authority, a decentralized network is not simply a single network, but multiple kinds of networks

324

connected on different levels of abstraction. Worse, the overly simplified models of decentralization presented in many papers and research prototypes do not take into account the changes produced by real-world usage into account. As shown by BitTorrent, simple decentralized networks tend to evolve from P2P into super-node systems. In general, as a system scales there is a tendency towards distribution, but not decentralization, in order to maintain efficiency. Using network science, one can show simple models such as random graphs with basic mechanism design such as preferential attachment scale into small-world systems over time, and these systems often simply transform into a federated client-server architecture or a simple centralized distributed system. In order to maintain decentralization as an emergent property, it appears that advanced hybrid and stratified system, e.g. Tor, are necessary to "unnaturally" maintain decentralization and the relevant privacy properties. Yet, the Tor network has many centralized technical (complete network information by directory authorities) and social assumptions (control by a core group of developers). The key point of a real measure of decentralization should be to take these more stratified designs into account. An ideal decentralized system would remove all centralized assumptions while maintaining the needed security and privacy properties.

The ultimate bet of decentralized systems is still open: is being vulnerable to a (possibly random) subset of decentralized authorities better than being vulnerable to a single centralized authority? Decentralization seems to be the result of a breakdown in trust in centralized institutions, but we do not yet understand how to build decentralized social institutions to support decentralized technical systems despite the promises of Bitcoin to produce algorithmic monetary policy, or the promise of Ethereum to support modern civilization with scripts with dubious security properties. Decentralization is a hard problem, but the fact that it is technically amendable to advanced techniques from distributed systems and cryptography should indicate that the social questions at the heart of decentralization are not unsolvable.

Acknowledgements. The authors would like to thank the reviewers for insightful comments that helped improving the paper, in particular Prateek Mittal for acting as shepherd. This work is supported by the EU H2020 project NEXTLEAP (GA 688722).

References

- B. Adida. Helios: Web-based open-audit voting. In 17th USENIX Security Symposium, 2008.
- [2] M. Akhoondi, C. Yu, and H. V. Madhyastha. LASTor: A low-latency as-aware tor client. In *IEEE Symposium on Security and Privacy*, 2012.
- [3] D. Alhadidi, N. Mohammed, B. C. M. Fung, and M. Debbabi. Secure distributed framework for achieving *e*differential privacy. In *12th Privacy Enhancing Technologies Symposium*, 2012.
- [4] M. AlSabah, K. S. Bauer, and I. Goldberg. Enhancing Tor's performance using real-time traffic classification. In 19th ACM Conference on Computer and Communications Security, 2012.
- [5] M. AlSabah, K. S. Bauer, I. Goldberg, D. Grunwald, D. McCoy, S. Savage, and G. M. Voelker. DefenestraTor: Throwing Out Windows in Tor. In *11th Privacy Enhancing Technologies Symposium*, 2011.
- [6] A&M Records, Inc. v. Napster, Inc., 239 F.3d 1004. https: //law.resource.org/pub/us/case/reporter/F3/239/239.F3d. 1004.00-16403.00-16401.html, 2001. Last accessed: June 20, 2017.
- [7] D. P. Anderson, J. Cobb, E. Korpela, M. Lebofsky, and D. Werthimer. Seti@ home: an experiment in publicresource computing. *Communications of the ACM*, 45(11):56–61, 2002.
- [8] R. Anderson. The Eternity service. In *Pragocrypt*, 1996.
- [9] P. S. Andre. IETF RFC 6120 Extensible Messaging and Presence Protocol (xmpp): Core. https://www.ietf.org/rfc/ rfc6120.txt, 2011. Last accessed: June 20, 2017.
- [10] M. Andrychowicz, S. Dziembowski, D. Malinowski, and L. Mazurek. Secure multiparty computations on Bitcoin. In *IEEE Symposium on Security and Privacy*, 2014.
- [11] M. S. Artigas and P. G. López. On routing in Distributed Hash Tables: Is reputation a shelter from malicious behavior and churn? In 9th IEEE Conference on Peer-to-Peer Computing, pages 31–40, 2009.
- [12] M. Backes, A. Kate, S. Meiser, and E. Mohammadi. (nothing else) MATor(s): Monitoring the anonymity of Tor's path selection. In 21st ACM Conference on Computer and Communications Security, 2014.
- [13] P. Baran et al. On distributed communications. Volumes I-XI, RAND Corporation Research Documents, August, 1964.
- [14] K. S. Bauer, D. McCoy, D. Grunwald, T. Kohno, and D. C. Sicker. Low-resource routing attacks against Tor. In ACM Workshop on Privacy in the Electronic Society, 2007.
- [15] L. Bauer, S. Garriss, and M. K. Reiter. Distributed proving in access-control systems. In *IEEE Symposium on Security* and Privacy, 2005.
- [16] M. Belenkiy, J. Camenisch, M. Chase, M. Kohlweiss, A. Lysyanskaya, and H. Shacham. Randomizable proofs and delegatable anonymous credentials. In 29th International Cryptology Conference Advances in Cryptology, 2009.
- [17] M. Belenkiy, M. Chase, C. C. Erway, J. Jannotti, A. Küpçü, A. Lysyanskaya, and E. Rachlin. Making P2P accountable without losing privacy. In ACM Workshop on Privacy in the

Electronic Society, 2007.

- [18] E. Ben-Sasson, A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, and M. Virza. Zerocash: Decentralized anonymous payments from Bitcoin. In *IEEE Symposium on Security and Privacy*, 2014.
- [19] A. Bielenberg, L. Helm, A. Gentilucci, D. Stefanescu, and H. Zhang. The growth of diaspora-a decentralized online social network in the wild. In *IEEE Conference on Computer Communications Workshops*, 2012.
- [20] A. Birgisson, J. G. Politz, Úlfar Erlingsson, A. Taly, M. Vrable, and M. Lentczner. Macaroons: Cookies with contextual caveats for decentralized authorization in the cloud. In *Network and Distributed System Security Symposium*, 2014.
- [21] A. Biryukov, D. Khovratovich, and I. Pustogarov. Deanonymisation of clients in bitcoin P2P network. In 21st ACM Conference on Computer and Communications Security, 2014.
- [22] A. Biryukov, I. Pustogarov, and R. Weinmann. Trawling for Tor Hidden Services: Detection, measurement, deanonymization. In *IEEE Symposium on Security and Privacy*, 2013.
- [23] J. Biskup and U. Flegel. Threshold-based identity recovery for privacy enhanced applications. In 7th ACM Conference on Computer and Communications Security, 2000.
- [24] BitTorrent. http://www.bittorrent.org/. Last accessed: June 20, 2017.
- [25] M. Blaze, J. Feigenbaum, and A. D. Keromytis. Keynote: Trust Management for Public-Key Infrastructures (position paper). In 6th International Workshop on Security Protocols, 1998.
- [26] D. Bogdanov, S. Laur, and J. Willemson. Sharemind: A framework for fast privacy-preserving computations. In 13th European Symposium on Research in Computer Security. 2008.
- [27] N. Borisov, G. Danezis, and I. Goldberg. DP5: A private presence service. *Proceedings on Privacy Enhancing Technologies*, 2015(2):4–24, 2015.
- [28] The Briar Project. https://briarproject.org. Last accessed: June 20, 2017.
- [29] S. Buchegger, D. Schiöberg, L. Vu, and A. Datta. Peer-SoN: P2P social networking: early experiences and insights. In 2nd ACM EuroSys Workshop on Social Network Systems, 2009.
- [30] J. Camenisch, S. Hohenberger, M. Kohlweiss, A. Lysyanskaya, and M. Meyerovich. How to win the clonewars: efficient periodic n-times anonymous authentication. In 13th ACM Conference on Computer and Communications Security, 2006.
- [31] J. Camenisch and A. Lysyanskaya. An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In *International Conference on* the Theory and Application of Cryptographic Techniques Advances in Cryptology, 2001.
- [32] J. Camenisch, A. Lysyanskaya, and M. Meyerovich. Endorsed e-cash. In 2007 IEEE Symposium on Security and Privacy, 2007.
- [33] M. Castro, P. Druschel, A. J. Ganesh, A. I. T. Rowstron, and D. S. Wallach. Secure routing for structured peerto-peer overlay networks. In 5th USENIX Symposium on

Operating System Design and Implementation, 2002.

- [34] H. Chan and A. Perrig. Efficient security primitives derived from a secure aggregation algorithm. In 15th ACM Conference on Computer and Communications Security, 2008.
- [35] F. Chang, J. Dean, S. Ghemawat, W. C. Hsieh, D. A. Wallach, M. Burrows, T. Chandra, A. Fikes, and R. E. Gruber. Bigtable: A distributed storage system for structured data. ACM Transactions on Computer Systems (TOCS), 2008.
- [36] D. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Commun. ACM*, 1981.
- [37] D. Chaum. The Dining Cryptographers problem: Unconditional sender and recipient untraceability. J. Cryptology, 1988.
- [38] Y. Chen, R. Sion, and B. Carbunar. XPay: practical anonymous payments for tor routing and other networked services. In ACM Workshop on Privacy in the Electronic Society, 2009.
- [39] B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan. Private information retrieval. J. ACM, 1998.
- [40] J. Claessens, C. Díaz, C. Goemans, J. Dumortier, B. Preneel, and J. Vandewalle. Revocable anonymous access to the Internet? *Internet Research*, 2003.
- [41] I. Clarke, O. Sandberg, B. Wiley, and T. W. Hong. Freenet: A distributed anonymous information storage and retrieval system. In *Designing Privacy Enhancing Technologies, International Workshop on Design Issues in Anonymity and Unobservability*, 2000.
- [42] H. Corrigan-Gibbs, D. Boneh, and D. Mazières. Riposte: An anonymous messaging system handling millions of users. In 2015 IEEE Symposium on Security and Privacy, 2015.
- [43] E. Damiani, S. D. C. di Vimercati, S. Paraboschi, P. Samarati, and F. Violante. A reputation-based approach for choosing reliable resources in peer-to-peer networks. In 9th ACM Conference on Computer and Communications Security, 2002.
- [44] G. Danezis, C. Díaz, C. Troncoso, and B. Laurie. Drac: An architecture for anonymous low-volume communications. In 10th Privacy Enhancing Technologies Symposium, 2010.
- [45] G. Danezis, R. Dingledine, and N. Mathewson. Mixminion: Design of a type III anonymous remailer protocol. In *IEEE Symposium on Security and Privacy*, 2003.
- [46] G. Danezis, C. Lesniewski-Laas, M. F. Kaashoek, and R. Anderson. Sybil-resistant dht routing. In *European* Symposium On Research In Computer Security, pages 305– 318. Springer, 2005.
- [47] G. Danezis and P. Mittal. Sybillnfer: Detecting sybil nodes using social networks. In *Network and Distributed System Security Symposium*, 2009.
- [48] Critical update re: Dao vulnerability. https://blog. ethereum.org/2016/06/17/critical-update-re-daovulnerability/. Last accessed: June 20, 2017.
- [49] J. Dean and S. Ghemawat. MapReduce: simplified data processing on large clusters. *Communications of the ACM*, 2008.
- [50] C. Decker, R. Eidenbenz, and R. Wattenhofer. Exploring and improving BitTorrent topologies. In 13th IEEE International Conference on Peer-to-Peer Computing, 2013.

- [51] diaspora*: The online social world where you are in control. https://diasporafoundation.org/. Last accessed: June 20, 2017.
- [52] C. Díaz, G. Danezis, C. Grothoff, A. Pfitzmann, and P. F. Syverson. Panel Discussion - Mix Cascades Versus Peerto-Peer: Is One Concept Superior? In *Privacy Enhancing Technologies*, pages 242–242, 2004.
- [53] C. Díaz, S. J. Murdoch, and C. Troncoso. Impact of network topology on anonymity and overhead in low-latency anonymity networks. In 10th Privacy Enhancing Technologies Symposium, 2010.
- [54] C. Diaz, O. Tene, and S. Gurses. Hero or villain: The data controller in privacy law and technologies. *Ohio St. LJ*, 74:923–963, 2013.
- [55] R. Dingledine and N. Mathewson. Anonymity loves company: Usability and the network effect. In 5th Workshop on the Economics of Information Security (WEIS), 2006.
- [56] R. Dingledine and N. Mathewson. Design of a blockingresistant anonymity system. *The Tor Project, Tech. Rep*, 1, 2006.
- [57] R. Dingledine, N. Mathewson, and P. F. Syverson. Tor: The second-generation onion router. In 13th USENIX Security Symposium, 2004.
- [58] Dot-Bit: Secure Decentralized DNS. https://bit.namecoin. info/. Last accessed: June 20, 2017.
- [59] J. R. Douceur. The sybil attack. In 1st International Worksop on Peer-to-Peer Systems, 2002.
- [60] Y. Duan, N. Youdao, J. Canny, and J. Z. Zhan. P4P: practical large-scale privacy-preserving distributed computation robust against malicious users. In 19th USENIX Security Symposium, 2010.
- [61] M. Edman and P. F. Syverson. AS-awareness in tor path selection. In 16th ACM Conference on Computer and Communications Security, 2009.
- [62] e-gold. http://e-gold.com/. Last accessed: June 20, 2017.
- [63] T. Elahi, G. Danezis, and I. Goldberg. PrivEx: Private collection of traffic statistics for anonymous communication networks. In 21st ACM Conference on Computer and Communications Security, 2014.
- [64] C. M. Ellison. Establishing identity without certification authorities. In 6th USENIX Security Symposium, 1996.
- [65] L. Eschenauer and V. D. Gligor. A key-management scheme for distributed sensor networks. In 9th ACM Conference on Computer and Communications Security, 2002.
- [66] Ethereum Project. https://www.ethereum.org/. Last accessed: June 20, 2017.
- [67] European Data Protection Supervisor. Opinion on privacy in the digital age (march 2010): "Privacy by Design" as a key tool to ensure citizen's trust in ICTS, 2010.
- [68] A. J. Feldman, A. Blankstein, M. J. Freedman, and E. W. Felten. Social networking with Frientegrity: Privacy and integrity with an untrusted provider. In 21th USENIX Security Symposium, 2012.
- [69] M. Feldotto, C. Scheideler, and K. Graffi. HSkip+: A selfstabilizing overlay network for nodes with heterogeneous bandwidths. In 14th IEEE International Conference on Peer-to-Peer Computing, 2014.
- [70] M. J. Freedman and R. Morris. Tarzan: A peer-to-peer anonymizing network layer. In 9th ACM conference on Computer and communications security, 2002.

