

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/312428172>

# Blockchain: A Primer

Working Paper · December 2016

---

CITATIONS

0

---

READS

307

1 author:



Gerald P Dwyer

Clemson University

119 PUBLICATIONS 1,571 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



Why are some countries poor and others not? [View project](#)



Financial crises, effects and regulatory responses [View project](#)

All content following this page was uploaded by [Gerald P Dwyer](#) on 16 January 2017.

The user has requested enhancement of the downloaded file. All in-text references [underlined in blue](#) are added to the original document and are linked to publications on ResearchGate, letting you access and read them immediately.

## Blockchain: A Primer

Gerald P. Dwyer  
Clemson University

### Abstract

The Bitcoin blockchain is the primary innovation in Bitcoin that makes it practical. Blockchains have applications in many contexts other than cryptocurrencies. This note is an introduction to blockchains that requires no prior knowledge, including of Bitcoin. Blockchains are ledgers of transactions kept by a set of participants, none of which is accorded special status as the “correct one.” Instead, agreement is reached by a process of consensus. I show how this works for Bitcoin, discuss applications in many alternative settings and provide some detail about a very different proof-of-concept application of blockchains by the Japan Exchange Group.

December 2016

This primer is based on an invited tutorial given at the Financial Association meeting in October 2016. I thank the Spanish Ministry of Economy and Competitiveness for support through the project ECO2013-42849-P at the University of Carlos III, Madrid.

People have diametrically opposed opinions about blockchain, just like they do about Bitcoin. Blockchain was the innovation that made Bitcoin work as currency (Dwyer 2015). Blockchains are much more prosaic than Bitcoin, simply being databases with a particular structure. That structure makes it possible to resolve multiple copies to consensus about a single chain being the correct one.

Richard Branson has suggested that Blockchain could bring an “economic revolution” to many developing countries where proving ownership of assets such as real estate is difficult or non-existent (Kharpal 2016). In particular, Hernando de Soto is involved in a trial to use blockchain as a property register for the Republic of Georgia (Prisco 2016).

In terms of U.S. developments, NASDAQ has implemented Linq which is intended to record trades of stock in liquidity events. Overstock.com has issued shares on the Bitcoin blockchain and these shares can be traded on a blockchain. Financial firms, including UBS, Wells Fargo, J.P. Morgan and many others are participating in blockchain developments and trials such as HyperLedger and R3. A Federal Reserve Governor gave a detailed speech (Brainard 2016) and the Federal Reserve Board issued a technical paper (Mills 2016).

There are detractors. Craig Pirrong (2016) has argued that much of the discussion of blockchains and how they will transform financial markets is nothing but “techno-utopianism”. The supposed roles for blockchains are nothing more than trivial, and sometimes worse, transformations of existing procedures.

It is easy to dismiss technological change in the future because it is unpredictable. In this context it is worth recalling Kenneth Olsen, president of the mini-computer maker DEC, who had the misfortune of saying in 1977 that “I see no reason why anyone would want a computer in their home.” He was right about one thing: Few would have wanted computers at that time in their homes. Computers were accessible by Hollerith-encoded cards or by expensive, dumb terminals. Today, a smart phone such as a Samsung S5 is over 20 times faster than a supercomputer in 1979.<sup>1</sup> The Internet now is widely available and the easy-to-access World Wide Web has been invented. Global Positioning System (GPS) also has been implemented and is accessible by these small supercomputers. Many might say that they can’t see leaving home without one of these supercomputer in their pockets. Let alone would they want to be home without one.

Will such transformations ultimately make blockchains more ubiquitous and important than seem plausible today? Maybe so. Maybe not. I will suggest in this essay that blockchains will have a role in our financial future. How big a role remains to be seen.

---

<sup>1</sup> The comparison is based on a standard test suite for computers, Linpack with n=100. The scores are 3.4 million floating point instructions per second (mflops) for a Cray I in 1979 (Dongarra 2007) and a typical result of 76.34 mflops running Mobile Linpack on my personal S5. This is a rough comparison (Dongarra, *et al.* No date). The best result of 100 evaluations was 116.47 mflops.

It is quite possible to describe blockchains and their implementation with little or no reference to any existing system. This essay is a primer and I will proceed on the opposite tack. I mainly will illustrate blockchains in two implementations, with attendant discussion of the general issues.

Blockchains can be quite different depending on the context. Another name for blockchains is “distributed ledger technology.” Blockchains generally involve a database distributed among participants, cryptography to secure important parts of the database, and some method to reach a consensus about which of alternative databases is the correct one. One can conceive of a variety of different implementations. It is simpler to examine two implementations in particular: 1. Bitcoin; and 2. A Japan Exchange Group Trial. Bitcoin’s implementation is well documented. Many of the proof-of-concept trials in progress are proprietary and not well documented. An exception is the Japan Exchange Group Trial, which is likely to have many elements of similar trials by large financial institutions.

