INSTITUTE FOR DEFENSE ANALYSES

On Blockchains: The Hype, the Technology, and Potential Applications for the Department of Defense

Steven P. Wartik, Project Leader
Michelle G. Albert

January 2024

IDA Product 3001138

INSTITUTE FOR DEFENSE ANALYSES
730 East Glebe Road
Alexandria, Virginia 22305

The Institute for Defense Analyses is a nonprofit corporation that operates three Federally Funded Research and Development Centers. Its mission is to answer the most challenging U.S. security and science policy questions with objective analysis, leveraging extraordinary scientific, technical, and analytic expertise.

About This Publication

This work was conducted by the IDA Systems and Analyses Center under contract HQ0034-19-D-0001, Project C520824, "Blockchain Applications to DoD," for the IDA. The views, opinions, and findings should not be construed as representing the official position of either the Department of Defense or the sponsoring organization.

Acknowledgements

Abdullah Naimzadeh and William S. Cunningham for their technical review of this paper.

For More Information

Steven P. Wartik, Project Leader
swartik@ida.org, 703-845-6646703-845-6646

Margaret E. Myers, Director, Information Technology and Systems Division
mmyers@ida.org, 703-578-2782

Copyright Notice

Rigorous Analysis | Trusted Expertise | Service to the Nation

Executive Summary

Blockchains are best known as the technology underpinning cryptocurrencies such as Bitcoin, but their potential uses range far beyond that. Additional blockchain applications are still emerging, but it seems that blockchains are becoming a foundational technology that may affect how the Department of Defense (DoD) operates.

Blockchains record transactions and other data in a secure, distributed ledger. Every user has a copy of the ledger, which helps ensure the integrity of the transactions and makes the ledger auditable, though blockchains can be configured to let users identify only the transactions to which they were a party. This ensures privacy for both users and transactions. A blockchain is comprised of blocks that users add to the overall ledger. Each block contains three items: the transactions associated with the block; a hash value associated with the previous block; and a cryptographic nonce, which is a value used to compute the hash. It is virtually impossible to undetectably modify the content of a block. As the hash value of a new block depends on the hash value of the previous block, any attempts to modify a transaction in a block would create a hash that does not fit with the previous sequence. This feature guarantees the immutability of data placed on a blockchain.

Many well-known blockchain applications, such as Bitcoin and other cryptocurrencies, employ a public model that has no barriers to entry. Anyone can participate in transactions, and every user can validate transactions. Users are anonymous or pseudonymous. This model creates trust in the system—the blockchain—rather than in any particular user. But not all blockchain applications are open. A private blockchain model requires users to get permission to access the ledger and sometimes to perform certain operations on the ledger. This model relies on a more traditional trust system: trust in the administrator granting permission and trust in the users who receive permission. Both models have benefits, though a permissioned system would likely provide a better fit for DoD needs and policy.

Blockchains employ different consensus models, which provide protocols for adding new blocks to a blockchain. The proof-of-work model, which Bitcoin uses, is based on the idea that the right to add a block should be earned by demonstrating that effort has been made. This effort, known as *mining*, requires the use of computational power to solve mathematical puzzles to verify transactions. Profitable, industrial-scale mining uses huge amounts of energy and is inefficient by design, which makes the proof-of-work model a poor choice for military purposes, especially in theater when real-time information is

needed and energy sources may be scarce. Other consensus mechanisms, such as proof of stake, require substantially less energy but still operate in near-real time, rather than real time. Mining or managing a blockchain is likely better suited for planning purposes rather than operational use.

However, there is a difference between having to manage or mine a blockchain and making use of blockchain infrastructure. Some blockchain applications provide functionality on top of the base infrastructure,[1] which can allow users access without imposing mining or other resourcing burdens. There are a number of potential blockchain applications that DoD may consider.

Supply chains, for example, present myriad security risks. DoD acquires products and systems comprised of multiple parts, many of which are complex systems in their own right. These products and systems often travel to multiple locations and pass through many hands before they are ready for DoD use, which leaves them open to potential tampering. Software supply chains are especially vulnerable, as software applications often use code pulled from open-source libraries, which is extremely difficult to vet. Blockchains can help track the provenance of each part or piece of code manufactured in a supply chain, providing a path from design to delivery.