- [71] J. Freudiger, M. H. Manshaei, J. Hubaux, and D. C. Parkes. On non-cooperative location privacy: a gametheoretic analysis. In 16th ACM Conference on Computer and Communications Security, 2009.
- [72] R. Geambasu, T. Kohno, A. A. Levy, and H. M. Levy. Vanish: Increasing data privacy with self-destructing data. In 18th USENIX Security Symposium, 2009.
- [73] Gnutella: File sharing and distribution network. http://rfcgnutella.sourceforge.net/. Last accessed: June 20, 2017.
- [74] G. Greenwald. No place to hide: Edward Snowden, the NSA, and the US surveillance state. Macmillan, 2014.
- [75] D. A. Gritzalis. Secure electronic voting, volume 7. Springer Science & Business Media, 2012.
- [76] E. Heilman, A. Kendler, A. Zohar, and S. Goldberg. Eclipse attacks on Bitcoin's peer-to-peer network. In 24th USENIX Security Symposium, 2015.
- [77] S. Helmers. A brief history of anon.penet.fi: the legendary anonymous remailer. CMC Magazine, 1997.
- [78] M. Herrmann and C. Grothoff. Privacy-implications of performance-based peer selection by onion-routers: A realworld case study using I2P. In *Privacy Enhancing Technologies*, 2011.
- [79] K. J. Hoffman, D. Zage, and C. Nita-Rotaru. A survey of attack and defense techniques for reputation systems. ACM Comput. Surv., 2009.
- [80] S. Hohenberger, S. Myers, R. Pass, and A. Shelat. AN-ONIZE: A large-scale anonymous survey system. In *IEEE Symposium on Security and Privacy*, 2014.
- [81] I2P: The invisible internet project. https://geti2p.net/en/. Last accessed: June 20, 2017.
- [82] R. Janakiraman, M. Waldvogel, and Q. Zhang. Indra: A peer-to-peer approach to network intrusion detection and prevention. In 12th IEEE International Workshops on Enabling Technologies, 2003.
- [83] A. Johnson, P. F. Syverson, R. Dingledine, and N. Mathewson. Trust-based anonymous communication: adversary models and routing algorithms. In 18th ACM Conference on Computer and Communications Security, 2011.
- [84] A. Johnson, C. Wacek, R. Jansen, M. Sherr, and P. F. Syverson. Users get routed: traffic correlation on Tor by realistic adversaries. In 20th ACM SIGSAC Conference on Computer and Communications Security, 2013.
- [85] P. C. Johnson, A. Kapadia, P. P. Tsang, and S. W. Smith. Nymble: Anonymous IP-address blocking. In 7th Privacy Enhancing Technologies Symposium, 2007.
- [86] J. Juen, A. Johnson, A. Das, N. Borisov, and M. Caesar. Defending Tor from network adversaries: A case study of network path prediction. *PoPETs*, 2015.
- [87] M. Kallahalla, E. Riedel, R. Swaminathan, Q. Wang, and K. Fu. Plutus: Scalable secure file sharing on untrusted storage. In USENIX Conference on File and Storage Technologies, 2003.
- [88] A. Kapadia and N. Triandopoulos. Halo: High-assurance locate for distributed hash tables. In *Network and Distributed System Security Symposium*, 2008.
- [89] G. Karame, E. Androulaki, and S. Capkun. Doublespending fast payments in Bitcoin. In 19th ACM Conference on Computer and Communications Security, 2012.
- [90] R. Kumaresan and I. Bentov. How to use Bitcoin to incentivize correct computations. In 21st ACM SIGSAC

Conference on Computer and Communications Security, 2014.

- [91] L. Lamport. Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM*, 21(7), 1978.
- [92] B. Laurie. Certificate transparency. Queue, 2014.
- [93] S. Le Blond, P. Manils, A. Chaabane, M. A. Kaafar, A. Legout, C. Castellucia, and W. Dabbous. Poster: Deanonymizing BitTorrent users on Tor. In 7th USENIX Symposium on Network Design and Implementation (NSDI'10), 2010.
- [94] C. Lesniewski-Laas, B. Ford, J. Strauss, R. Morris, and M. F. Kaashoek. Alpaca: extensible authorization for distributed services. In 14th ACM Conference on Computer and Communications Security, 2007.
- [95] F. Lesueur, L. Mé, and V. V. T. Tong. An efficient distributed PKI for structured P2P networks. In 9th International Conference on Peer-to-Peer Computing, 2009.
- [96] N. Li, W. H. Winsborough, and J. C. Mitchell. Distributed credential chain discovery in trust management. *Journal of Computer Security*, 2003.
- [97] Z. Li, S. Alrwais, Y. Xie, F. Yu, and X. Wang. Finding the linchpins of the dark web: a study on topologically dedicated hosts on malicious web infrastructures. In *IEEE Symposium on Security and Privacy*, 2013.
- [98] Y. Liu and J. Pan. The impact of NAT on BitTorrent-like P2P systems. In 9th International Conference on Peer-to-Peer Computing, 2009.
- [99] J. Maheswaran, D. I. Wolinsky, and B. Ford. Crypto-book: an architecture for privacy preserving online identities. In Proceedings of the Twelfth ACM Workshop on Hot Topics in Networks, page 14. ACM, 2013.
- [100] J. McLachlan, A. Tran, N. Hopper, and Y. Kim. Scalable onion routing with Torsk. In 16th ACM Conference on Computer and Communications Security, 2009.
- [101] M. S. Melara, A. Blankstein, J. Bonneau, E. W. Felten, and M. J. Freedman. CONIKS: bringing key transparency to end users. In 24th USENIX Security Symposium, 2015.
- [102] B. Mitra, F. Peruani, S. Ghose, and N. Ganguly. Analyzing the vulnerability of superpeer networks against attack. In 14th ACM Conference on Computer and Communications Security, 2007.
- [103] P. Mittal and N. Borisov. Information leaks in structured peer-to-peer anonymous communication systems. In 15th ACM Conference on Computer and Communications Security, 2008.
- [104] P. Mittal and N. Borisov. ShadowWalker: peer-to-peer anonymous communication using redundant structured topologies. In 16th ACM Conference on Computer and Communications Security, 2009.
- [105] P. Mittal, M. Caesar, and N. Borisov. X-Vine: Secure and pseudonymous routing in DHTs using social networks. In 19th Network and Distributed System Security Symposium, 2012.
- [106] P. Mittal, F. G. Olumofin, C. Troncoso, N. Borisov, and I. Goldberg. PIR-Tor: Scalable anonymous communication using private information retrieval. In 20th USENIX Security Symposium, 2011.
- [107] P. Mittal, C. Papamanthou, and D. Song. Preserving link privacy in social network based systems. In 20th Network

and Distributed System Security Symposium(NDSS). Internet Society, 2013.

- [108] P. Mittal, M. K. Wright, and N. Borisov. Pisces: Anonymous communication using social networks. In 20th Network and Distributed System Security Symposium, 2013.
- [109] F. Monrose and S. Krishnan. DNS prefetching and its privacy implications: When good things go bad. In 3rd USENIX Workshop on Large-Scale Exploits and Emergent Threats, 2010.
- [110] S. J. Murdoch and R. N. M. Watson. Metrics for security and performance in low-latency anonymity systems. In 8th Privacy Enhancing Technologies Symposium, 2008.
- [111] S. J. Murdoch and P. Zielinski. Sampled traffic analysis by Internet-exchange-level adversaries. In 7th International Symposium on Privacy Enhancing Technologies, 2007.
- [112] S. Nakamoto. Bitcoin: A peer-to-peer electronic cash system, 2008.
- [113] A. Nambiar and M. K. Wright. Salsa: a structured approach to large-scale anonymity. In 13th ACM Conference on Computer and Communications Security (CCS, 2006.
- [114] A. Narayanan and V. Shmatikov. De-anonymizing social networks. In 30th IEEE Symposium on Security and Privacy, 2009.
- [115] A. Narayanan, V. Toubiana, S. Barocas, H. Nissenbaum, and D. Boneh. A critical look at decentralized personal data architectures. arXiv preprint arXiv:1202.4503, 2012.
- [116] M. A. U. Nasir, S. Girdzijauskas, and N. Kourtellis. Socially-aware distributed hash tables for decentralized online social networks. In *IEEE International Conference on Peer-to-Peer Computing*, 2015.
- [117] S. Nilizadeh, S. Jahid, P. Mittal, N. Borisov, and A. Kapadia. Cachet: a decentralized architecture for privacy preserving social networking with caching. In *Conference on emerging Networking Experiments and Technologies*, 2012.
- [118] A. Oram. Peer-to-Peer: Harnessing the power of disruptive technologies. O'Reilly, 2001.
- [119] B. Parno, A. Perrig, and V. D. Gligor. Distributed detection of node replication attacks in sensor networks. In *IEEE Symposium on Security and Privacy*, 2005.
- [120] T. Paul, A. Famulari, and T. Strufe. A survey on decentralized Online Social Networks. *Computer Networks*, 2014.
- [121] A. Pfitzmann and M. Hansen. Anonymity, unlinkability, unobservability, pseudonymity, and identity management – a consolidated proposal for terminology. Technical report, 2005.
- [122] Pluggable transports. https://obfuscation.github.io/. Last accessed: June 20, 2017.
- [123] J. Postel. IETF RFC 821 Simple Mail Transfer Protocol. https://www.ietf.org/rfc/rfc821.txt, 1982. Last accessed: June 20, 2017.
- [124] J. A. Pouwelse, P. Garbacki, J. Wang, A. Bakker, J. Yang, A. Iosup, D. H. J. Epema, M. J. T. Reinders, M. van Steen, and H. J. Sips. Tribler: A social-based peer-to-peer system. In 5th International workshop on Peer-To-Peer Systems (IPTPS), 2006.
- [125] T. Pulls, R. Peeters, and K. Wouters. Distributed privacypreserving transparency logging. In 12th ACM Workshop on Privacy in the Electronic Society, 2013.
- [126] M. A. Rajab, F. Monrose, and A. Terzis. On the effectiveness of distributed worm monitoring. In 14th USENIX

- [127] M. Raya, M. H. Manshaei, M. Félegyházi, and J. Hubaux. Revocation games in ephemeral networks. In 15th ACM Conference on Computer and Communications Security, 2008
- [128] Redecentralize.org. http://redecentralize.org/. Last accessed: June 20, 2017.
- [129] M. K. Reiter and A. D. Rubin. Crowds: Anonymity for web transactions. ACM Trans. Inf. Syst. Secur., 1998.
- [130] Y. Rekhter, T. Li, and S. Hares. A border gateway protocol 4 (BGP-4). Technical report, 2005.
- [131] Reproducible Builds Provide a verifiable path from source code to binary. https://reproducible-builds.org/. Last accessed: June 20, 2017.
- [132] R. L. Rivest and B. Lampson. Sdsi-a simple distributed security infrastructure. Crypto, 1996.
- [133] P. Rogaway and M. Bellare. Robust computational secret sharing and a unified account of classical secret-sharing goals. In 14th ACM Conference on Computer and Communications Security, 2007.
- [134] C. Rossow, D. Andriesse, T. Werner, B. Stone-Gross, D. Plohmann, C. J. Dietrich, and H. Bos. SoK: P2PWNED - modeling and evaluating the resilience of peer-to-peer botnets. In 2013 IEEE Symposium on Security and Privacy, 2013
- [135] J. M. Rushby. Design and verification of secure systems, volume 15. ACM, 1981.
- [136] P. Schaar. Privacy by design. Identity in the Information Society, 3(2):267-274, 2010.
- [137] S. Schiffner, A. Pashalidis, and E. Tischhauser. On the limits of privacy in reputation systems. In 10th ACM workshop on Privacy in the electronic society, 2011.
- [138] B. Schmidt, R. Sasse, C. Cremers, and D. A. Basin. Automated verification of group key agreement protocols. In 2014 IEEE Symposium on Security and Privacy, 2014.
- [139] M. Selimi and F. Freitag. Tahoe-LAFS distributed storage service in community network clouds. In 2014 IEEE Fourth International Conference on Big Data and Cloud Computing, BDCloud 2014, Sydney, Australia, December 3-5, 2014, pages 17-24, 2014.
- [140] S.-W. Seong, J. Seo, M. Nasielski, D. Sengupta, S. Hangal, S. K. Teh, R. Chu, B. Dodson, and M. S. Lam. PrPI: a decentralized social networking infrastructure. In 1st ACM Workshop on Mobile Cloud Computing & Services: Social Networks and Beyond, 2010.
- [141] A. Shamir. How to share a secret. Commun. ACM, 1979.
- [142] R. Sharma and A. Datta. SuperNova: Super-peers based architecture for decentralized online social networks. In 4th International Conference on Communication Systems and Networks, 2012
- [143] M. Sherr, M. Blaze, and B. T. Loo. Scalable link-based relay selection for anonymous routing. In 9th Privacy Enhancing Technologies Symposium, 2009.
- [144] R. Snader and N. Borisov. A tune-up for Tor: Improving security and performance in the tor network. In 15th Network and Distributed System Security Symposium, 2008.
- [145] E. Sparrow, H. Halpin, K. Kaneko, and R. Pollan. LEAP: A next-generation client VPN and encrypted email provider. In International Conference on Cryptology and Network Security, pages 176-191. Springer, 2016.