### **Bitcoin and its blockchain**

The key elements of Bitcoin are: 1. A database of transactions; 2. Available on a peer-to-peer network; 3. With blocks added to the blockchain to record new transactions; and 4. A way to reach consensus on the current valid transactions.<sup>2</sup>

An important difference between the Bitcoin network and most apparent networks is their topology, or organization. Typical networks, as at universities and firms, are client-server networks. As Figure 1 (a) shows, the network has a central server and clients are nodes that interact with other computers through the server. Whether a client wants to interact with another computer on the same network or a computer on a different network, interaction is executed through the server.

Online banking provides a useful example. When you connect to your bank, you interact with the bank’s computer system. That computer system maintains a copy of your bank balance and transactions. While you may check the bank’s balance for your account periodically, the bank’s summary of your account is the correct one, not one you might have, and you can correct errors by bringing the error to the bank’s attention. If you want to send funds from your account to another customer’s account at the bank, that transfer is routed through the server. You do not simply connect to the other customer’s account and transfer the funds.

The Bitcoin network is quite different; it is a peer-to-peer network. Figure 1 (b) shows such a network. The network is called “peer to peer” because all nodes on the network are peers with equal standing. There is no server which contains the accepted database and provides services to clients. Instead, in a peer-to-peer network, some or all of the nodes in the network maintain

---

<sup>2</sup> Much of this section relies on material available in Dwyer (2015).

copies of the database. As in Figure (b), not all of the nodes are necessarily connected to all other nodes.

There are at least a couple of questions raised by this organization into a peer-to-peer network.

Why does anyone maintain a copy of the database? It would seem lower cost to inquire from the database held by some trusted party rather than maintain a copy oneself. In Bitcoin, as we shall see, the database is kept at least by “miners” who use the current copy of the database and computations to generate a flow of income.

Second, and maybe more disconcertingly, what is to keep the copies of the database the same? It would be easy to see one node add a transaction and another not, or for one node to add two transactions in one order and another peer to add the same two transactions in the opposite order. If the database is to have some integrity, it must be identical across the peers, at least eventually. Agreement on the state of the database is called “consensus.” Reaching consensus might be guaranteed in a short period of time or it might only be guaranteed asymptotically.

A digression is in order here. Functions called “hashes” are an important part of Bitcoin and its blockchain. A hash is a function that takes an input message  $m$  and converts it into an output message  $h$  of a certain length through the hashing function  $H$ :

$$h = H(m)$$

A hashing function takes an input  $m$  of some length and produces an output  $h$  of a specific, typically shorter length.<sup>3</sup> Not just any function can be a cryptographic hashing function. The cryptographic hash functions used in Bitcoin also are: 1. One way with it difficult or impossible to find input data that can produce a particular hash value. 2. Computationally infeasible given an input to find another input that produces the same hash value; and 3. Computationally infeasible to find two hash values that produce the same hash value (Paar and Pelzl 2010, pp. 296-303.) To put it less precisely in terms of standard mathematical functions, cryptographic hashing functions are highly nonlinear without regular features; for example, they are not concave or invertible. The hashing function used in Bitcoin is based on the cryptographic function SHA-256.<sup>4</sup> The SHA-256 hash of the prior sentence is eb1f3c16444cbb59a8a58626366a2726335b45c2f1c811bfcf3c149fb943f560.<sup>5</sup> The same text with the letter “f” deleted from the final “function” generates the totally different sequence 9c29d6bffc8cab8b0807e99754edb9cde780adb676de3c5a9e97b5debaff673b.

---

<sup>3</sup> Length is measured in terms of number of bytes.

<sup>4</sup> The hashing function “H” generally involves a substantial number of complicated operations, and SHA-256 does involve quite a few operations.

<sup>5</sup> The representation is hexadecimal, the base 16 with the 16 “digits” 0123456789abcdef. The hashes are computed by the website <http://www.xorbin.com/tools/sha256-hash-calculator>.

Bitcoin operates according to design rules laid out by Nakamoto (2008) which can be revised by common agreement of the team responsible for the software and the miners (Dwyer 2015). The design rules specify the rule for the creation of bitcoins, how blocks are added and how often they are added. Bitcoins are created at a pre-specified rate with an upper limit of 21 million bitcoins, which will be reached in 2041. After 21 million bitcoins have been created, no new bitcoins will be created. Blocks are added at a pre-specified average rate of every ten minutes with the new bitcoins are created.

Transactions are added to the blockchain in a block, hence the “block” in blockchain. Each block in the Bitcoin blockchain includes a header and a set of transactions.<sup>6</sup>

The header includes: 1. the hash of the previous block’s header; 2. a time stamp; 3. the target difficulty for this block; and 4. an open field to alter the hash of the header of this block.<sup>7</sup>

The first transaction is a singular transaction called the “coinbase.” This transaction records the reward for creating the block. It includes the address of the miner receiving the new bitcoins for creating the block plus the total transaction fees for the block.

Each transaction includes 1. The address or addresses to which the bitcoins are sent; 2. The digital signature of the party sending the bitcoins;<sup>8</sup> and 3. The number of bitcoins sent. Figure 2 illustrates the header and the sequence of transactions.