Modern database management systems support write-once information storage for command and control (C2) and command, control, and communications (C3) systems, but they require a data administrator with special permissions. If an adversary breaches a system and assumes an administrator role, they could tamper with or delete data. System logs give an indication whether data has been changed, but that is a reactive approach. Blockchains provide a proactive one: Any attempt to delete or modify an existing record on a blockchain yields an incorrect hash value, which flags illegal modifications immediately.

Blockchains and smart contracts could help DoD automate certain workflows. Smart contracts are programs that, when executed, automatically carry out a preset sequence of events. For DoD, smart contracts could support workflows that require meeting specific conditions before moving to the next phase. Acquisition is one example: Smart contracts could facilitate a workflow by triggering a notice for review once the required criteria have been created and recorded. Contracting processes are another. A smart contract may generate payment for services rendered once proof of completion is received and verified.

Electronic health records (EHR) pose interoperability issues among different healthcare providers and even with patients trying to access their data. The DoD has seen

---

[1] A blockchain's base network and underlying infrastructure (layer 1) can support additional functionality (known as layer 2). Layer-2 functionality relies on layer 1 for security and a consensus mechanism. See https://academy.binance.com/en/articles/what-is-layer-1-in-blockchain for more information.

success with its EHR modernization program, but interoperability outside the Department is still an issue. Moving patient records from DoD to the Department of Veterans Affairs (VA) when military personnel retire is cumbersome. Blockchain-based EHRs are meant to support interoperability, easing the administrative burden of sharing records and allowing patients easier access to their information.

Blockchains may also provide a useful tool for tracking media use and potentially countering disinformation. A blockchain's distributed ledger serves as an immutable archive that can be searched and used to prove data authenticity and integrity. Using a blockchain to track DoD-generated images, audio, and video could make it easier to determine if and when such media has been manipulated. This type of system could be part of a suite of tools to help debunk attempted disinformation campaigns and build trust in DoD as the source of the original media.

Blockchains can support system and information confidentiality, integrity, and availability, also known as the *CIA triad*. The DoD places high value on attaining and maintaining systems with those traits. Blockchain-based systems can securely store data, support data sharing and interoperability among organizations, help automate workflows, and record and track item provenance, all of which could provide increased efficiencies across the Department.

Table of Contents

Table of Figures

1. Introduction

Blockchain and the cryptocurrencies it enables surely rank among the most-hyped technologies of the 21st century. Bitcoin was announced to the world in a 2008 paper, and pundits seized on it, promoting it as the future of currency during its booms and disparaging it as just another fad during its busts. Gartner expresses the state of a technology using a "hype cycle," depicting inflated expectations, subsequent disillusionment, and eventual realization of practical use [1]. Gartner has published several hype cycles for blockchain. All clearly show, based on extensive research, that much of the excitement surrounding blockchain has yet to yield tangible positive results and that quantifiable justifications for some blockchain applications are still years away.

Whatever the future of cryptocurrencies, it is important to understand them and the technologies that implement them, especially blockchains. Blockchains are best known as the ledgers that record cryptocurrency transactions, but they have applications beyond cryptocurrencies. A *blockchain* is a technology for securely recording any kind of data that guarantees with an immensely high probability that the recorded data has not been corrupted and that any corruption can be easily detected. A blockchain can provide this guarantee across every node in a network. All participants can have a copy of an entire blockchain and examine every bit of data recorded on it. At the same time, a blockchain can preserve the privacy of its participants. Only those who place data on a blockchain know they placed that particular data. Everyone else can see the data, but they remain unaware of the actors involved. And, because every blockchain user has their own copy, the absence of a few users (caused by, say, system failure) does not affect the other users' ability to make use of the blockchain.

In other words, a blockchain can be used to achieve:

- Confidentiality: Information is protected from access by unauthorized users.
- Integrity: Information at rest is protected from change, and information in motion is received in the same state in which it was sent.
- Availability: Information can be accessed as needed.