- [146] E. Stefanov and E. Shi. Multi-cloud oblivious storage. In ACM SIGSAC Conference on Computer and Communications Security, 2013.
- [147] I. Stoica, R. Morris, D. R. Karger, M. F. Kaashoek, and H. Balakrishnan. Chord: A scalable peer-to-peer lookup service for Internet applications. In SIGCOMM, 2001.
- [148] M. W. Storer, K. M. Greenan, E. L. Miller, and K. Voruganti. POTSHARDS: secure long-term storage without encryption. 2007.
- [149] R. Süselbeck, G. Schiele, P. Komarnicki, and C. Becker. Efficient bandwidth estimation for peer-to-peer systems. In IEEE International Conference on Peer-to-Peer Computing, 2011
- [150] P. F. Syverson, D. M. Goldschlag, and M. G. Reed. Anonymous connections and onion routing. In IEEE Symposium on Security & Privacy, 1997.
- [151] Taler: Taxable anonymous libre electronic reserve. https: //taler.net/. Last accessed: June 20, 2017.
- [152] C. Tang and I. Goldberg. An improved algorithm for tor circuit scheduling. In 17th ACM Conference on Computer and Communications Security, 2010.
- [153] A. Tran, N. Hopper, and Y. Kim. Hashing it out in public: common failure modes of DHT-based anonymity schemes. In ACM Workshop on Privacy in the Electronic Society, 2009
- [154] P. P. Tsang, M. H. Au, A. Kapadia, and S. W. Smith. Blacklistable anonymous credentials: blocking misbehaving users without ttps. In 14th ACM Conference on Computer and Communications Security, 2007.
- P. P. Tsang, M. H. Au, A. Kapadia, and S. W. Smith. [155] PEREA: towards practical TTP-free revocation in anonymous authentication. In 15th ACM Conference on Computer and Communications Security, 2008.
- [156] P. P. Tsang, A. Kapadia, C. Cornelius, and S. W. Smith. Nymble: Blocking misbehaving users in anonymizing networks. IEEE Trans. Dependable Sec. Comput., 2011.
- [157] E. Y. Vasserman, R. Jansen, J. Tyra, N. Hopper, and Y. Kim. Membership-concealing overlay networks. In 16th ACM Conference on Computer and Communications Security, 2009.
- [158] J. Verkamp and M. Gupta. Inferring mechanics of web censorship around the world. In 2nd USENIX Workshop on Free and Open Communications on the Internet, 2012.
- [159] C. Wacek, H. Tan, K. S. Bauer, and M. Sherr. An empirical evaluation of relay selection in Tor. In 20th Network and Distributed System Security Symposium, 2013.
- [160] M. Wachs, F. Oehlmann, and C. Grothoff. Automatic transport selection and resource allocation for resilient communication in decentralised networks. In 14th IEEE International Conference on Peer-to-Peer Computing, 2014.
- [161] M. Wachs, M. Schanzenbach, and C. Grothoff. A censorship-resistant, privacy-enhancing and fully decentralized name system. In 13th International Conference on Cryptology and Network Security, 2014.
- M. Waldman and D. Mazières. Tangler: a censorship-[162] resistant publishing system based on document entanglements. In 8th ACM Conference on Computer and Communications Security, 2001.
- [163] M. Waldman, A. D. Rubin, and L. F. Cranor. Publius: A robust, tamper-evident, censorship-resistant, and source-

anonymous web publishing system. In 9th USENIX Security Symposium, 2000.

- [164] L. Wang and J. Kangasharju. Measuring large-scale distributed systems: case of BitTorrent mainline DHT. In 13th IEEE International Conference on Peer-to-Peer Computing, 2013.
- [165] Q. Wang, Z. Lin, N. Borisov, and N. Hopper. rBridge: User reputation based Tor bridge distribution with privacy preservation. In 20th Network and Distributed System Security Symposium, 2013.
- [166] Q. Wang, P. Mittal, and N. Borisov. In search of an anonymous and secure lookup: attacks on structured peer-to-peer anonymous communication systems. In 17th ACM Conference on Computer and Communications Security, 2010.
- [167] X. Wang, S. Chen, and S. Jajodia. Tracking anonymous peer-to-peer VoIP calls on the Internet. In 12th ACM Conference on Computer and Communications Security, 2005.
- [168] B. Wilcox-O'Hearn. Experiences deploying a large-scale emergent network. In *International Workshop on Peer-to-Peer Systems*, pages 104–110. Springer, 2002.
- [169] M. Winslett, C. C. Zhang, and P. A. Bonatti. PeerAccess: a logic for distributed authorization. In 12th ACM Conference on Computer and Communications Security, 2005.
- [170] E. Wobber, M. Abadi, M. Burrows, and B. Lampson. Authentication in the taos operating system. ACM Transactions on Computer Systems (TOCS), 12(1):3–32, 1994.
- [171] E. Wobber, M. Abadi, M. Burrows, and B. W. Lampson. Authentication in the Taos operating system. In 14th ACM Symposium on Operating System Principles, 1993.
- [172] S. Wolchok, O. S. Hofmann, N. Heninger, E. W. Felten, J. A. Halderman, C. J. Rossbach, B. Waters, and E. Witchel. Defeating Vanish with low-cost sybil attacks against large DHTs. In *Network and Distributed System Security Symposium*, 2010.
- [173] D. I. Wolinsky, H. Corrigan-Gibbs, B. Ford, and A. Johnson. Dissent in numbers: Making strong anonymity scale. In 10th USENIX Symposium on Operating Systems Design and Implementation, 2012.
- [174] M. K. Wright, M. Adler, B. N. Levine, and C. Shields. An analysis of the degradation of anonymous protocols. In *Network and Distributed System Security Symposium*, 2002.
- [175] M. K. Wright, M. Adler, B. N. Levine, and C. Shields. The predecessor attack: An analysis of a threat to anonymous communications systems. ACM Trans. Inf. Syst. Secur., 2004.
- [176] J. J. Wylie, M. W. Bigrigg, J. D. Strunk, G. R. Ganger, H. Kiliççöte, and P. K. Khosla. Survivable information storage systems. *IEEE Computer*, 2000.
- [177] YaCy: The Peer to Peer Search Engine. http://yacy.net/ en/index.html. Last accessed: June 20, 2017.
- [178] B. Yang and H. Garcia-Molina. PPay: micropayments for peer-to-peer systems. In 10th ACM Conference on Computer and Communications, 2003.
- [179] youbroketheinternet. http://youbroketheinternet.org/. Last accessed: June 20, 2017.
- [180] M. Young, A. Kate, I. Goldberg, and M. Karsten. Practical robust communication in DHTs tolerating a Byzantine adversary. In *ICDCS*, 2010.

- [181] H. Yu, P. B. Gibbons, M. Kaminsky, and F. Xiao. Sybil-Limit: A near-optimal social network defense against Sybil attacks. *IEEE/ACM Trans. Netw.*, 2010.
- [182] H. Yu, P. B. Gibbons, and C. Shi. DCast: sustaining collaboration in overlay multicast despite rational collusion. In 19th ACM Conference on Computer and Communications Security, 2012.
- [183] D. J. Zage and C. Nita-Rotaru. On the accuracy of decentralized virtual coordinate systems in adversarial networks. In 14th ACM Conference on Computer and Communications Security, 2007.
- [184] E. Zhai, R. Chen, Z. Cai, L. Zhang, E. K. Lua, H. Sun, S. Qing, L. Tang, and Z. Chen. Sorcery: Could we make P2P content sharing systems robust to deceivers? In 9th IEEE International Conference on Peer-to-Peer Computing, 2009.
- [185] E. Zhai, D. I. Wolinsky, R. Chen, E. Syta, C. Teng, and B. Ford. Anonrep: Towards tracking-resistant anonymous reputation. In 13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16), pages 583–596. USENIX Association, 2016.
- [186] X. Zhang, H.-C. Hsiao, G. Hasker, H. Chan, A. Perrig, and D. G. Andersen. SCION: Scalability, control, and isolation on next-generation networks. In *IEEE Symposium on Security and Privacy*, 2011.
- [187] B. Zhu, S. Setia, and S. Jajodia. Providing witness anonymity in peer-to-peer systems. In 13th ACM Conference on Computer and Communications Security, 2006.
- [188] S. Zhu, S. Setia, and S. Jajodia. LEAP+: Efficient security mechanisms for large-scale distributed sensor networks. *TOSN*, 2006.
- [189] G. Zyskind, O. Nathan, and A. Pentland. Enigma: Decentralized computation platform with guaranteed privacy. *CoRR*, abs/1506.03471, 2015.

EXHIBIT H



The encrypted threat: Bitcoin's social cost and regulatory responses



By Ulrich Bindseil, Patrick Papsdorf and Jürgen Schaaf¹ European Central Bank

JEL codes: E42, G10, G18, G28. Keywords : Bitcoin, crypto-assets, illicit payments, proof of work.

While Bitcoin raised the attention for the potential of distributed ledger technology (DLT), it fails to deliver on its promises but comes at high costs. It is unfitted and inefficient as a means of payment but used extensively for illicit activities. It is unsuitable as an investment asset and neither empowers, nor relieves the sovereign individual from the state. While so far authorities seemed to have insufficiently addressed the negative effects of Bitcoin for society, this is eventually changing. Illicit usage will be further hindered, and compliance costs added to the Bitcoin ecosystem. Likewise, growing concerns on Bitcoin's climate footprint have now led to calls of some authorities to address or even ban essential elements of Bitcoin's technology. Nevertheless, Bitcoin has reached new valuation records in November 2021, maybe also because of perceived or actual supportive legislative measures facilitating investment inflows into Bitcoin. As it is difficult to find arguments supporting the sustainability of Bitcoin, and as the social fall-out of its collapse would be significant, authorities should (1) strengthen global implementation of AML/CFT standards and broaden measures to stop Bitcoin being a vehicle for illicit purposes; (2) avoid measures that invite additional investment flows into Bitcoin.

¹ Views expressed in this paper are the ones of the authors and not necessarily the ones of the ECB. We would like to thank Fiona van Echelpoel, Anton van der Kraaij, Mirjam Plooij and Pedro Miguel Bento Pereira Da Silva for useful comments. Thanks to Ines Rossteuscher and Pierfrancesco Zeoli for their research assistance. All remaining errors are ours.