The blockchain is a chain because each block contains the hash of the previous block’s header. Hence the “chain in “blockchain.” It is trivial to compute the hash of the previous block’s header and verify that this block is the successor of the previous block. This makes it easy to verify if a block is in the correct order in the Bitcoin blockchain. That is one part of keeping consensus on the correct copy of the blockchain.

---

<sup>6</sup> It also includes the size of the block in bytes and the number of transactions. The description of a block that follows is based on Antonopoulos (2015, Chs. 7 and 8) and Franco (2016, Ch. 7).

<sup>7</sup> It also include a version number to track software and protocol upgrades and a hash of the Merkle tree of the block’s structure. A Merkle tree of the transactions is created to speed up finding transactions and make it possible to verify a transaction is in a block without downloading the entire block (Antonopoulos 2015, pp. 164-71.)

<sup>8</sup> The digital signature is based on private-key public-key cryptography, as is the address. Every user has at least one public-key private-key pair. The address in the transaction is a hash of the public key and the bitcoins are sent to address. The digital signature is created by hashing the sender’s public key with the sender’s private key. The identity of the sender can be verified because public key applied to this private key of the public key is the public key. This might seem to be reversibility but not in a sense that the private key can be inferred from the public key given long enough keys and current computer technology. In short, let  $\text{pub}(\cdot)$  denote application of a hash function with the public key and  $\text{pri}(\cdot)$  denote application with the private key. Then  $\text{sig} = \text{pri}(\text{public key})$  and  $\text{public key} = \text{pub}(\text{pri}(\text{public key}))$ . Knowing  $\text{pub}(\cdot)$  does not make it computationally feasible to generate  $\text{pri}(m) = \text{pub}^{-1}(m)$ . The private key cannot be inferred from the public key and the private key is known only by the encryptor. Franco (2015) provides a readable discussion and Paar and Pelzl (2010) provide a more thorough introduction.

Who adds a new block?

Miners work to add a new block. They create a header and assemble transactions to create a block. If this were all there were to it, obviously almost anyone and therefore everyone would be able to do it. Instead, miners compete to add a block. A new block is added when a hash  $h$  of the block header  $b$  is less than or equal to a target value  $t$ , i.e.

$$h = H(b) \leq t$$

where  $H$  is the pre-assigned function for doing this hash. The target value is created by having leading zeroes. Finding a hash less than or equal to  $t$  is not a matter of skill but of sheer computing power and luck. A solution value of  $h$  less than or equal to  $t$  is best found by a random search. In this instance, a random search means randomly changing the nonce part of the header and some other open spaces to achieve a hash value less than or equal to  $t$ . The goal for the system is to add a new block every ten minutes and the difficulty,  $t$  is set to add a block every ten minutes. The more computer power mining, the faster any given hash less than or equal to some value can be achieved.<sup>9</sup> The difficulty is increased if blocks are added more often than every ten minutes and decreased if blocks are added less often than every ten minutes.

What happens if two miners find a solution at or about at the same time? This is possible and happens. As Figure 1 shows, not all miners are directly connected to all other miners. When a miner broadcasts he has found a solution, it takes time for the broadcast and the solution to propagate in the system. It is quite possible to have more than one solution being propagated at the same time. After receiving a message that a solution has been found and the blockchain is longer by one block, each miner puts transactions together for the next block and searches for the next block. It is quite possible, and apparently happens roughly daily, that miners are working on adding a block to two different blockchains of the same length. In reasonably short order, one of the two blocks becomes the longest and all miners will switch to that block.<sup>10</sup>

This resolution of different versions of the blockchain is called reaching “consensus.” Consensus in Bitcoin is agreement on the longest blockchain. This rule is fundamental to maintaining consistency of copies of the blockchain across nodes on the network given the lack of a trusted central server to guarantee the correct blockchain.

---

<sup>9</sup> There are issues associated with mining that are outside the scope of this primer. Miners can be, and mostly have been for some time, organized into mining pools in which miners are rewarded for participating in a joint search. There is an incentive to join a pool because members of pools face less idiosyncratic risk than miners operating alone. Issues can arise if a pool becomes dominant, in particular controlling more than 50 percent of the network’s computing power. These issues are not entirely trivial because a pool consisting of all miners would face no risk because, on average, a block is added every ten minutes. Put more succinctly, there is no aggregate risk and idiosyncratic risk arises because miners do not all pool their efforts together in one pool.

<sup>10</sup> This relies on the variance of the time to find a solution hash being large enough that one or the other blockchain becomes the longer one. Once a block becomes the longer block, it can be expected to remain the longer block.

This resolution with a peer-to-peer network sometimes is called “trustless” because there is no central server that is trusted to have the correct blockchain. Instead though, any particular node must trust nodes to which it is connected and it must trust the rules governing Bitcoin will achieve consensus on a single blockchain as being the correct one. By design, Bitcoin has no central point of failure. In a traditional client-server network, if the central server is corrupted or compromised, then all clients are dealing with these false data. In a peer-to-peer network, if not too many peers are corrupted or compromised, then all peers can eventually have a correct copy of the data.