These three properties are known as the *CIA triad*[2] and represent essential characteristics of a modern, secure system. In this age of nation-state hacking, the ability to build systems with the CIA traits cannot be overvalued.

The Department of Defense (DoD) puts tremendous effort into ensuring its systems achieve the CIA triad. Incorporating blockchains into DoD processes and systems where applicable may help safeguard sensitive data and support strategic and operational planning.

A. Purpose

There has been a lot of hype around blockchains—Gartner's blockchain hype cycles, mentioned earlier, provide an illustration—but blockchain is a relatively new technology, and the full scope of blockchain applications has not yet been realized. This paper attempts to cut through the hype and look at blockchain not as a basis for cryptocurrencies, but as a foundational technology that may affect how DoD does business now and in the future. This paper presents an overall survey of blockchains, covering the history of blockchains, blockchain concepts and technology, blockchain types and associated strengths and weaknesses, and potential applications for DoD.

[2] It is not known who coined the phrase, but the first paper to use the three words together [57] appeared in 1977.

2. Blockchain Concepts and Principles

The purpose of a blockchain is to create a record of a series of transactions.

There is a lot to unpack in that simple sentence, as blockchains touch on history, culture, trade, business, and technology. This section will cover each of these areas as they relate to blockchain, with the intent of showing laws, policies, and practices that motivated blockchain, along with how blockchain implements these laws, policies, and practices or, in some cases, circumvents them.

A. Ledgers

A blockchain can be most simply defined as a distributed ledger. A *ledger* is an ancient concept. According to the Oxford English Dictionary (OED), a ledger is "the principal book of the 'set of books' ordinarily employed for recording mercantile transactions."[3]

The OED's first citation for this sense of "ledger" is 1588, but the concept is much older. The earliest known recorded mercantile transactions occurred in the 4th century BCE.[4] Predictably, ledgers have evolved since that time, although the fundamental data remains the same: a collection of records in which each record describes an item, a quantity of that item, and an amount paid or received for that item in consequence of a transaction. Simple innovations include recording the transaction date and time (not so obvious in 3000 BCE) and standardizing item nature and quantity (e.g., "50-lb sack of barley" instead of "sack of barley").

Ledgers are, historically, maintained by accountants. The owner of a business, if not themselves an accountant, faces an age-old question: How do I know whether my accountants are honest? In the absence of safeguards, it is no difficult matter for an accountant to falsify ledger entries, hoarding material sold or currency received for themselves.

The basic solution, it would appear, took over four millennia to conceive fully but has proven remarkably durable since its introduction in medieval Italy [3]. An organization

---

[3] See https://www.oed.com/view/Entry/106898?rskey=ZM7gHs&result=1&isAdvanced=false#eid, definition A.1.d (retrieved 2020-01-07).

[4] The Cuneiform Digital Library Initiative (https://cdli.ucla.edu/) has an online selection of ancient ledgers. See, for example, https://cdli.ucla.edu/search/archival_view.php?ObjectID=P000735, which describes grain transactions. (URLs retrieved 2020-01-07.)

establishes double-entry bookkeeping, wherein every transaction requires an entry in two ledgers. An account entry in one ledger must have a corresponding and opposite entry in another ledger. If one business ledger records the business received 50 ducats, there had better be an entry in another ledger stating that those 50 ducats were deposited somewhere. Account auditing is suddenly much easier because the two ledgers can be compared. Inadvertent errors are undetectable only if two parties make exactly the same mistake, which is unlikely. Fraud requires complicity between two parties. Complicity, while hardly inconceivable, entails more effort and more risk.

Discouraging (if not eliminating) complicity relies on ledgers being under the authority of different parties. In double-entry accounting, there is one ledger for credits and another for debits. If the same individual maintains both ledgers, simultaneously manipulating both is simple enough. But if different individuals are responsible for each ledger, fraud requires complicity. The greater the difficulty for the individuals to communicate, the harder it is to establish complicity. It follows that the larger the business, the easier it is to establish an environment reducing the likelihood of complicity.