1. Introduction

In November 2021, the market capitalisation of crypto assets exceeded for the first time USD 3 Trillion, of which around USD 1.3 trillion were contributed by Bitcoin (see Figure 1). This article restates the reasons why the observed Bitcoin valuation is unlikely to be sustainable. Moreover, it emphasises that, even if financial stability risks of a Bitcoin collapse might be contained, the Bitcoin life cycle will likely have implied painful losses for many retail Bitcoin investors and a significant enrichment for early investors who liquidate their position in time. Beyond the negative effects of a perceived unjustified redistribution of wealth, Bitcoin will have represented a significant negative-sum game as it will have come with large costs in the form of hardware investments and energy consumption. The article therefore concludes that public authorities should not contribute to scale up the eventual damage of Bitcoin to society. Instead they should, first, treat the Bitcoin network as rigorously as the conventional financial industry in terms of prevention of illicit payments, money laundering and terrorist financing, second, address the negative externalities of Bitcoin's energy consumption, and third, deny recognition of Bitcoin as an investment and not allow it to become incrementally part of the regular financial system without strictest safeguards. In the rest of this introduction (section 1), we will briefly recall the origins and the principles of the functioning of the Bitcoin network. Section 2 turns to the vulnerability and inefficiency of the Bitcoin technology. Section 3 explains why Bitcoin is not a suitable means of payment, and section 4 why it is neither an investment asset. Based on sections 2-4, section 5 concludes that Bitcoin is unlikely to be sustainable. Section 6 argues that contrary to one common narrative, Bitcoin does not help the sovereign individual to regain its liberty, and section 7 recalls the misuse of Bitcoin for criminal activities. Section 8 explains how all these issues can be mapped into private and social costs of Bitcoin and concludes that the net welfare effects of Bitcoin over its life cycle will have been significantly negative. Section 9 turns to recent measures by regulators and public authorities, noting that the latter are becoming tougher on Bitcoin's use for illicit payments and its other shortcomings, while some ambiguous regulatory measures facilitate Bitcoin's recognition as an investment asset.



Figure 1: Market capitalisation of selected crypto-assets

Sources: Coincodex, TradingView and authors' calculations

As summarised for example in Schär and Berentsen (2020), in 2007 a group of software developers invented a completely decentralized booking concept. Under the pseudonym Satoshi Nakamoto, a white paper, and the source code for a "digital cash" were published (Nakamoto 2008); in January 2009, the first fifty Bitcoin were generated. To date, the identity of Satoshi Nakamoto has not been disclosed. The Bitcoin system allows its holders to be anonymous through encryption, although the Bitcoin blockchain² is transparent in terms of what addresses hold which amounts of Bitcoin and on the related transaction flows. In addition, transactions are considered irreversible, regardless of the reason.

An overview of the functioning and governance of the Bitcoin network is provided e.g. by Böhme et al. (2015). The underlying technology and the conceptional setup can be summarized as follows: There is no central authority, but a global network of computers controls, monitors, and stores the system information. New Bitcoins are coined by decentralized "mining" by users and their computers. New data packets are added to the blockchain every few minutes. The maximum total number of Bitcoins is technically limited to about 21 million, of which just under 19 million are already in circulation. When this limit is reached – the transaction fees become the only source of income for the miners, on whose existence Bitcoin depends in the long run. To prove the correctness of the entire blockchain and its extensions, computers must solve a mathematical puzzle for each block. The so-called miners validate the transactions by entering them into a public ledger. Currently rewards include transaction fees as well as seignorage from newly created Bitcoins, i.e. the market value of a bitcoin minus the mining costs.

This proof-of-work method has a scalable difficulty level and aims to keep the incentive for miners to keep running the system sufficiently high. The more computing capacity and the faster the validation process takes place, the safer the whole system will be. Such dynamic and decentralized protection leads to an exponential increase in the power demand of the computers, which means a huge energy demand for the system. Bitcoin's price directly affects the value of the mined coins and therefore the amount of resources miners can afford to spend on mining (see e.g. the simple model of de Vries, 2021). With a higher Bitcoin price, more producers are incentivised to compete for new coins. This in turn requires to make the encryption puzzle more difficult. By consequence, the miners will require more electricity to solve the puzzle and will consume more electricity and increase carbon emissions.

While some technological development occurred in the blockchain since its inception in 2009 through forks and upgrades (e.g. Segwit, Lightning Networks and Taproot), which try to address some aspects like scalability and cost, it remains that the Bitcoin blockchain itself implies the above shortcomings.

2. Vulnerability and inefficiency of the Bitcoin technology

The durability, stability and scalability of the Bitcoin network is noteworthy. Moreover, as stated e.g. by Auer (2021), blockchain and the distributed ledger technology are rapidly becoming an industry standard for digital assets and in other applications. The entire potential of these technologies has still not fully been explored.

Still, several authors have raised serious doubts on Bitcoin's underlying technology and concept (for example Taleb, 2021; Avoca 2021; Acemoglu, 2021; Kolbert, 2021). The proof-of-work concept, which is a constituting feature of the Bitcoin system, is generally recognised as cumbersome and slow: it can only handle seven to ten transactions per second. This results in long transaction processing time as found by Avoca (2021).

² Blockchain is a sub-category of the distributed ledger technology (DLT). The various DLT concepts differ mainly in how transactions are validated and stored.

For comparison: The Visa network is said to be able to process an estimated 24,000 transactions per second (Avoca, 2021, p.4), i.e. the scalability and efficiency of well-designed conventional centralised payment systems is far less constrained.

It may also be noted that slow and opaque pricing networks have traditionally attracted predatory highfrequency algorithm traders and are vulnerable to related market stress. The flash crash of 6 May 2010 was a point in case (although unrelated to Bitcoin). As Baqer (2016) showed Bitcoin itself has suffered from attacks by high frequency trading firms, too. Avoca (2021) stress that the Bitcoin network is also vulnerable because of its reliance on a single security technology that experts consider to be outdated by advances in computing. Bitcoin uses the secure hash algorithm (SHA) which is more than twenty years old. While the U.S. Department of defence and many leading IT firms like Microsoft found the SHA-1 standard too weak for cyber-protection and decommissioned its use in the early 2010s. Researchers believe that the technology will not be able to keep up in a quantum computing environment. In the absence of a central legitimized management it is hard to see how the fundamental security technology could be replaced to withstand the challenges of future technological advances of others.

The Bitcoin network has also been reported for a long time to have another technical vulnerability of conceptional nature. It is prone to a so-called 51 percent attack, which occurs when miners (potentially malicious) gain control of more than 51 percent of the network's hash-rate: they could then issue coins twice. While Bitcoin is in principle less exposed to the risk of a 51 percent attack because of its vast network of 1,000 nodes, a problematic concentration would actually have occurred in 2014: In June 2014, the mining pool GHash.IO reached a share of about 55 percent of the Bitcoin hashrate over 24-hours. Although a month later GHash.IO's share of the network's hashrate had dropped to just over 38 percent, the risk remained that a single miner or mining pool could again take control. GHash.IO voluntarily committed to stay far below 40 percent (Hern, 2014).

Moreover, the Bitcoin network is already now increasingly run by supercomputers and server farms and the incentive structure of retail miners might take a hit once all Bitcoins are minted and the reward system will rely on fees only. In consequence, the hash-rate is not unlikely to be increasingly concentrated in the hands of a few.

As importantly, the Bitcoin network comes with **a large energy hunger** due to its reliance on proof-of-work (see figure 2). It wastes power and is therefore an immense environmental polluter. The reason is the power demand of the proof-of-work concept - which is a necessary condition for the security of the system. According to the Cambridge Centre for Alternative Finance Bitcoin computers use around 140 terawatt hours of electricity per year - about a quarter of Germany's electricity consumption. Digiconomist (2021) estimates that the entire Bitcoin network consumes 201.894 TWh per year. This would be close to the amount of energy all data centres consume globally. The consumed energy further results in 95.9 metric tons of CO2, comparable to the carbon footprint of metropolitan London. The more energy the Bitcoin network uses, the more secure it is. A lower energy demand of the Bitcoin system is therefore neither expected nor desired – rather, Bitcoin is sometimes justified by the fact that it would on balance be beneficial for planet earth and humanity as argued e.g. by Vukolic (2021). And even if alternative sources of energy were used or disused power plants revived, the network would still waste energy that could be used for other purposes, as convincingly argued recently by the Swedish Financial Supervisory Authority and the Swedish Environmental Protection Agency (2021).

Figure 2: estimated energy consumption of the Bitcoin network



3. Bitcoin is not a currency

Nakamoto (2008) presented Bitcoin as useful for society through its payment function, but his related arguments were already rather unclear at that time. There is today in any case a broad consensus that Bitcoin fails in its original objective of being a currency. Bitcoin is too volatile to fulfil the classic functions of money: unit of account, means of payment, store of value (see figure 3 illustrating the exceptional volatility of Bitcoin). Moreover, the system is too slow and expensive to compete with established payment systems and currencies. Incentivizing system maintenance without central authority is challenging and expensive. The lack of acceptance by merchants due to long settlement times and high fees (currently between USD 2,5 and 4 per transaction) already shows that Bitcoin cannot be understood as a means of payment outside of niches. Therefore, Bitcoin's business model as a global means of payment is not plausible.

The latest attempt to make the vision of Nakamoto (2008) reality on a larger scale was El Salvador trying to introduce Bitcoin as a second legal tender alongside the US Dollar on 7 September 2021. The launch was bumpy largely because there was no popular acceptance of the new means of payment. On the day of introduction, the Bitcoin exchange value plummeted by 15 percent, accompanied by protests targeted against President Nayib Bukele as reported by BBC News (2021).

Nevertheless, the number of Chivo Bitcoin wallets has expanded to more than 4 million. This however might be related to the USD 30 (its equivalency in BTC) given by the government to Salvadoran citizens to download the Chivo wallet as suspected by Fitch (2021). It is also important to note that payments through the Chivo wallet are actually layered and not settled in the Bitcoin network. Instead, they are just internally settled by the wallet provider, who acts as custodian (Merten, 2021). Therefore, at best, the Chivo Wallet is a payment system backed by Bitcoin, but fully betraying the idea of Nakamoto of overcoming the dependence of payments on centralised intermediaries, even if this betrayal has good reasons (the Bitcoin network being too slow, insufficiently scalable, and too costly for payments). Whether Chivo Wallets are fully backed or possibly underfunded is not fully transparent (although there is no indication that they are underfunded).

In the meantime, President Bukele has launched new plans to push Bitcoin's use and mining in El Salvador with a new city built around a Bitcoin industry. The construction financing and maintenance of the "Bitcoin City" would be based on new Bitcoin bonds; and the required energy taken from a volcano in the proximity.³





³ See e.g. Reuters, 22 November 2021, "El Salvador plans first Bitcoin City, backed by bitcoin bonds", by Nelson Renteria

4. Bitcoin does not appear to be a sustainable investment

One of the most popular arguments among Bitcoin supporters is that the limited supply of Bitcoin would make it a great asset to protect investors against inflation, while fiat money, which can be multiplied at will, would increasingly lose value.

However, even if one were to assume that Bitcoin could become the new global money, its technically fixed "money supply" would turn out to be a weakness on closer inspection: the world would be led into a deflation trap in a growing economy. In a deflation, falling prices of goods and services tempt citizens to postpone less urgent purchases into the future. This is reasonable for individuals, but aggregate demand suffers which slows down the economy. ⁴

The advocates of gold as a weapon against inflation – and those who praise Bitcoin for the same as reason as the new gold - should remember the reasons for the abolition of the gold standard. While the gold peg could indeed offer protection against inflation, the flip side is the above-mentioned increased risk of deflation: In 1931 major currencies gave up the gold peg after years of painful recession, deflation, and financial instability.

Similarly, the indirect gold standard of the Bretton Woods monetary system after the end of World War II failed. During that time currencies were no longer tied directly to gold but to the US Dollar (at a fixed parity of 35 US Dollar per ounce of gold). The reason for the failure was that the U.S. could not keep money tight enough to maintain the gold parity as credible and at the same time provide the dynamically growing world economy with sufficient liquidity.

But the often-used comparison to gold also fails for more basic reasons. As Taleb (2021) argues, gold is both used industrially and has been appreciated as jewellery for centuries before it became a store of value, an investment asset, or a reserve currency. Moreover, it does not degenerate over time and retains its value even in chaotic or degenerative states of the world like natural catastrophes or in the case of a temporary or lasting failure of the electric or digital infrastructure.

Finally, the objection that the fiat money of modern central banks also has no intrinsic value falls short: because in deliberately moving away from the gold standard, sovereigns and central banks have put in place clearly defined mandates, legal guarantees, institutional and operational arrangements (independence as well as loans against collateral) to be able to release the gold brake without losing stability (see e.g. Bindseil and Fotia, 2021, 103-107).

Last but not least, the alternative of Bitcoin as a store of value is not predominantly central bank fiat money, but the financing through equity and/or debt of real economic projects which serve needs of society and generate a cash flow which allows positive yields to be sustained, nchoring the value of the investment assets in its real productivity. Investors' worries about the stability of certain fiat currencies can be legitimately expressed by allocating their wealth into equity, commodities, real estate, human capital, or other productive assets. As Adam Tooze formulates, fiat money is backed by "nothing' other than the trifling matter of tens of trillions of dollars in private credit, the rule of law and the power of the state, itself inserted into a state system. In other words, the entire structure of global macrofinance" (Tooze, 2021).

⁴ Probably the alleged inflation protection provided by a fixed Bitcoin supply is illusory: the number of possible crypto assets that can rival Bitcoin is, after all, unlimited.

Some have also argued that the spike in Bitcoin valuation is due to the low interest rate policies of central banks. These would force investors to seek yield as assumedly offered by Bitcoin, a sort of digital commodity that would be able to escape from financial repression. While the wish for high nominal and real yields and the dissatisfaction about reality in many advanced economies is comprehensible and legitimate, it should not be a reason for shifting savings into highly speculative investments. If central banks were setting excessively low nominal interests rates, i.e. not well justified by monetary policy considerations, then investors should seek to fund real assets with sustainable values because of a proven contribution to the needs of society, and should at the same time support policy makers that commit to take measures supporting real economic growth and thereby real rates of returns of the capital stock of the economy.