The protocol underlying the Bitcoin blockchain is not the only possible way to organize additions to a database in a peer-to-peer network. The particular protocol in Bitcoin is called “proof of work” because the miners must do computational work in order to add a block. By the characteristics of the cryptographic hash function used, there is a lot of work involved in finding a solution but it is easy to verify that a solution has been found once presented with the block including the nonce and the solution. The hash need only be run once to verify that the proposed solution is a solution for that block.

The miner who solves the problem receives the new bitcoins created along with the addition of the block. The miner also receives transactions fees which are computed as change created in the transactions in the block.

Proof of stake is another mechanism for reaching consensus often mentioned as an alternative to proof of work. This is used by some other cryptocurrencies (Franco 2015, pp. 234-36). The amount of resources used in mining bitcoins is controversial but there is no doubt that electricity consumption in mining is nontrivial and increasing, possibly rivaling the electricity consumption of countries such as Ireland or Denmark (O’Dwyer and Malone 2014; Deetman 2016). Alternative consensus mechanisms could have a high payoff.

### **Generalizations of Blockchain’s Usefulness**

It quickly became apparent to many that the blockchain used in Bitcoin can be generalized and used in other contexts. I have glossed over many important details about the Bitcoin blockchain. The blockchain is a record of all bitcoins created and bought and sold including the address and the associated public keys which receive new bitcoins and buy and sell bitcoins. Using the record of transactions, it is possible to compute the ownership of all bitcoins at any moment. By keeping a record of all transactions, an ownership registry is created.

Ownership registries are useful in a variety of contexts. While not obvious from the perspective of someone in a high-income country, property registries and ownership are not readily available in many countries (de Soto 1989; 2000). A blockchain for property registry in a country could be both very worthwhile and relatively inexpensive. This is the basis of the blockchain project in Georgia organized by BitFury, Hernando de Soto and the Republic of Georgia (Prisco 2016).

Blockchains have been suggested for use in securities transactions. It is not hard to envision trading securities on the blockchain, although this is not necessarily a substitute for trading on organized exchanges. Bitcoins trade on exchanges as well as directly on the blockchain in person-to-person transactions ([Brandvold \*et al.\* 2015](#); Garcia and Schweitzer 2015; and Pieters and Vivanco). These trades on exchanges are executed without brokers or intervening intermediaries other than the exchange itself.

NASDAQ has created a trial blockchain Linq for public offerings of securities (NASDAQ 2015). At the end of 2015, the firm chain.com recorded the first offering of securities to a private investor on a blockchain. This initial proof of concept worked (Rizzo 2015).

Overstock.com has created a subsidiary tØ that provides a trading platform using a blockchain. Trading on that platform began in December 2016.

The Australian Stock Exchange is building a blockchain to evaluate a blockchain to replace its current technology for settling trades. Its current announced plan is to build a production version and then evaluate in 2017 whether to use the blockchain for recording and settling trades (Eyers 2016; Reichert 2016). The Australian Stock Exchange is by no means the only exchange evaluating blockchains. The Japan Exchange Group's efforts are discussed in detail below.

One purpose of using blockchains is to shorten the cycle between agreement on a transaction and final settlement. At the end of 2016, settlement of a stock purchase required three days. While under pressure to shorten the time before settlement, participants in the stock market are not inclined to shorten it to same-day or one-day settlement (Boston Consulting Group, 2012).

A major possible benefit of a blockchain for stock ownership is a relatively simple one: standardization across firms. This would shorten the settlement time and lower costs substantially. The costs and benefits from making the change all the way to virtually instantaneous settlement are not clear, perhaps especially since real-time gross settlement would be part of the result. Still, real-time gross settlement is not an untried system, being currently in use in the United Kingdom between banks

In banking, trade finance is a focus of a great deal of interest (IBM 2016). Trade finance tends to be involved and difficult to standardize. The blockchain, with its ability to provide an immutable record and standardize processes, has the potential to substantially lower the costs of providing trade finance.

International payments have been touted as a way that Bitcoin and blockchains could lower the costs of transactions substantially. The costs of remittances sent by migrants to relatives is substantial, on average 7.4 percent of the value transmitted (World Bank 2016). While much of the cost is the physical cost of maintaining offices and providing physical cash, Bitcoin has the

potential of lowering those costs. There is no systematic data on the extent to which cryptocurrencies are used in international transactions but the usage appears to be increasing. (Faife 2016; Alegado 2016).

Blockchains also can be used in multi-stage transactions, which are part of what is included in “smart contracts” based on a blockchain (Franco 2015, Ch. 12). For example, a potential purchase of a security can be recorded on the blockchain by locking up the use of a cryptocurrency. Then later, when the securities are delivered, the transfer of the cryptocurrency can be made final.

Some have suggested more ambitious smart contracts, for example contracts for difference Franco, 2015, pp. 190-91).<sup>11</sup> In a contract for difference, two parties take opposite positions concerning the change in the price of an asset, say the price of Apple stock. The contract pays off on a set date and one party pays the other.

There are a variety of ways to set up such contracts in a blockchain. The more plausible way is to have an initial entry that locks up funds as margin. On the expiration date of the contract, the two parties agree to transfer the funds based on the actual change in price.