1. Distributed Ledgers

In the modern world, every business uses ledgers. If law enforcement agencies suspect a business of illicit actions, they can use ledgers of other businesses with whom the alleged offender has conducted transactions to determine monetary inflows and outflows and verify whether the inflows of one business match the outflows of another. The double-entry accounting principle holds for transactions between businesses as well. Figure 2-1 depicts an example of double-entry accounting for a manufacturing business borrowing money from a bank. The bank has a ledger of assets; someone debits this ledger, and someone else simultaneously credits the ledger of withdrawals. Then, when the actual exchange of funds occurs between the two businesses, someone at the bank debits the withdrawals ledger, and someone at the business credits the cash-received ledger. Next, somebody debits the cash-received ledger, and somebody else at the business credits the assets ledger. Along with timestamps on each ledger entry, there now exists a sequenced list of the transactions that led to the business receiving cash. The bank can audit what happened to the funds (they were withdrawn), the company can audit what happened to the cash it received (it has become part of the company's assets), and investigators can audit financial transactions (the business received money from the bank).

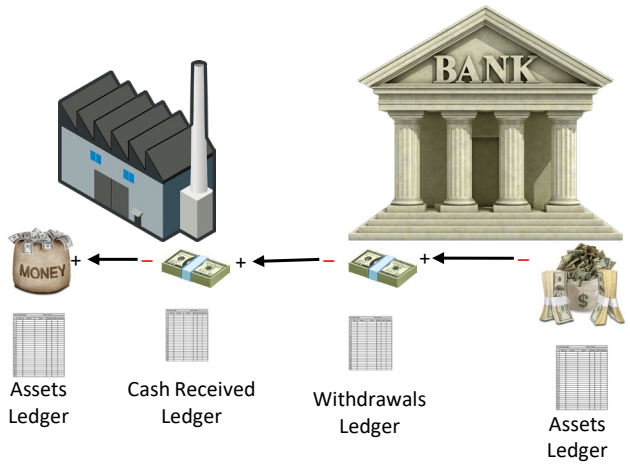

Figure 2-1. An Inter-Business Transaction

Having received money, the business puts it to use, spending it internally (paying employees, say) or externally (buying new equipment or purchasing raw materials). In either case, the resulting transactions will be recorded on other ledgers—some within the business and some outside it.

The world, then, operates as a collection of distributed ledgers: distributed within businesses and distributed across businesses. It is this insight that brings us back to blockchains. A blockchain preserves a record of one or more transactions. It ensures transaction integrity: A transaction is only recorded if it occurred. It also ensures auditability: An auditor can examine the blockchain to determine all transactions it records. At the same time, a blockchain intends to guarantee anonymity: Agents examining a blockchain can only identify transactions in which they were a party. No one can determine both parties involved in a transaction. There is no central blockchain administrator or "superuser" with special authority to identify participants.

B. Blockchain Basics

A blockchain is an ever-growing sequence of blocks in which each block is related to the previous block in the sequence. Figure 2-2 shows the start of a blockchain. The green box represents the first block, sometimes known as the genesis block. This block is generated to serve as a starting point. Black boxes represent subsequent blocks. The direction of the arrows means the contents of a block partially derive from the previous block in the chain. Specifically, a previous block's bit-string representation is used to compute a value that is stored in the next block. This is not to be confused with a linked-list data structure in which the next block in the chain can be determined but the previous block cannot.

Figure 2-2. A Simple Blockchain

Each black block encapsulates a record of one or more transactions; or, more generally, a record of one or more instances of what the blockchain records. Blockchains were originally developed to support transactions involving cryptocurrencies, and it is still commonplace to think of each block recording transactions. As subsequent chapters will show, blocks can record things other than transactions. This chapter will refer to transactions to keep the discussion focused.

1. Block Structure

In the simplest form, a block contains three items:

1. The transactions associated with the block.
2. A hash value computed from the previous block. (The first block in the chain does not, of course, compute a hash value from the previous block, but a hash value is generated specially based on the blockchain's implementation.)
3. A cryptographic nonce whose purpose is explained below.