5. Mounting doubts about the sustainability of Bitcoin

Because Bitcoin is neither efficient nor suitable as a means of payment, it is not competitive for legal payments. Moreover, Bitcoin has no intrinsic value and does not generate a cash flow or dividends. Hence, the market valuation of Bitcoin is purely based on speculation. As Diehl (2021) puts it: "[...] crypto morphed into a pure speculative mania which attracted a fanatic quasi-religious movement fuelled by gambling addiction and the pseudo-intellectual narrative economics of the scheme." This market rally only works as long as the Bitcoin community's beliefs about Bitcoin's alleged advantages as a means of payment or that the market value can rise forever can be maintained. The Bitcoin hype has all the characteristics of a speculative bubble along the so-called greater fool theory (Oxford Business Review, 2020). Accordingly, the value rises if there is still a "greater fool" who assumes he can sell at an even higher price later. But much like the number of Bitcoins is ultimately limited, "eventually, one runs out of greater fools" (Malkiel, 1973).

The enthusiasm for Bitcoin alone is not enough in the long run, especially as Bitcoin is in the end only a number chain and technologies are replaced by better technologies; with the newer soon displacing the new. In fact, Bitcoin remains the dominant crypto-asset but its market share has declined sharply in 2021 from more than 70 percent to less than 45 percent. Market interest has grown for newer blockchains that use smart contracts and aim to solve the challenges of earlier blockchains by introducing features to ensure scalability, interoperability, and sustainability. The biggest among the newer crypto-assets is Ether, which surpassed Bitcoin trading volumes earlier in 2021 (IMF, 2021). Finally, Bitcoin's stellar long-term market performance that continues to attract investors was largely contingent on the timing of the initial investment. According to Wewel (2021) crypto returns do not even deviate markedly from traditional assets on a risk-adjusted basis, which is attributable to their substantially higher volatility.⁵

6. The illusion of liberation

For all its economic shortcomings, there remains the vision of Bitcoin to restore freedom from government control and from centralized entities that abuse their power. Bitcoin, with its decentralized organization, promises the emancipation of the individual and the ultimate democratization of the monetary system as stated in Omarova (2021). However, even the case of Bitcoin freedom needs rules, otherwise there is a threat of anarchy and the law of the strongest. The fact that the economy and financial markets in developed market economies are not purely decentralized and spontaneously organized, but rely on central institutions and distributed nodes with internal hierarchies (firms) and within set rules, has long been recognized in economic literature, at least since the work of Nobel laureates Ronald Coase (1937) and Oliver E. Williamson (1975). Firms and incomplete

⁵ This critical assessment on Bitcoin is obviously not implying a more favourable verdict on other crypto-assets.

contracts help deal with uncertainty and complexity and reduce transaction costs - and are by no means secondbest solutions in the absence of appropriate technologies. Mechanistic rules, as Bitcoin appears to create, are not an appropriate solution for a changing world. Therefore, the recent more ambitious attempt to make Bitcoin a means of payment unavoidably betrayed its libertarian principles, including the core idea of Nakamoto (2008) to overcome the role of central payment intermediaries.

Bitcoin is also by no means as grassroots democratic as its community may have believed, at least in the early days, but is shaped by financial interests and powerful shareholders and, relatedly, the exposure to concentration risks, given its large reliance on a few entities, like custodial wallets and exchanges (for example, Binance handles more than half of trading volumes according to the IMF (2021)). The majority, 75 percent of the addresses, holds just over 0.2 percent of the market share; the hundred largest Bitcoin shareholders hold more value than the smallest 38 million combined (Dunn, 2021; although these numbers may be impacted by exchanges and wallet providers holding "omnibus accounts" for small holders).

Finally, Bitcoin offers a vision of a global means of payment without national jurisdictions to overcome borders quite unlike conventional cross-border payments. People could send value across borders for free and unhindered to anyone with a Bitcoin wallet. This view ignores that the high cost of conventional cross-border payments is not only due to the inefficiency of payment instruments, but in significant part to costs of market and liquidity risk management and regulatory requirements to combat money laundering and terrorist financing. However, the cost of complying with these requirements, and provisions for legal and exchange rate risks only affect the regulated financial sector. The fact that some bitcoin transactions, e.g. like peer-to-peer, have been able to escape this entirely so far is a regulatory gap, not a technological achievement. It is however not denied that the area of cross-border payments needs improvements in terms of cost, speed, transparency and inclusiveness. The authorities, in the form of the Financial Stability Board (FSB, 2020) published an ambitious roadmap to enhance cross-border payments in October 2020 that is being thoroughly followed up.

7. The use of Bitcoin for illicit activities

Bitcoin has been successful as payment means for criminal usages. A distinction must be made between market manipulation and dubious activities of exchange operators and the use for money laundering and drug trafficking, terrorist financing and extortion and ransom below the radar of law enforcement and regulatory authorities.

Dunn (2021) presents a long list of shady operators and market manipulation that have marked Bitcoin's history on the supply side. The first bubble in 2013 was fuelled by the Mt Gox exchange which hosted about 70 percent of Bitcoin trading. The exchange lost 650,000 Bitcoins of its users and went bankrupt. Studies by Gandal et al. (2021) suggest that the first boom – a rise from USD 100 to USD 1,000 in just two months - was due to manipulation of a trading software.

Griffin (2019) found that the second and third booms were associated with the launch and rise of Tether. Tether is a so-called stablecoin, i.e. a type of crypto-asset that aims to maintain a stable value by being backed by fiat currency or other assets. Tether is, according to the issuer, nominally pegged one-to-one to the US Dollar and is backed entirely by cash-like assets. Griffin's investigations during the 2017 boom suggested that 50 percent of the sharp price increase was due to manipulation with Tether.

Bitcoin has also been popular for financing criminal activities. Drug trafficking, money laundering, terrorist financing and extortion are the most popular areas of use. In recent years, exit scams have dominated crypto-

assets crimes according to Cybertrace (2021). For instance, in May 2021, a ransomware attack on a US energy pipeline was carried out by a group operating out of Russia who received ransomware payment of approximately USD 5 million in Bitcoin (Reuters, 2021d), in June 2021 the meat company JBS paid USD 11 Million in Bitcoin to avoid further disruptions (WSJ, 2021), in March 2021, another attack made CNA Financials pay USD 40 million in Bitcoin to get data back (Simpson, 2021). Suspicious activity reported on ransomware attacks accounted only in the first half of 2021 for 590 USD million in the U.S. with crypto assets being the vehicle for ransom payments (Fincen, 2021). While this amount of ransom attacks may present a limited share of global money laundering activity, it is certainly strongly growing and of high concern, also due to the high indirect cost an damage, e.g. through operational disruptions or confidential data leakages. In addition, Bitcoin is also one of the main crypto-assets used in the Darknet (65 percent in Q1 2020) (Crystal, 2020).

The share of illicit payments in total Bitcoin transactions are disputed: While Foley (2019) estimates that some 45 percent are for illegal use, the Chainalysis' 2021 crypto crime report finds less than 1 percent for 2020. As suggested by Green (2021), such small ratio could be because the denominator confuses trade volume (mostly relating to investment flows) with payments matching an economic transaction. FATF (July 2021) reports variations in identified illicit Bitcoin transactions from 2016 - 2020 to range between 0.6 and 9.9 percent (in proportion to the number of transactions) and 0.1 and 5.1 percent (in proportion to the USD value of transactions). Moreover, it was also found that in total, illicit transactions were identified to occur typically without an intermediary (wallet provider or exchange). Bitcoin's design attracts illicit usages as it allows to hide identities, to transact entirely within the darknet or on-chain without reliance on regulated entities, to use mixing services to obscure the trail of a transaction or to use exchanges that have not yet adopted the AML/CFT standards of FATF (Crystal, 2020).

However, Bitcoin's set-up can support forensic analysis in tracing illegal activities as transactions never disappear from the blockchain. While this may allow at times to recover some of the paid ransom, it remains a complex, time-consuming and disproportionate exercise, as the US Department of Justice has revealed in recent years. ^{6,7}

8. The high private and social cost of the Bitcoin network

The longer the boom lasts and the more money flows into the system before the music may stop, the higher are the risks and costs for invested individuals and the society at large. Often, different kinds of social costs of the Bitcoin network are not well-distinguished in the debate. Consider the following issues:

1. Bitcoin comes with significant **private costs** in the form of high energy and hardware consumption of the Bitcoin network. If it were true that Bitcoin is eventually unsustainable and will not persist, and will not have generated value for society apart from temporary hopes of speculative gains which eventually are disappointed, then these private costs will however have represented a net loss for society. This argument holds regardless of the potential negative externalities of energy consumption.

⁶ See Department of Justice Journal of federal law and practice, Forfeiting Cryptocurrency: Decrypting the Challenges of a Modern Asset Neal B. Christiansen Assistant United States Attorney Western District of Washington Julia E. Jarrett Assistant United States Attorney District of Oregon, September 2019.

⁷ See US Department of Justice "Department of Justice Seizes \$2.3 Million in Cryptocurrency Paid to the Ransomware Extortionists Darkside", 7 June 2021.

- 2. The question arises if the **negative externalities** of energy consumption are really priced in through adequate taxes. Geographical arbitrage of Bitcoin mining will lead to a further concentration of mining in locations where this is least the case, such that the internalization of negative externalities can be escaped.
- 3. Some have argued to **locate Bitcoin mining to locations where energy is quasi free** and therefore leaves no CO2 footprint. For example, El Salvador envisages to build a "Bitcoin City" close to a Volcano and use its energy. Similarly, Iceland has for long attracted mining operations with its abundance of cheap geothermal energy before its national energy company decided in December 2021 to cut power to new Bitcoin miners (Cointelegraph, 2021). The question arises why such a simple solution would not attract any other energy intensive activity with a priori limited geographical constraints? Moreover, the energy consumption of the Bitcoin network is inversely proportional to the cost of energy. This means that if mining farms move massively to areas where energy is cheaper, then the logic of the proof of work mechanism requires that more energy will be consumed in mining for a given price of Bitcoin.
- 4. The high social cost of Bitcoin and its **negative net social value** is currently not perceived by Bitcoin investors who believe them to be **covered by current and future speculative gains**. However, pure speculative gains are not a basis for sustainable price increases, and therefore the bill for the private costs of the Bitcoin network will eventually be paid. The full social cost settlement is due once the music would stop playing and the Bitcoin valuation would have collapsed.
- 5. A significant additional component of the ultimate social costs will be the **societal damage when many will have realized that they lost their hard-earned savings for the benefits of smarter Bitcoin investors who bought at low and sold at high prices**. Those who lost money, in particular retail investors who naively put a large share of their eggs in the crypto basket will not appreciate the huge welfare redistribution at their expense and will put into question the functioning of society which permitted such unfairness to happen. While the celebrities of the system can withdraw themselves from the centerstage, societal consensus and trust takes another hit. The bigger the burned market value will be, the more dramatic the social backslash will be.

On balance, societies will eventually have to write off the cumulative energy consumption (including unpriced negative externalities), investment costs of hardware, the built-up human capital of the Bitcoin ecosystem, the cumulated work and a good dose of societal consensus. Moreover, in the meantime the Bitcoin network will have facilitated criminal activities by providing a means of illicit payments. All these costs are broadly proportional to the market capitalization that Bitcoin will reach and are moreover driven by the overall duration of the Bitcoin cycle. McCauley, (2021) concludes from similar observations that Bitcoin is a negative-sum game for society even worse than a Ponzi scheme.

9. Regulatory mindset is changing

The broad use of Bitcoin for illicit activities was recognised early. The shutting down of the darknet illicit marketplace Silkroad in 2013 (Time, 2013) revealed the extensive use of Bitcoin for illicit purposes – just five years after the white paper of Satoshi Nakamato was published. In 2014, the money laundering and terrorist financing related to crypto assets started to be picked up by the FATF (2014) and in 2019, it issued its guidance for a "Risk-Based Approach for Virtual Assets and Virtual Asset Service Providers" demanding national implementation and enforcement (FATF, 2019). If fully and consistently implemented and enforced, it means that providers of services in crypto assets would apply AML/CFT measures, including customer due diligence or the checking and reporting of suspicious transactions. As a result, the illicit usage of Bitcoin would become much more difficult in particular when exchanging it into fiat currency or using it for purchases of goods and services. Despite the advancements in the field of AML/CFT, regulators have been somewhat slow in addressing the above mentioned most obvious societal problems of Bitcoin.