Some have suggested an alternative, in which the entry in the database contains latent code that goes to the agreed source of the price data and then automatically transfers the funds based on the price. While possible, there are issues with this arrangement in a distributed database with multiple copies of the code. Do all nodes check and perform the transfer? This seems rather wasteful especially since only one transfer is desired. The bitcoin blockchain is not equipped to deal with external checks of data or feeds of data, although some other cryptocurrencies are (Franco 2015, pp. 194-207).

This list of possible uses of blockchains is far from exhaustive.<sup>12</sup> Some have suggested that blockchains could be used as records to start Distributed Autonomous Organizations (DAOs). The complications involved in trying to organize an organization became evident when the cryptocurrency Ethereum had funds stolen from its first substantial DAO (Mizrahi 2016).

It is difficult to forecast which of these uses will turn out to be very successful, somewhat successful or abject failures. Innovation by its very nature is unpredictable. A set of related innovations is even harder to predict than the usefulness of that set of innovations in the long run. Trying to predict what will be successful in the long run with much confidence would be akin to trying to predict in 1977 the success of navigation on smart phones.

### **Japan Exchange Group**

---

<sup>11</sup> While problematic for legal reasons in the United States, these contracts are legal in many other countries.

<sup>12</sup> Swan (2015), Tapscott and Tapscott (2016) and Raval (2016) include a variety of applications underway.

While these generalizations give an overview of how blockchains could be used in contexts besides Bitcoin, it is impossible to provide a detailed but non-technical overview of the implementation issues that arise because of the different contexts. In addition, many implementations of blockchains are proprietary and details are not in the public domain. Indeed, totally inconsistent with the spirit of Bitcoin, some parties have filed applications for patents related to blockchains.<sup>13</sup>

In this section, I cover in some detail two proof-of-concept trials by the Japan Exchange Group. The reason for choosing this particular effort is the availability of details about the proof of concept documented in ““Applicability of Distributed Ledger Technology to Capital Market Infrastructure” (Santo *et al.* 2016). The term “distributed ledger technology” is another term for blockchain if one wants to reserve the term blockchain for distributed ledgers with all of the characteristics of the Bitcoin blockchain, which is not the usage in this chapter.<sup>14</sup> By the term “blockchain”, I mean not only the Bitcoin blockchain but also any sequential database that has distributed copies at various locations with consensus reached in one way or another by the various parties. As emphasized by Wattenhofer (2016), there is no reason to limit discussion of consensus to the particular mechanism used by Bitcoin.

While not emphasized earlier in this chapter, there is little doubt that speeding up and lowering the cost of settlement is a major possible application of blockchains and a major point of interest by financial intermediaries. The cost of clearing and settlement around the world is substantial. Mainelli and Milne (2016, p. 8) estimate the annual worldwide back-office and middle-office costs of securities clearing and settlement at \$100 billion or more. Important drivers of these costs are the inconsistent private databases used by participants and manual intervention required in clearing and settlement. The three-day settlement for securities trades in the United States has been a point of contention if not ridicule by many (Boston Consulting Group 2012).

The purpose of the Japan Exchange Group’s proof of concept was to determine whether a blockchain could speed up clearing and settlement of stock trades and do so at a lower cost than the current methods. The Japan Exchange Group is the operator of the Tokyo Stock Exchange, the Osaka Exchange, the Japan Securities Clearing Corporation and Japan Exchange Regulation. This proof of concept was a joint effort by the Japan Exchange Group, IBM Japan Ltd., Nomura Research Institute, Ltd., CurrencyPort Limited and participants from industry.

---

<sup>13</sup> This is not particularly surprising; there is no reason that a bank or investment bank should care about the ethos underlying Bitcoin or not benefit from someone else’s failure to file for a patent and then file for patents itself. It may well be unfortunate in the long run though for all concerned, an issue discussed in the conclusion.

<sup>14</sup> This term seems to be used more in official government papers and talks than in the computer-science and programming literature. Mills *et al.* (2016, p. 3) define “distributed ledger technology” for their paper as “some combination of components including peer-to-peer networking, distributed data storage, and cryptography that, among other things, can potentially change the way in which the storage, recordkeeping, and transfer of a digital asset is done.”

Unlike the Bitcoin blockchain, the database of securities trades in the proof of concept was not publicly available to just anyone.<sup>15</sup> Instead, it is a “permissioned” blockchain in which participation is allowed only by those with permission. In addition, because traders do not necessarily want information about securities trades made available to other trading firms, the structure of the database keeps some information private from other participants except Japan Exchange Group. The network topology is given in Figure 2.

The Japan Exchange Group (JPX) has a superior position in the graph because it provides trusted information on trades to all nodes involved in a trade. The purpose of the network is to clear and settle the trade, not to certify an agreement to execute a trade. All participants are connected to all other participants. A trade might involve traders at any two firms and so all of them are connected.