Figure 2-3 depicts these items in two blocks, at the point where the *n*th block has just been added to a blockchain. The transactions are the record of whatever transactions an entity wants to record on this blockchain, and in a format specific to the blockchain. The previous hash value is computed by converting the contents of block $n-1$ to a string and applying a cryptographic hashing function to that value. All three items in the block are used as the input to the hashing function.

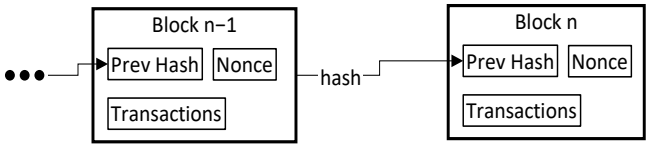

Figure 2-3. Block Content

A cryptographic hashing function is, by definition, fairly simple to compute but very hard to reverse. The value of Prev Hash in Block *n* can be computed quickly from Block

$n-1$—in microseconds or less on today's fastest computers[5]—but the value of Block $n-1$ cannot be determined from Prev Hash except by exhaustive search, a process likely to end after the universe does.[6]

This guarantees non-repudiability. It is not possible, or at least extremely improbable, to undetectably modify the content of a block. If a transaction in Block $n-1$ is altered, its hash will not match the value recorded in Block n. Users must preserve Prev Hash or it will not match the hash of Block $n-2$, and timestamps are ordered, so options are limited. As such, users can only modify the nonce in addition to transactions, and doing so requires an exhaustive search.

2. Mining

Each block contains a nonce, a 32-bit value used in computing the block's hash. A nonce is hard (e.g., time consuming) to discover by design. Requiring time to discover the nonce ensures that blocks are infrequently added to a blockchain, giving transactions and new blocks time to be propagated throughout the network. All ledgers have an opportunity to add the new block to their chain.

For blockchains using the proof-of-work consensus mechanism (see Section 2.E.1 for more information), the mining process works as follows. Suppose a blockchain has n blocks, and therefore every ledger in the blockchain has an identical copy of it (at least, the block headers[7] are identical). Certain nodes in the network conduct transactions and broadcast them to every other node. When a preset number of transactions are received or a predefined time limit is reached, nodes form a Merkle Tree[8] from the transactions, then hash the Merkle Tree's root (itself a hash), the hash of block n, and a nonce—with the requirement that the hash value's bit representation must start with some number of zeros. Of these three quantities, the nonce is the only one a node can vary, so mining is the task of finding a suitable nonce. The more zeros needed, the more work is required to find a suitable nonce. The first node to find a suitable nonce gains the right to add a new block to the blockchain. It broadcasts the new node to all nodes in the network. These nodes add

---

5 See https://en.bitcoin.it/wiki/Non-specialized_hardware_comparison#CPUs.2FAPUs.

6 See https://bitcoin.stackexchange.com/questions/41829/wont-asic-miners-eventually-break-sha-256-encryption/41842.

7 A block header contains metadata about the block and (except for the first block in a blockchain) a cryptographic hash that links the block to the previous block's header. Transactions stored in a block are not part of the header. See Figure 2-3 for an illustration.

8 The National Institute of Standards and Technology (NIST) defines a Merkle Tree as "a data structure where the data is hashed and combined until there is a singular root hash that represents the entire structure" (https://csrc.nist.gov/glossary/term/merkle_tree). See Section 2.B.3 for more about Merkle Trees, including an illustration.

the block to their blockchains, and the process begins again as the nodes await new transactions.

The requirement that a nonce begins with a certain number of zeros is the key to making mining difficult. Insofar as is known, the only way to find a suitable nonce is to try all possible values. A hash being apparently random, it follows that each bit of the hash is random, and so the probability of the hash beginning with 0 is ½; the probability of the hash beginning with 00 is ¼; and the probability of the hash beginning with n zeros is $1/2^n$. If the number of zeros is small, a computer performing mining can expect to find a suitable nonce easily, but if the number of zeros is large, a single computer is unlikely to find a suitable nonce quickly or even tractably. However, a blockchain network comprises many nodes, a large number of which engage in mining. The number of zeros can be dynamically adjusted such that the probability of at least one node "mining" a nonce in a suitable time is high. Bitcoin's objective is for a node to mine a nonce every 10 minutes. As of this writing, a Bitcoin hash is a string of 256 bits that must begin with 56 zeros (or more conveniently, 64 hex digits,[9] the first 19 of which must be zero). Figure 2-4 shows real-time results from a web page that reports Bitcoin blocks as they are created.[10] The time needed to find nonces and create blocks varies widely—sometimes it is less than 10 minutes, sometimes more.