Several explanations may be considered for this. *First*, the potential development of social risks may have been underestimated because of the relatively small size and unleveraged nature of the crypto assets market, which was assessed to not represent a fundamental threat to global financial stability. Second, regulatory responsibilities for Bitcoin seem somewhat fragmented as it raises multi-facetted threats and involves multiple actors. Moreover, the risks and concerns related to Bitcoin were first mainly related to money laundering and terrorist financing, while ransomware attacks occurred more recently, and with the surge in Bitcoin activities led to consumer and investor protection concerns. *Third*, many aspects of Bitcoin are fundamentally new and difficult to comprehend. Furthermore, they do not easily fit into existing regulation and raise regulatory challenges: Bitcoin operates borderless, misses a national anchor and was not perceived as a legal entity that could be addressed by regulation and incrimination (ECB, 2019). Also, regulators need to seek the right timing and suitable design of financial regulation to address risks and avoid gaps as well as unintended consequences, like stifling innovation (Warren, 2021). Given the global nature of Bitcoin, global cooperation amongst regulators is of importance to avoid regulatory gaps and arbitrage, as pointed out by IMF (2018), which is a time-intensive process. And it is not atypical that once a need for regulation has been identified, it can take years until regulation is finalised and applied. Fourth, the vested interests of large Bitcoin holders and financial intermediaries seeking for investment and business opportunities might have led to increased lobbying activities. The Economist (2021) warns that crypto lobbying was going ballistic, as companies were hoping to influence where the rules end up while regulators were toughening up their approaches.

In light of the continued reporting on illicit usages and climate implications, growth of the crypto-asset markets and its increased integration with financial markets] the threats are recognised (ECB, 2019, FSB, 2021, Cunliffe, 2021) and more accentuated calls for addressing the risks of crypto-assets are made (ECB, 2019; Lagarde, 2017). In a speech on 10 December 2021, Panetta (2021) has been explicit *that "the value of crypto-assets is growing rapidly and currently stands at over 2,500 billion dollars*'. *That is a significant figure with the potential to generate risks to financial stability that shouldn't be underestimated."*

Moreover, a number of jurisdictions have taken or are preparing measures to regulate Bitcoin alongside other crypto-assets. The spectrum of regulatory approaches is hereby wide reaching from criminalising crypto-asset business to more inclusive approaches of licensing and supervising intermediaries.

Some jurisdictions have banned Bitcoin (and similar crypto-assets), e.g.: In December 2021, Reuters (2021a) reported that the **Indian government** is considering prohibiting crypto-asset activities of individuals including a use as store of value, unit of account or means of transfer with violations by individuals being possibly sanctioned by arrests without bail options. Notably, reportedly the bill would also include non-custodial wallets, an area of the Bitcoin network that is largely unregulated. However, the bill has not yet been presented to the Parliament (Business Insider India, 2021). In November, the religious leaders in **Indonesia**, the National Ulema Council (MUI), have forbidden Muslims (almost 90 percent of the population) to use Bitcoin and other crypto assets. The MUI deemed crypto assets as "haram", i.e. banned, as it had elements of "uncertainty, wagering and harm", as reported by (Bloomberg, 2021). In June 2021, the Chinese central bank announced that all transactions of crypto-assets were illegal, effectively banning Bitcoin and other crypto-assets entirely (BBC News, 2021 a). In November 2021, Sweden proposed an EU wide ban of proof-of work crypto-assets like Bitcoin due to their energy consumption. Crypto asset producers were increasing their presence in the Nordic region to search more renewable energy sources, the heads of the Swedish Financial Supervisory Authority and the Swedish Environmental Protection Agency (2021) stated in an open letter. But Sweden would need the renewable energy for the climate transition to meet the Paris Agreement. Energy-intensive mining of crypto-assets should therefore be prohibited.

Other jurisdictions have taken a less rigorous stance and have primarily aimed at bringing crypto-assets "within the regulatory perimeter" to address risks but also support possible benefits of innovation (e.g. Cunliffe, 2021). Besides this call by Cunliffe, the UK's FCA (2021) prohibited activities of crypto-exchange Binance and issued a warning to consumers and on crypto-assets. Australia, in December 2021, introduced a draft legislation aiming at licensing crypto-exchanges and activities in crypto-assets (Reuters (2021b)). For the U.S., Reuters (2021c) reports that the regulation and enforcement across different authorities is shaping for each's respective area of responsibility. The approach of regulating crypto assets is accompanied by several enforcement actions (Vox, 2021). The President's Working Group on Financial Markets (US Treasury, 2021), comprising the Secretary of the Treasury and the Heads of all the key US financial regulators, call to speed up efforts on regulation and guides federal agencies to use their existing powers. The group of legislators call for more federal oversight of custodial wallet providers, i.e. firms that offer products that allow users to hold their crypto tokens. Moreover, the SEC has rejected a bitcoin-based exchange traded fund (ETF) in November 2021 due to concerns of possible price manipulation (FT, 2021a). The OCC (2021) requires from banks to have controls prior to engaging in cryptoassets business and must receive a non-objection. And the U.S. Infrastructure Bill of November 2021 calls cryptoexchanges to notify the tax authority of crypto asset transactions (Time, 2021). In the EU, the European Commission (2020) proposed the regulation for Markets in Crypto-Assets (MiCA). In the absence of a central issuer, MiCA will not regulate Bitcoin and other crypto-assets, but target intermediaries, offering services in crypto-assets (crypto asset service providers). This is in line with other approaches seen for decentralised crypto-assets, e.g. by the FATF, and such approach was also suggested by the ECB's crypto asset task force (2019). Besides MiCA, the European Commission (2021) presented a draft legislative proposal to enhance the EU's framework for AML/CFT. Like MiCA, it requires intermediaries to apply AML/CFT measures and forbids the opening of anonymous crypto asset wallet accounts. These regulations, once they apply, will likely address several of the societal issues related to Bitcoin – but not all of them. The rules will not cover Bitcoin transactions that happen without any regulated intermediaries, namely using non-custodial wallets or on-chain peer-to-peer transactions, or if service providers and countries with low compliance levels are used. For professional criminals, using the Bitcoin network through on-chain peer-to-peer payments does not seem to be particularly challenging, even if the regulation of service providers will make the laundering of Bitcoin received through illicit activities more difficult.

These examples indicate that regulators are progressing in addressing the risks posed by Bitcoin and cryptoassets. At the same time, they illustrate that stances differ (because of different assessments of the value of crypto-assets for society) and initiatives are at different levels of maturity. It can also be noted that regulations, apart from those criminalizing Bitcoin, face limits due to Bitcoin's decentralized and global set-up.

To forestall or limit global regulatory gaps and arbitrage, international cooperation on crypto-assets amongst regulators is important, as the IMF (2018) called for. FSB (2019) demonstrated the manifold initiatives and activities related to crypto-assets at international level. In parallel to initiatives at country level, the international bodies have amplified their efforts in addressing the risks posed by crypto-assets over the last years and some of those international actions have guided national implementations.

An example is FATF, which issued global, binding standards to prevent the use of crypto-assets for money laundering and terrorist financing. As outlined above, FATF focuses on the providers offering crypto-asset services to apply the same safeguards as the financial sector. However, in its progress report of July 2021, FATF (2021) indicates deficits in the implementation in jurisdictions, concluding "that there is not yet a global regime to prevent the misuse of virtual assets and VASPs for money laundering or terrorist financing." Moreover, the report acknowledges that a significant value for peer-to-peer crypto-asset transactions may be operating outside the FATF standards.

A further example is the G7 (2020). The G7 raised concerns that ransomware payments are regularly made in crypto-assets and demanded the implementation of the FATF standards. Furthermore, G7 Finance Ministers and Governors (2021) stated that volatile unbacked crypto-assets were not suitable for payments.

Another international standard setter, the Basel Committee on Banking Supervision (BCBS), consulted on a preliminary proposal for the prudential treatment of banks' crypto-asset exposures, distinguishing crypto-assets that may be generally eligible for falling within the current Basel requirements; and other crypto-assets, such as Bitcoin, that would require a more conservative prudential treatment (BCBS, 2021).

All these examples show that the international community aims at harmonised international standards for addressing the risks of crypto assets. Of course, the translation of those into national rules can be a years-long process with fragmented implementation.

While there has been significant progress towards a consistent and effective regulation of crypto assets, Bitcoin prices and market capitalization have still reached new peaks in November 2021. Some measures by public authorities may have contributed to these new peaks by supporting renewed investment flows into Bitcoin. For example, the US SEC (2021) recently gave the green light for a first futures-based Bitcoin ETF (while it has repeatedly rejected Bitcoin spot market ETF)⁸ or the German legislator has adopted in July 2021 a "Fondstandortgesetz" which allows German investment funds for institutional investors ("Spezialfonds") to invest up to 20 percent into crypto assets. For hesitant investors, such public measures seem to legitimize Bitcoin without necessary safeguards; they could be interpreted as signals that public authorities do not doubt on the sustainability and rationale of Bitcoin. Moreover, these measures facilitate investment flows and the integration of Bitcoin in the traditional financial systems. Finally, any signal from public authorities through measures about Bitcoin are considered indications of the future policy stance. This reduces investors' uncertainty vis-à-vis an asset of which the price is not anchored in any real contribution to society. Overall, it may be concluded that the net effect of authorities' recent measures on Bitcoin were ambiguous. This ambiguous net effect therefore also could apply to Bitcoin's eventual negative consequences for society, which go beyond its use for illicit payments.

10. Conclusion

As also argued in detail elsewhere (Taleb, 2021; Dunn 2021; Green, 2021; Roubini, 2021; Bindseil and Schaaf, 2021) the sustainability of Bitcoin is questionable. It is difficult to find good arguments that support its soundness as a medium of exchange or as form of investment. If Bitcoin eventually collapses, the net social cost of the Bitcoin life cycle will be very large. And it will be the larger the longer it lasts, and the higher Bitcoin's maximum market capitalisation will be. In the absence of a positive contribution of Bitcoin for society, the gross and net social costs will be equal - and they will encompass the energy consumption and hardware usage of the Bitcoin network, the human and technical capital that will have to be written off. What is not quantifiable is the damage to the social fabric that will occur when retail investors find that their savings are lost, while some early investors who got out before the music stopped playing have enriched themselves at their expense.

⁸ The U.S. has introduced many bills in recent years that affect the crypto ecosystem, be it tax requirements or securities law. However different states may have their own regulatory requirements, i.e. the US lacks a comprehensive framework.

Public policymakers have not been fast to address all problems related to Bitcoin. Although its usage for illicit payments has been noted early, slow global implementation and enforcement of AML/CFT rules for Bitcoin based payments has undermined the huge efforts made to prevent illicit payments through regulated industries and allowed regulatory arbitrage by criminal actors. Moreover, Bitcoin has become, also somewhat through the benevolence of public authorities, an asset class that everyone can now easily invest into and that "looks like a security, swims like a security, and quacks like a security, but is not regulated as a security" (Diehl, 2012) and even more importantly, that lacks a plausible underlying contribution to society justifying its valuation.

More recently, many public authorities, have taken or plan to take strong measures against Bitcoin, after concluding that its societal value is negative. Also, regulators of advanced western economies have launched significant implementation measures to fight the reliance on Bitcoin for illicit purposes, although the non-intermediated use of the Bitcoin network is still largely out of regulators' actions. Further regulatory efforts are therefore needed that effectively address all kinds of illicit payments through Bitcoins. The principle of "same function - same risks - same rules" is to apply consistently if global efforts against illicit payments are to be successful, regardless of the unique nature of Bitcoin.

Legislators and authorities need to be careful to not at the same time contribute to renewed momentum of investment flows into Bitcoin that will contribute to increase the market capitalisation of Bitcoin and to the scale of the eventual cumulated social cost of the Bitcoin network. The year 2021 has seen several such developments, and the spike of Bitcoin valuations in November 2021 is likely also attributable to investment inflows that were supported by such measures. For example, the news that the trading of futures-based bitcoin ETFs would (or could) not be prohibited, or the German Fondstandortgesetz effective as of 1 July 2022 allowing institutional investors funds to invest into Crypto-assets are being mentioned as drivers for the Bitcoin price dynamics in autumn 2021.

Last but not least, doubts on the sustainability of Bitcoin and the related social costs does not mean that DLT, blockchain and decentralised finance have no merits as innovative technological approaches. What remains unclear is if crypto coins other than stablecoins (or non-fungible tokens representing ownership of some other assets) can represent a meaningful *investment asset*. In the case of Bitcoin these doubts are particularly strong because of Bitcoin's reliance on the inefficient proof-of work concept and its poor performance as means of payment.

References

Acemoglu, D (2021), "The Bitcoin Fountainhead", in: Project Syndicate, 5 Oct 2021, available: https://www.project-syndicate.org/commentary/bitcoin-an-appealing-distraction-by-daron-acemoglu-2021-10.