The consensus algorithm used in the proof of concept is quite different than the protocol used by Bitcoin.<sup>16</sup> The consensus protocol used in the proof of concept requires only that just over two-thirds of the nodes agree on the correct blockchain.<sup>17</sup> This particular algorithm does not use proof of work, as does Bitcoin. This can make the process faster and eliminates the substantial electricity and computing costs in Bitcoin. A permissioned blockchain makes it possible to rely on trust in other participants, something avoided in Bitcoin’s blockchain.

The central party, JPX, sends information to the participants about a trade. This is a new block and it is trusted by all parties. This starts the execution of a smart contract at each of the nodes, which requires each participant to validate the trade and execute code. The trade information included ownership information by investors. When messages were received from just short of two-thirds of the other participants that the block was valid, then a block was added to the blockchain. Cash payment was recorded on the blockchain as a token transfer with the necessary messages passed to an off-blockchain payment system.

The validating nodes were member institutions. Other firms with listed stocks could participate in the proof of concept but did not participate in reaching consensus.

---

<sup>15</sup> Due to privacy concerns, detailed information on all trades is not available even to all of the participants.

<sup>16</sup> The consensus protocol used in the proof of concept goes by the name “Practical Byzantine Fault Tolerance” (Castro and Liskov 1999). The original problem of reaching agreement among nodes was posed by Leslie Lamport and co-authors as a co-ordination problem faced by Byzantine generals laying siege to a city (Lamport, Shostak and Pease 1982). The name of the solution is derived from this problem statement.

The Byzantine generals command divisions of an army surrounding a city. There is no central command of the armies and the generals can communicate only by messengers. The generals want to reach a decision about what course of action to take. If a few generals storm the city and most do not, the outcome might be quite serious losses. If they all storm the city, victory is more likely. The problem is even more complicated because some of the generals may be traitors.

<sup>17</sup> More precisely, the algorithm with  $n$  participants allows for  $(n-1)/3$  of the nodes to be “faulty.” Consensus is reached when  $2(n-1)/3+1$  of the nodes agree. This is exactly  $(2/3)n-(2/3)+1$  nodes, or rounded up in integers, one more than  $2/3$  of the nodes. There are of course further details about communication and timeliness.

Actions permitted on the blockchain included more than just reconciliation of trades and settlement. Securities issuance, dividends and stock splits could be recorded and, because ownership information was included in trades, the blockchain became an ownership registry.

The overall conclusions in the working paper are:

1. The process is capable of processing tens to a hundred transactions per second. It is not currently feasible for high volume trading on exchanges. It is sufficient for over-the-counter markets, at least in the Japanese environment.
2. It is feasible for post-trade processing where millisecond and microsecond speeds are not necessary.
3. The study's authors recommend a central administrator with nodes mutually validating correctness of the state.
4. The procedure was lower cost than the current procedure.

An unresolved issue is the implied use of real time gross settlement of securities trades in the procedure. Market participants generally prefer netting, which reduces the amount of liquid assets that have to be held.

Were the trials successes? Not surprisingly, the Japan Exchange Group's trial left the exchange at the center of the trading process. Whether exchanges might be replaced in some contexts remains to be seen. Widespread adoption of blockchains for clearing and settlement will take thought and require innovations in the way trades are executed and settled. Legal changes are likely to be necessary in addition to technical innovations. The potential for much faster settlement definitely is there.

## **Conclusion**

Blockchains are a particular implementation of ledgers of transactions. The most distinguishing characteristic is the distribution of the database among participants with no particular copy the correct one. Instead, consensus methods are used. Possibly anything that can be done with a ledger of transactions can be done with a blockchain.

A major advantage of a blockchain is the agreed common platform on which transactions are recorded. This is no small thing. Agreeing to use the same platform is an agreement not to make marginal changes in one's own ledger that create incompatible with other ledgers. Such a change and the resulting incompatibility may be useful and otherwise unimportant today. It can become quite important in the future and require substantial effort at reconciliation.

Ironically given concern about blockchains being too slow, one of the consistent complaints about blockchains is the inconsistency of their rapid settlement with current ways of doing some things in financial markets, such as short sales. Whether such inconsistencies will

seriously the eventual adoption of blockchains in financial markets remains to be seen. As long as competition is allowed to affect the process and it is not throttled by regulation, there is every reason to think blockchains will be adopted where they improve trading or lower costs or both and not where they are worse than current practice. Regulation may well seriously hinder the adoption of blockchains.

Many aspects of blockchains are inconsistent with current regulations. For example, contracts for difference may be executed only on an organized exchange in the United States under current interpretations of the rules. The evidentiary value of a blockchain in a court of law is unknown.

There is no reason to think that currently known protocols for reaching consensus include all possible protocols. It is likely that research into alternative mechanisms for reaching consensus will have high value.