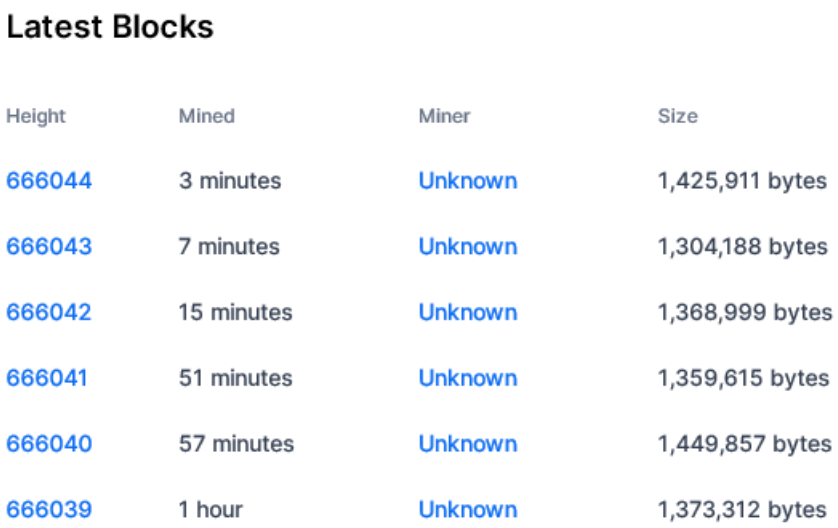

Latest Blocks

Height	Mined	Miner	Size
666044	3 minutes	Unknown	1,425,911 bytes
666043	7 minutes	Unknown	1,304,188 bytes
666042	15 minutes	Unknown	1,368,999 bytes
666041	51 minutes	Unknown	1,359,615 bytes
666040	57 minutes	Unknown	1,449,857 bytes
666039	1 hour	Unknown	1,373,312 bytes

Note. Image scraped from Blockchain.com as of January 15, 2021.

Figure 2-4. Real-time Bitcoin Mining Results

---

[9] Hexadecimal is a numbering system using base 16. This system uses the digit values of 0-9 and A-F, respectively, to represent numbers from 0–15. Combinations of digits are used to represent larger numbers. See https://www.techtarget.com/whatis/definition/hexadecimal for more information.

[10] The image is scraped from https://www.blockchain.com/explorer, a site that reports Bitcoin statistics.

3. Distributed Blockchains

If a blockchain is a kind of ledger, and good accounting practices necessitate multiple ledgers, the next problem is to figure out who maintains the ledgers. Once again, the issue boils down to trust. A blockchain-based ledger may be secure, but any user who submits a transaction to someone who maintains a blockchain-based ledger needs the following guarantees:

1. Their transaction is not viewed or modified en route.
2. The maintainer of the ledger does not modify the transaction before entering it in the blockchain.

Regarding item 1, any transaction submitted over an unsecured network is vulnerable to interception. The user and the receiver could set up a cryptographic exchange infrastructure, which is in fact how most transactions occur over the Internet. Establishing the infrastructure is costly, time-consuming, and requires considerable technical knowledge. It also requires the user to trust in the receiver, though there have been enough password breaches and infrastructure failures to demonstrate that such trust is unfounded. For this same reason, item 2 is hard to guarantee. Surreptitious malware on a company's web server, and for that matter unscrupulous employees, have been used to corrupt transactions.

The blockchain solution to these potential problems is straightforward: Every user has a copy of the blockchain. There is no centralized authority. If anyone questions whether a user was involved in a transaction or the details of that transaction, the user has the data to respond definitively.