Aggarwal, D et al. (2017), "Quantum attacks on Bitcoin, and how to protect against them", Submitted on 28 Oct 2017; https://arxiv.org/abs/1710.10377

Avoca Global Advisors (2021), "Bitcoin: a trojan horse", wcg avoca 14 Oct 2021, available: https://www.linkedin.com/posts/activity-6854411140617818112-NpKx

Auer, R, C Monnet and H S Shin (2021), "Distributed ledgers and the governance of money", BIS Working Papers No 924, January 2021 (revised November 2021), available at SSRN: https://ssrn.com/abstract=3770075 or http://dx.doi.org/10.2139/ssrn.3770075

Baqer, K, Huang, D Y (2016), "Stressing Out: Bitcoin 'Stress Testing'", Published in Financial cryptography, 22 February 2016, http://diyhpl.us/~bryan/papers2/bitcoin/Stressing%20out:%20Bitcoin%20stress%20testing%20-%20analysis%20-%202016.pdf

BBC News (2021), "Bitcoin protests in El Salvador against cryptocurrency as legal tender", 16 September 2021, https://www.bbc.com/news/world-latin-america-58579415.

BBC News (2021)a, "China declares all crypto-currency transactions illegal", published 24 Sept 2021, https://www.bbc.com/news/technology-58678907

BCBS (2021), "Prudential treatment of cryptoasset exposures", published on 10 June 2021, https://www.bis.org/bcbs/publ/d519.htm

Bindseil, U and A Fotia (2021), "Introduction to Central Banking", Berlin: Springer, https://link.springer.com/content/pdf/10.1007%2F978-3-030-70884-9.pdf

Bindseil, U and J Schaaf (2021), "Nicht vom Bitcoin narren lassen", in: Frankfurter Allgemeine Zeitung, 18 Sept 2021

Bloomberg (2018), "IMF Calls for Global Talks on Digital FX as Bitcoin Whipsaws", published on 18 Jnaury 2018, https://www.bloombergquint.com/global-economics/imf-calls-for-global-talks-on-digital-fx-as-bitcoin-whipsaws

Bloomberg (2021), "Crypto Is Forbidden for Muslims, Indonesia's National Religious Council Rules", published 11 Nov 2021, https://www.bloomberg.com/news/articles/2021-11-11/crypto-is-forbidden-for-muslims-says-indonesia-s-ulema-council

Böhme, R, Christin, N., Edelman, B., and Moore, T. (2015), "Bitcoin: Economics, Technology, and Governance", in: Journal of Economic Perspectives, 29, 2 (2015), 213–238.

Business Insider India (2021), "India ends Winter Session of Parliament with no crypto bill in sight", published on 23 December 2021, https://www.businessinsider.in/cryptocurrency/news/no-india-crypto-bill-in-sight-even-after-winter-session-of-parliament-comes-to-a-close/articleshow/88449161.cms.

Coase, R H (1937), "The Nature of the Firm," Economica 4 (November): 386–405.

Cointelegraph (2021), "Iceland cuts power to new Bitcoin miners", 8 Dec 2021, https://cointelegraph.com/news/iceland-cuts-power-to-new-bitcoin-miners

Corporate Compliance Insights (2021), "The Complete Guide to Compliance and Compliance Risk Management", in Resource Library, Whitepapers, 17 November 2021, https://www.corporatecomplianceinsights.com/banks-15b-in-fines-in-2020/

Cunliffe, J (2021), "Is 'crypto' a financial stability risk?", speech given at Sibos, published on 13 October 2021, https://www.bankofengland.co.uk/speech/2021/october/jon-cunliffe-swifts-sibos-2021.

Cybertrace (2021), "Cryptocurrency Crime and Anti-Money Laundering Report", February 2021; available: https://ciphertrace.com/2020-year-end-cryptocurrency-crime-and-anti-money-laundering-report/

Crystal analytic (2020), "Darknet Use and Bitcoin — A Crypto Activity Report by Crystal Blockchain", blog, 19 May 2020, https://crystalblockchain.com/articles/darknet-use-and-bitcoin-a-crypto-activity-report-by-crystal-blockchain/

Diehl, S (2021), "The Token Disconnect", blog, https://www.stephendiehl.com/blog/disconnect.html.

continued

Digiconomist (2012), "Bitcoin may consume as much energy as all data centers globally", blog, March 10, 2021, https://digiconomist.net/bitcoin-may-consume-as-much-energy-as-all-data-centers-globally

Dunn, W (2021), "Bitcoin's gold rush was always an illusion", in: The New statesman, 20 July 2021, https://www.newstatesman.com/business/finance/2021/07/bitcoins-gold-rush-was-always-illusion

ECB Crypto-Asset Task Force (2019)," Crypto-Assets: Implications for financial stability, monetary policy, and payments and market infrastructures", Occasional paper 223, May 2019.

European Commission (2020), "Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL on Markets in Crypto-assets, and amending Directive (EU) 2019/1937", published 24 September 2020, https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52020PC0593

European Commission (2021), "Anti-money laundering and countering the financing of terrorism legislative package", published 20 July 2021, https://ec.europa.eu/info/publications/210720-anti-money-laundering-countering-financing-terrorism_en

FATF (2014)," Virtual Currencies: Key Definitions and Potential AML/CFT Risks", June 2014, https://www.fatf-gafi.org/publications/methodsandtrends/documents/virtual-currency-definitions-aml-cft-risk.html.

FATF (2019), "Guidance for a Risk-Based Approach to Virtual Assets and Virtual Asset Service Providers", June 2019, https://www.fatf-gafi.org/publications/fatfrecommendations/documents/guidance-rba-virtual-assets.html

FATF (2021), "Second 12-Month Review of the Revised FATF Standards on Virtual Assets and Virtual Assets service providers", July 2021, https://www.fatf-gafi.org/media/fatf/documents/recommendations/Second-12-Month-Review-Revised-FATF-Standards-Virtual-Assets-VASPS.pdf

FCA (2021), "Consumer warning on Binance Markets Limited and the Binance Group", published on 26 June 2021, https://www.fca.org.uk/news/news-stories/consumer-warning-binance-markets-limited-and-binance-group

Fincen (2021), "Financial Trend Analysis January to June 2021". See Y. Bizouati-Kennedy "US Government is Seizing so Many Cryptos, It's Enrolling Private Contractors", 5 August 2021 in gobankingrates.com

Fitch (2021), "El Salvador Bank Bitcoin Risk to Depend on Adequacy of Regulation", in: Fitch Wire, 11 Nov 2021, https://www.fitchratings.com/research/banks/el-salvador-bank-bitcoin-risk-to-depend-on-adequacy-of-regulation-11-11-2021

Foley, S, JR Karlsen and TJ Putniņš (2019), "Sex, Drugs, and Bitcoin: How Much Illegal Activity Is Financed through Cryptocurrencies?", in: The Review of Financial Studies, Volume 32, Issue 5, May 2019, Pages 1798–1853, https://doi.org/10.1093/rfs/hhz015

Fitch (2021), "El Salvador Bank Bitcoin Risk to Depend on Adequacy of Regulation", in: Fitch Wire, 11 Nov 2021, https://www.fitchratings.com/research/banks/el-salvador-bank-bitcoin-risk-to-depend-on-adequacy-of-regulation-11-11-2021

Foley, S, JR Karlsen and TJ Putniņš (2019), "Sex, Drugs, and Bitcoin: How Much Illegal Activity Is Financed through Cryptocurrencies?", in: The Review of Financial Studies, Volume 32, Issue 5, May 2019, Pages 1798–1853, https://doi.org/10.1093/rfs/hhz015

FSB (2019), Crypto-assets Work underway, regulatory approaches and potential gaps, published on 31 May 2019, https://www.fsb.org/wp-content/uploads/P310519.pdf

FSB (2021), "Crypto-assets and Global "Stablecoins", published on 11 October 2021, https://www.fsb.org/work-of-the-fsb/financial-innovation-and-structural-change/crypto-assets-and-global-stablecoins/

FSB (2020), "Enhancing Cross-border Payments Stage 3 roadmap", published on 13 October 2020, https://www.fsb.org/wp-content/uploads/P131020-1.pdf

G7 (2020), "Ransomware Annex to G7 Statement", publised on October 13, 2020, https://home.treasury.gov/system/files/136/G7-Ransomware-Annex-10132020_Final.pdf

continued

G7 (2021), "G7 Finance Ministers and Central Bank Governors' Statement on Central Bank Digital Currencies (CBDCs) and Digital Payments", published on 13 October 2021, https://www.bundesfinanzministerium.de/Content/EN/Downloads/G7-G20/2021-10-13-g7-central-bank-digital-currencies.pdf?_blob=publicationFile&v=2

Gandal, N, JT Hamrick, T Moore and T Obermana (2018), "Price manipulation in the Bitcoin ecosystem", Journal of Monetary Economics, Volume 95, May 2018, Pages 86-96, https://www.sciencedirect.com/science/article/pii/S0304393217301666

Green, M (2021), "The Case Against Bitcoin", in "Common sense", 14 May 2021, https://bariweiss.substack.com/p/the-case-against-bitcoin.

Hern, A (2014), "Bitcoin currency could have been destroyed by '51%' attack", The Guardian, 16 June 2014.

IMF (2021), "The Crypto ecosystem and financial stability risks", in: Global Financial Stability report, Chapter 2, October 2021.

John M. Griffin, JF and A Shams (2019), "Is Bitcoin Really Un-Tethered?", (October 28, 2019); https://ssrn.com/abstract=3195066 or http://dx.doi.org/10.2139/ssrn.3195066

Kolbert, E (2021), "Why Bitcoin Is Bad for the Environment", in: The New Yorker, published 22 April 2021, https://www.newyorker.com/news/daily-comment/why-bitcoin-is-bad-for-the-environment

Lagarde, C (2017), "Central Banking and Fintech — A Brave New World?", Speech, London, 28 September 2017, https://www.imf.org/en/News/Articles/2017/09/28/sp092917-central-banking-and-fintech-a-brave-new-world.

Malkiel, B (1973), "A Random Walk Down Wall Street", W. W. Norton & Company, 1973.

McCauley, Robert (2021), "Why bitcoin is worse than a Madoff-style Ponzi scheme", Financial Times, 22 December 2021.

Merten, L (2021): "Chivo wallet: bitcoin Verlust durch Sicherheitsluecke?" published on 29 December 2021, https://bitcoin-2go.de/chivo-wallet-sicherheitsluecke-bitcoin-verlust/.

Nakamoto, S (2008), "Bitcoin: A Peer-to-Peer Electronic Cash System", Bitcoin Foundation, November 2008

OCC (2021), "OCC Clarifies Bank Authority to Engage in Certain Cryptocurrency Activities and Authority of OCC to Charter National Trust Banks" published on 23 November 2021, https://www.occ.gov/news-issuances/news-releases/2021/nr-occ-2021-121.html

Omarova, S T (2021), "The People's Ledger: How to Democratize Money and Finance the Economy", Cornell law school research paper No. 20 – 45; 74 Vanderbilt Law Review 1231 (2021); https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3715735

Oxford Business Review (2020), "The Greater Fool Theory", Oxford Business Review, 2020-12-30.

Panetta, F (2021), "The present and future of money in the digital age", Lecture, Rome, 10 December 2021, https://www.ecb.europa.eu/press/key/date/2021/html/ecb.sp211210~09b6887f8b.en.html#footnote.2

Reuters (2021)a, "Proposed India bill banning crypto payments could mean jail for violations", published on 7 December 2021, https://www.reuters.com/markets/currencies/proposed-india-bill-banning-crypto-payments-could-mean-jail-violations-document-2021-12-07/

Reuters (2021)b, "Australia proposes new laws to regulate crypto, BNPL Proposed India bill banning crypto payments could mean jail for violations", published on 8 December 2021, https://www.reuters.com/markets/currencies/australia-plans-update-regulatory-framework-payment-systems-2021-12-07/

Reuters (2021)c, "US crypto framework begins to evolve: A Special Report update", published on 22 October 2021, https://www.reuters.com/legal/transactional/us-crypto-framework-begins-evolve-special-report-update-2021-10-22/

Reuters (2021)d, "U.S. seizes \$2.3 million in bitcoin paid to Colonial pipeline hackers", published on 7 June 2021, https://www.reuters.com/article/us-cyber-colonial-justice-idTRNIKCN2DJ2BN

Roubini, N (2021), "Bitcoin is not a hedge against tail risk", in Financial Times, published on 10 Feb 2021, https://www.ft.com/content/9be5ad05-b17a-4449-807b-5dbcb5ef8170

continued

Schär, F, A Berentsen (2020), "Bitcoin, Blockchain, and Cryptoassets: A Comprehensive Introduction", MIT press, 2020.

Simpson, A (2021), "Memo cites lessons from ransomware payments by CAN, JBS and Colonial pipeline", in Insurance Journal (2021), 20 November 2021, https://www.insurancejournal.com/news/national/2021/11/29/643569.htm

Swedish Financial Supervisory Authority and the Swedish Environmental Protection Agency (2021), "Crypto-assets are a threat to the climate transition – energy-intensive mining should be banned", published by Finansinspektionen on 5 Nov 2021: https://www.fi.se/en/published/presentations/2021/crypto-assets-are-a-threat-to-the-climate-transition-energy-intensive-mining-should-be-banned/".

Taleb, N (2021), "Bitcoin, currencies, and fragility", published online: 22 Jul 2021: https://www.tandfonline.com/doi/full/10.1080/14697688.2021.1952702?scroll=top&needAccess=true.