## References

- Alegado, Siegfried. 2016. "Philippines mulling Bitcoin regulation as remittance use surges." Bloomberg, December 21. Accessed at <https://www.bloomberg.com/professional/blog/philippines-mulling-bitcoin-regulation-remittance-use-surges/>.
- Anonymous. "Samsung Galaxy S5 Review." 2014. Gsmarena.com. Accessed at [http://www.gsmarena.com/samsung\\_galaxy\\_s5-review-1064.php](http://www.gsmarena.com/samsung_galaxy_s5-review-1064.php).
- Antonopoulos, Andreas M. 2015. *Mastering Bitcoin*. Sebastopol, California: O'Reilly Media Inc.
- Boston Consulting Group. 2012. *Cost Benefit Analysis of Shortening the Settlement Cycle*. Available at [www.dtcc.com/~/.BCG\\_Shortening\\_the\\_Settlement\\_Cycle\\_October2012.pdf](http://www.dtcc.com/~/.BCG_Shortening_the_Settlement_Cycle_October2012.pdf).
- Brainard, Lael. 2016. "Distributed Ledger Technology: Implications for Payments, Clearing and Settlement." Speech delivered at the International Institute of Finance Annual Meeting, October 7. Available at <https://www.federalreserve.gov/newsevents/speech/brainard20161007a.pdf>.
- Brandvold, Morten, Peter Molnár, Kristian Vagstad, Ole Christian Andreas Valstad, 2015. Price discovery on bitcoin exchanges. *Journal of International Financial Markets, Institutions and Money* 36 (May), 18-35.
- Castro, Miguel and Barbara Liskov. 1999. "Practical Byzantine Fault Tolerance." *Proceedings of the Third Symposium on Operating Systems Design and Implementation*.
- de Soto, Hernando. 2000. *The Mystery of Capital*. New York: Basic Books.
- de Soto, Hernando. 1989. *The Other Path*. New York: Harper & Row Publishers.
- Deetman, Sebastiaan. 2016. "Bitcoin Could Consume as Much Electricity as Denmark by 2020." *Motherboard*, March 20. Accessed at <http://motherboard.vice.com/read/bitcoin-could-consume-as-much-electricity-as-denmark-by-2020>.
- Dongarra, Jack. 2007. "Frequently Asked Questions on the Linpack Benchmark and Top 500." Accessed at <http://www.netlib.org/utk/people/JackDongarra/faq-linpack.html>.
- Dongarra, Jack, Reed Wade and Paul McMahon. No date. "Linpack Benchmark – Java Version." Accessed at <http://www.netlib.org/benchmark/linpackjava/>.
- Dwyer, Gerald P. 2015. "The Economics of Bitcoin and Similar Private Digital Currencies." *Journal of Financial Stability*, 17 (April 2015), 81-91.

Eyers, James. 2016. "ASX Builds Blockchain for Australian Equities." *Sydney Herald*. January 22. Accessed at <http://www.smh.com.au/business/banking-and-finance/asx-builds-blockchain-for-australian-equities-20160121-gmbic0.html>.

Faife, Corin. 2016. "Why Bitcoin's Remittance Disruption Slowed to a Crawl." CoinDesk. December 11. Accessed at <http://www.coindesk.com/why-bitcoins-remittance-disruption-slowed-to-a-crawl/>.

Franco, Pedro. 2015. *Understanding Bitcoin*. Chichester, West Sussex: John Wiley & Sons Ltd.

Garcia, David, and Frank Schweitzer. "Social Signals and Algorithmic Trading of Bitcoin." *Royal Society Open Science* 2: 150288. <http://dx.doi.org/10.1098/rsos.150288>

IBM Institute for Business Value. 2016. *Blockchain Rewires Financial Markets*. Somers, NY: IBM.

Kharpal, Arjun. 2016. "Richard Branson: Blockchain could create 'economic revolution' in emerging markets." CNBC, October 3. Accessed at <http://www.cnbc.com/2016/10/03/richard-branson-blockchain-could-create-economic-revolution-in-emerging-markets.html>.

Mainelli, Michael, and Alistair Milne. 2016. "The Impact and Potential of Blockchain on the Securities Transaction Lifecycle." Swift Institute Working Paper No. 2015-07.

Mills, David, Kathy Wang, Brendan Malone, Anjana Ravi, Jeff Marquardt, Clinton Chen, Anton Badev, Timothy Brezinski, Linda Fahy, Kimberley Liao, Vanessa Kargenian, Max Ellithorpe, Wendy Ng, and Maria Baird. 2016. "Distributed ledger technology in payments, clearing, and settlement," Finance and Economics Discussion Series 2016-095. Washington: Board of Governors of the Federal Reserve System, Available at <https://doi.org/10.17016/FEDS.2016.095>.

Mizrachi, Avi. 2016. "2016 – a Year of Institutional Adoption, Hype and Drama for Blockchain." Finance Magnates. December 26. Accessed at <http://www.financemagnates.com/cryptocurrency/education-centre/2016-a-year-of-institutional-adoption-hype-and-drama-for-blockchain/>.

Nakamoto, Satoshi, 2008. Bitcoin: "A Peer-to-peer Electronic Cash System: Accessed at <http://bitcoing.org/bitcoin.pdf>

NASDAQ. 2015. "NASDAQ Linq Enables First Ever Securities Issuance Documented with Blockchain Technology." Press release accessed at <http://ir.nasdaq.com/releasedetail.cfm?ReleaseID=948326>.