Blockchain implementations address item 1 by using a variant of the block structure shown in Figure 2-3. A block is split into header and data sections. The header contains a hash of the data section. The validity of the blockchain is therefore tested by examining only the header, not the entire block.

The data section, which records the transactions, is organized as a Merkle Tree [3]. In these Merkle Trees, each leaf is a hash of a transaction and each interior node is a hash of its children. Merkle had the interesting observation that, for cryptographic signature purposes, all one really cares about is the hash value of the root node. Other nodes in the tree can be deleted if their leaves are not of interest. Once a block is no longer the head of a blockchain, users are typically only interested in transactions in which they participated. On a given computer, the majority of blocks will only contain a header.

The left side of Figure 2-5 shows a block containing three transactions. The hash of each transaction is a leaf node. The hash of the hashes of Tx1 and Tx2 is an interior node, and that hash, combined with the hash of Tx3, is the root node. On the right, two transactions

have been pruned from the tree. The tree still contains enough information to compute the root node's hash and therefore allows verification that Tx3 was properly placed in the block.

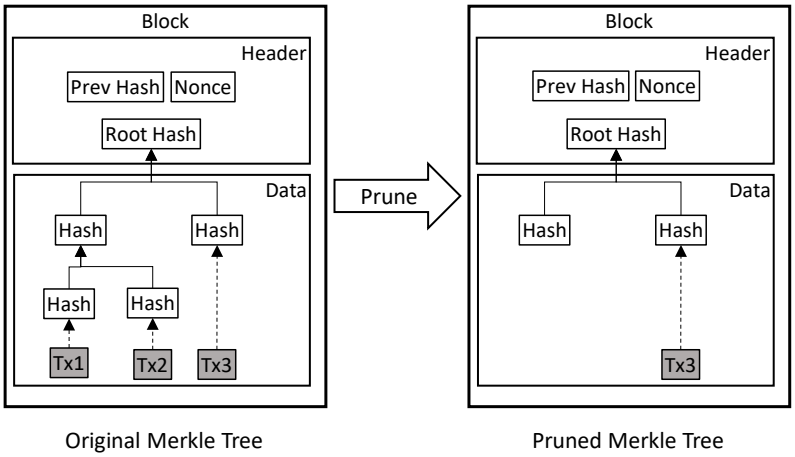

Figure 2-5. A Block as a Merkel Tree

In the original Bitcoin implementation, a header was about 80 bytes [4]. Based on the assumption that a new block is added to the blockchain every 10 minutes (see Section 3.A for more about transaction times), many transactions can be added to a blockchain before storage becomes a problem.

Propagating transactions to every blockchain user is another conceptually simple but potentially massive undertaking. The technical term is network broadcasting, and IP-based communication offers a mechanism to send a packet to every node in a network. The maximum latency depends on the degree of node connectivity within the network; but, as blockchains typically use the Internet, the latency is likely no worse than the underlying Internet infrastructure. As the most well-known blockchain, Bitcoin provides an idea of transmission magnitude. The number of nodes in the Bitcoin network has fluctuated wildly, with a high of over 200,000 in early 2018 [42]. As of April 2023, the network had approximately 17,800 reachable nodes.[11] Even at its largest, there were no reports of Bitcoin problems due to latency.

4. Blockchain Infrastructure

Setting up a blockchain is not particularly difficult, and there are many free and open-source blockchain implementations (see Appendix A for examples). The Openchain

[11] The website https://bitnodes.io/nodes/ was developed to provide statistics on the Bitcoin blockchain. It reports, among other things, the number of reachable nodes.

website, for example, promises "anyone can spin up a new Openchain instance within seconds."[12]

Simply creating a blockchain instance does not address whether using that instance will be practical. As such, it is worth examining some of the underlying considerations.

1. ***The underlying network.*** A blockchain is intended for peer-to-peer transactions. In the network on which it operates, a node should have multiple paths to communicate with other nodes. A network with a hub is a poor choice, because it introduces a single point of failure. Hardware failure aside, the greater the degree of network connectivity, the greater the difficulty of staging a denial-of-service attack.