Time (2013), "Everything You Need to Know About Silk Road, the Online Black Market Raided by the FBI ", published on 4 Oct 2013, https://nation.time.com/2013/10/04/a-simple-guide-to-silk-road-the-online-black-market-raided-by-the-fbi/.

Time (2021), "Two Things Crypto Investors Should Know About the Infrastructure Bill President Biden Signed" published on 29 November 2021, https://time.com/nextadvisor/investing/cryptocurrency/infrastructure-bill-crypto-taxes/The Economist (2021), "Crypto lobbying is going ballistic - As regulators toughen up, companies hope to influence where the rules end up", 12 December 2021; https://www.economist.com/finance-and-economics/crypto-lobbying-is-going-ballistic/21806674.

The Guardian (2021), "India to ban private cryptocurrencies and launch official digital currency", published on 24 Nov 2021, https://www.theguardian.com/world/2021/nov/24/india-to-ban-private-cryptocurrencies-and-launch-official-digital-currency.

Tooze, A (2021), "Chartbook newsletter #15: Talking and reading about Bitcoin", accessed online on 15 November 2021 : https://adamtooze.substack.com/p/chartbook-newsletter-15.

US Treasury (2021), "President's Working Group on Financial Markets Releases Report and Recommendations on Stablecoins", 1 November 2021, https://home.treasury.gov/news/press-releases/jy0454.

Violation Tracker (2021), Violation Tracker Parent Company Summary, BNP Paribas, https://violationtracker.goodjobsfirst.org/parent/bnp-paribas.

Vox (2021), "Biden's SEC is ready to regulate cryptocurrency", published on 9 September 2021, https://www.vox.com/recode/22663312/coinbase-sec-cryptocurrency-bitcoin

Vukolic, M (2021), "On the Future of Decentralized Computing", accessed online 15 November 2021 on: https://vukolic.com/on-the-future-of-decentralized-computing.pdf

Wewel, C (2021), "Crypto returns: Some stylised facts", FX Atlas, J Safra Sarasin, November 10, 2021: https://publications.jsafrasarasin.com/publ-dl-ch/dldiscl?dl=5DDA631652B5F3C6B91E3CB5DA01A36F3E352BDB2E6853B68DFEDFB963377AF39E9341956D11B781.

Williamson, O E. (1975), "Markets and Hierarchies: Analysis and Antitrust Implications", New York: Macmillan Publishers.

WSJ (2021), "JBS paid \$11 million to resolve ransomware attack", Wallstreet Journal, published on 9 June 2021, https://www.wsj.com/articles/jbs-paid-11-million-to-resolve-ransomware-attack-11623280781.

About the authors

Ulrich Bindseil is Director General Market Infrastructure and Payments at the European Central Bank (ECB), a post he has held since November 2019. Previously, he was Director General Market Operations (2012-2019) and head of the Risk Management Division (2005-2008). Mr Bindseil first entered central banking in 1994, when he joined the Economics Department of the Deutsche Bundesbank, having studied economics. His publications include, among others, *Monetary Policy Operations and the Financial System*, OUP, 2014, *Central Banking before 1800 – A Rehabilitation*, OUP, 2019, and *Introduction to Central Banking, Springer*, 2021.

Patrick Papsdorf is Head of Section of Payments Oversight at the European Central Bank. Prior to this, he has been Adviser in the Directorate General Market Infrastructure and Payments for Eurosystem operated market infrastructures and related analytics. Earlier positions include a commercial bank, the Deutsche Bundesbank and the Federal Reserve Bank of New York. Patrick holds a Bachelor in Business Administration from the University of the Bundesbank and a Master in Global Management

Jürgen Schaaf has been Advisor to the Senior Management of Market Infrastructure and Payments at the European Central Bank since November 2019. He is currently focusing on central bank digital currencies and retail payments strategies. Before that, he was Counsellor to the Executive Board of the ECB (Dec 2012 - October 2019) and Secretary of the Single Supervisory Mechanism (SSM) Project Team Dec 2012 – Dec 2013). Before he joined the ECB, he was Personal Adviser to the Governor of Banque centrale du Luxembourg. In previous occupations he worked ECB watcher at Börsen-Zeitung and Senior Economist at Deutsche Bank. He studied economics in Marburg and Canterbury/Kent and holds a Ph.D. in Economics from Phillips University Marburg.

SUERF Publications

Find more SUERF Policy Notes and Policy Briefs at www.suerf.org/policynotes



SUERF is a network association of cenand tral bankers regulators, academics, and practitioners in the financial sector. The focus of the association is on the analysis, discussion and understanding of financial markets and institutions, the monetary economy, the conduct of regulation, supervision and monetary policy. SUERF's events and publications provide a unique European network for the analysis and discussion of these and related issues.

SUERF Policy Notes focus on current financial, monetary or economic issues, designed for policy makers and financial practitioners, authored by renowned experts.

The views expressed are those of the author(s) and not necessarily those of the institution(s) the author(s) is/are affiliated with.

All rights reserved.

Editorial Board: Ernest Gnan Frank Lierman David T. Llewellyn Donato Masciandaro Natacha Valla

SUERF Secretariat c/o OeNB Otto-Wagner-Platz 3 A-1090 Vienna, Austria Phone: +43-1-40420-7206 www.suerf.org • suerf@oenb.at

EXHIBIT I





Seems he proved that the ledger never halted even though, from what I understand, an employee of Ripple stated that it did....

anybody watching this on Twitter? It's fascinating! Also, I have to admit, I never followed Galgitron very closely, but that guy is a genius!

Anybody agree, disagree, or just want to say something?

Also, I can't get any work done because of this..... Posted October 13, 2020 Okay....so the plot thickens... here is the response backing up the claim that the ledger halted: **Forum Statistics**

Started February 23

By Eric123

Epic Pennant on BTC Chart

Started January 24, 2019

Total Topics	32.6k
Total Posts	887.6k

FOOD





In case anybody missed it, the "sailboat" is also "Galgitron".... okay so I am trying to think this through objectively and the one thing that sticks out is that the party stating/hinting at/misleading people to believe the ledger was halted was openly encouraging people to "figure it out on their own, as all the information, including dates is readily available".... but it doesn't seem that even the savviest of technical prowess would be able to spot the "halt"...IMO 11,647

I'll stop.



nikb

Posted October 13, 2020 Popular Post •••• The data that Galgitron is showing would not represent the issue: he's post-processing ledgers based on the **close_time** field.

Here are the facts: sometime around midnight (UTC time!) on 2018-11-14, the validator operated by **data443** stopped issuing validations. This persisted for almost 30 minutes and was probably caused by a network or other outage as opposed to routine system maintenance.

Shortly after, the **rabbitkick.club** validator operated by @ScottBranson desynced (that's generally not a problem, the Internet isn't a perfect network) and sporadically issued several validations which propagated poorly and which other validators were generally unable to acquire from the network. This lasted for almost 2 hours.

During this time, other validators continued to issue validations, but due a bug, the "healthy" validators repeatedly switched working branches for consensus and only issued partial validations, which was what they were supposed to do per the protocol design as the XRP Ledger values safety. Even though a quorum of validators were periodically in agreement on the proper ledger, the partial status on the validations prevented them from fully validating.

I personally restarted several of Ripple's validators, and other validator operators restarted theirs. The restart effectively reset the LedgerTrie state (it's ephemeral by design), which, in turn, stopped those validators from switching working branches postrestart, resolving the issue as the restarted validators began issuing full, instead of partial validations for the ledgers during this incident. This, by the way, is why the timeline, as represented by the close_time field appears unaffected.

Post-incident, the team at Ripple invested a significant amount of time troubleshooting the issue and proposed several improvements, including a <u>commit</u> that improved the calculation of the preferred ledger branch, added additional diagnostic checks to help making troubleshooting easier, and introduced several unit tests to try and exercise the LedgerTrie code more deeply.

hallwaymonitor, karstnDE, QWE and 13 others 🛛 🕥



 $\mathbf{\Psi}$

Posted October 13, 2020 Thanks @nikb

but to be clear, there really would not be a way for an ordinary community member like myself (or the majority of xrp twitter community) to locate and confirm the halt despite being given the date and time it might have happened, correct?

I ask this because it appeared to me that the original intention (of the hints and questions on twitter) was for the "xrp community" to do the due diligence and figure this out on their own... almost seemed that it was a surprise that nobody already did it as it was sooooo easy (admittedly my own interpretation of what I've ingested on twitter) to figure out.... BUT, there would be no way in hell anybody would have been able to figure that out.... Or (and wouldn't surprise me) I totally misread/misinterpreted a lot of stuff....

FOOD


maybe at the beginning of this thread?

Thanks again for the quick and detailed information!

nikb



Posted October 13, 2020 Just now, EcneitapLatnem said: Thanks @nikb

but to be clear, there really would not be a way for an ordinary community member like myself (or the majority of xrp twitter community) to locate and confirm the halt despite being given the date and time it might have happened, correct?

I ask this because it appeared to me that the original intention (of the hints and questions on twitter) was for the "xrp community" to do the due diligence and figure this out on their own... almost seemed that it was a surprise that nobody already did it as it was sooooo easy (admittedly my own interpretation of what I've ingested on twitter) to figure out.... BUT, there would be no way in hell anybody would have been able to figure that out.... Or (and wouldn't surprise me) I totally misread/misinterpreted a lot of stuff....

Seems a lot of confusion could have been avoided a long time ago if what you had written above was posted days/weeks/months ago.... maybe at the beginning of this thread?

Thanks again for the quick and detailed information!

If anyone was monitoring the validation stream they could have detected it in real-time.

Galgitron dismisses the issue and traduces the good name of others by relying on incorrect information and a flawed understanding of the protocol and the incident.

By way of analogy, this would be like going through a diary and claiming that just because a page is marked as "January 29, 2020" the page must have, necessarily, been written on January 29, 2020 as opposed to a week later, and the timestamp only serves to indicate that the events being described took place at that time.

RecentChange, WrathofKahneman and QWE

3

...



Posted October 13, 2020

39 minutes ago, nikb said:

The data that Galgitron is showing would not represent the issue: he's post-processing ledgers based on the **close_time** field.

Here are the facts: sometime around midnight (UTC time!) on 2018-11-14, the validator operated by **data443** stopped issuing validations. This persisted for almost 30 minutes and was probably caused by a network or other outage as opposed to routine system maintenance.

Shortly after, the **rabbitkick.club** validator operated by @ScottBranson desynced (that's generally not a problem, the Internet isn't a perfect network) and sporadically issued several validations which propagated poorly and which other validators were generally unable to acquire from the network. This lasted for almost 2 hours.

During this time, other validators continued to issue validations, but

...

which was what they were supposed to do per the protocol design as the XRP Ledger values safety. Even though a quorum of validators were periodically in agreement on the proper ledger, the partial status on the validations prevented them from fully validating.

I personally restarted several of Ripple's validators, and other validator operators restarted theirs. The restart effectively reset the LedgerTrie state (it's ephemeral by design), which, in turn, stopped those validators from switching working branches postrestart, resolving the issue as the restarted validators began issuing full, instead of partial validations for the ledgers during this incident. This, by the way, is why the timeline, as represented by the close_time field appears unaffected.

Post-incident, the team at Ripple invested a significant amount of time troubleshooting the issue and proposed several improvements, including a commit that improved the calculation of the preferred ledger branch, added additional diagnostic checks to help making troubleshooting easier, and introduced several unit tests to try and exercise the LedgerTrie code more deeply.

No company policy about airing dirty laundry? Posted October 13, 2020 (edited) 2 hours ago, Soup said: No company policy about airing dirty laundry?

Dirty laundry? What?

I answered a question people had about an incident that occurred on the open, permissionless and publicly accessible XRP Ledger.

I don't see this as "airing dirty laundry" nor do I feel it's a bad idea for incidents to be discussed and evaluated so lessons can be learned.

Edited October 13, 2020 by nikb

Wrote "good idea" instead of "bad idea" accidentally reversing my intent.

Ryyy20, QWE, LetHerRip and 2 others



nikb

Ripple Employee

9 1.1k

Posted October 13, 2020 Just now, nikb said: Dirty laundry? What?

I answered a question people had about an incident that occurred on the open, permissionless and publicly accessible XRP Ledger.

I don't see this as "airing dirty laundry" nor do I feel it's a good idea for incidents to be discussed and evaluated so lessons can be learned.

Your casting doubt on the health of the company's cash cow. I doubt Ripple loves the clarification.

King34Maine and PlanK

2

Posted October 13, 2020

Just now, Soup said:

Your casting doubt on the health of the company's cash cow. I doubt Ripple loves the clarification.

I am not casting doubt on anything. I'm explaining an incident which occurred on a public blockchain.

You may prefer to stick your head in the sand, but I don't and I doubt



Ripple Employee

9 1.1k



Go to topic listing

«

	Theme 🔻	Privacy Policy	Contact	
Home > XRP > Technical Discussion >				🖭 All Activit

We have placed cookies on your device to help make this website better. You can adjust your cookie settings, otherwise we'll assume you're okay to continue.