O'Dwyer, Karl J., and David Malone. "Bitcoin Mining and its Energy Footprint." Accessed at [https://karlodwyer.github.io/publications/pdf/bitcoin\\_KJOD\\_2014.pdf](https://karlodwyer.github.io/publications/pdf/bitcoin_KJOD_2014.pdf).

Paar, Christof, and Jan Pelzl. 2010. *Understanding Cryptography*. Berlin: Springer-Verlag.

Pieters, Gina, and Sofia Vivanco. 2016. "Financial Regulations and Price Inconsistencies across Bitcoin Markets." Accessed at <http://www.netlib.org/utk/people/JackDongarra/faq-linpack.html>.

Prisco, Giulio. 2016 "BitFury Announces Blockchain Land Titling Project with the Republic of Georgia and Economist Hernando de Soto", Bitcoin Magazine (April 27). Accessed at <https://bitcoinmagazine.com/articles/bitfury-announces-blockchain-land-titling-project-with-the-republic-of-georgia-and-economist-hernando-de-soto-1461769012>.

Raval, Siraj. 2016. *Decentralized Applications*. Sebastopol, California: O'Reilly Media Inc.

Reichert, Corinne. 2016. "ASX Completes Blockchain Trading Platform Prototype." ZDNet. September 28. Accessed at <http://www.zdnet.com/article/asx-completes-blockchain-trading-platform-prototype/>.

Rizzo, Pete. 2015. "Hands On with Linq, NASDAQ's Private Markets Blockchain Project." *CoinDesk*, November 21. Accessed at <http://www.coindesk.com/hands-on-with-linq-nasdaqs-private-markets-blockchain-project/>.

Santo, Atsushi, Ikuo Minowa, Go Hosaka, Satoshi Hayakawa, Masafumi Kondo, Shingo Ichiki, and Yuki Kaneko. 2016. "Applicability of Distributed Ledger Technology to Capital Market Infrastructure." Japan Exchange Group Working Paper, vol. 15. August 30.

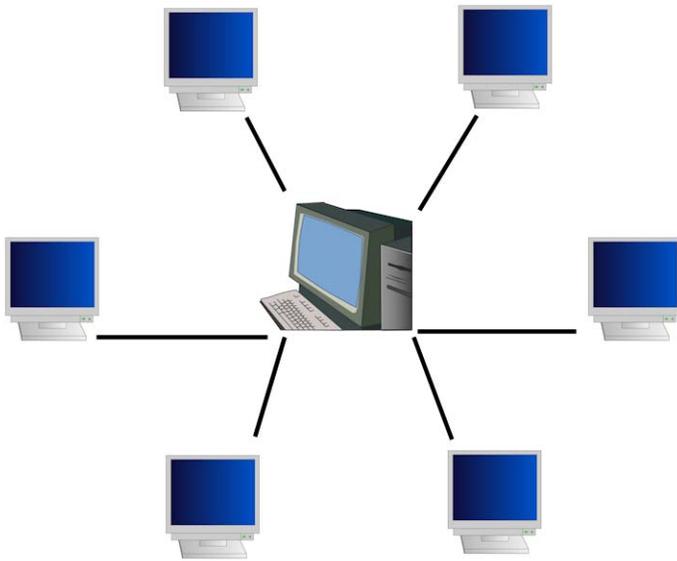
Swan, Melanie. 2015. *Blockchain: Blueprint for a New Economy*. Sebastopol, California: O'Reilly Media Inc.

Tapscott, Don, and Alex Tapscott. 2016. *Blockchain Revolution*. New York: Portfolio/ Penguin.

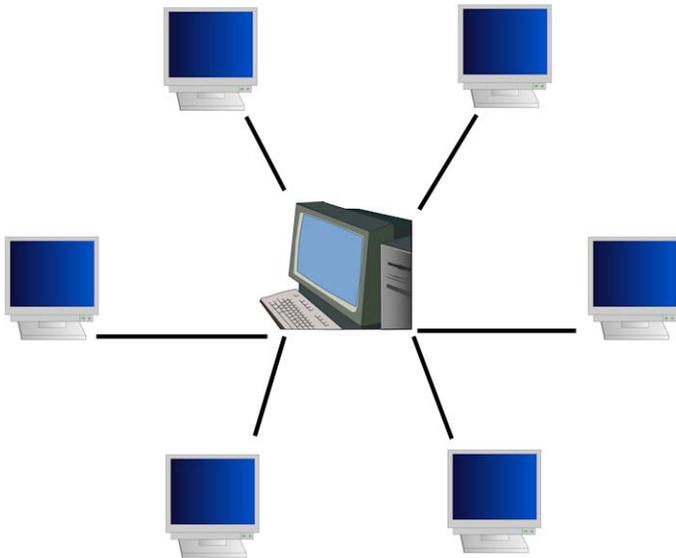
Wattenhofer, Roger. *The Science of Blockchain*. Zurich: Inverted Forest Publishing.

World Bank. 2016. "Remittance Prices Worldwide." Accessed at <https://remittanceprices.worldbank.org/en>.

Figure 1: Client-Server and Peer-to-Peer Networks



(a) Client-server network



(b) Peer-to-Peer Network

Figure 2: JPX Network