2. ***The resources available for mining.*** Whatever approach is used to mine blocks, it is important to remember that a blockchain exists to support some larger objective (in the case of Bitcoin, financial transactions) and that mining drains resources that might otherwise be used to achieve that objective.

3. ***Storage.*** Section 2.B.3 discusses the storage requirements for a blockchain, notes that a single block header can be stored in a relatively small space, and describes pruning block data. The figures, however, are for Bitcoin and its blockchain. Some blockchains have larger headers,[13] and pruning transactions from the data sections in some blockchains would cause unacceptable information loss.

   Storage is more than writing to disk, of course, and it is important to know the expectations of analyzing stored data. Blockchains can be used to manage and enhance large-scale storage, but depending on the nature and structure of transactions, it may be worthwhile using a sophisticated storage mechanism such as a relational database management system with a powerful and expressive query language like Structured Query Language (SQL).[14] An index-based storage tool queried by NoSQL[15] might be more appropriate. Whatever the mechanism, both free and commercial versions exist, and money often buys improved performance, not to mention professional, on-demand support.

---

[12] See https://docs.openchain.org/en/latest/index.html.

[13] See [6] for a discussion of fields commonly included in block headers.

[14] SQL is a widely accepted relational database query language with standards specified by the American National Standards Institute (ANSI) and International Organization for Standardization (ISO), and it includes vendor-specific features in almost every implementation.

[15] NoSQL is an acronym with no single agreed-upon meaning; common ones are "Non-SQL" or "Not only SQL." It refers to a database management system whose primary, often only, query language is not SQL.

4. ***Security.*** Anyone implementing a blockchain should assume compromise is a possibility, understand the implications of compromise, perform auditing to detect compromise, and have tools and a strategy to recover from compromise.

A blockchain guarantees its transaction data is preserved unmodified. It can also ensure that each transaction is executed exactly once.

However, there is no reason a blockchain could not be used to record something other than transactions. Subsequent sections go into more detail, but some of the other types of data that could be in a blockchain include:

1. ***Contracts.*** These are agreements between two or more parties stating that one party agrees to perform services for another party. In more advanced form, a contract can automatically call upon parties to perform the work and can be used to query the contract's status and the degree to which parties have fulfilled their obligations under the contract.
2. ***Unmodifiable records.*** Some systems, intending to support auditing, do not allow data to be modified once it has been entered. Logging software, for example, has functionality to append data to a log but offers no capability to delete or modify data. Most simple logging systems write text to a file. If an intruder modifies that text to erase records of their penetration, detecting the change can be difficult, if not impossible. If the records were stored in a blockchain, modifying any record would make the rest of the blocks invalid.

5. Blockchain as Infrastructure

Section 2.B.4 discusses the infrastructure a blockchain requires. From another perspective, a blockchain *is* the infrastructure. It exists to make some capability possible: financial transactions *à la* Bitcoin, secure storage, identity management, and many others. In all cases, a blockchain is used to decentralize a once-centralized capability. This architectural perspective is sometimes called Web 3.0 or Web3. Proponents of Web 3.0 argue that it represents the next step in the evolution of the World Wide Web from Web 2.0, which started around 2005 and resulted in companies like Amazon, Google, and Facebook (now Meta) centralizing user capabilities and providing access at the expense of privacy. Web 3.0, its proponents claim, will return the Internet to its original vision, one in which everyone controls their personal information and establishes peer-to-peer communications rather than relying on a middleman to coordinate, and potentially examine, information exchanges.

Web 3.0 is in its infancy. Companies that advocate it must devise incentives for using it. Often these incentives take the form of tokens that offer financial rewards for participation (Section 2.G). It remains to be seen whether incentives will catch on enough

to justify participation in a blockchain and generate sufficient revenue streams for companies promoting a blockchain.

C. Blockchain Types and Trust Models

In every transaction, each participant must decide how much they trust the other participants, and that trust may or may not be reciprocated. In this age of remote transactions, trust is very difficult to establish. Third parties promote trust—a bank, for example, can act as a trusted third party for two participants exchanging some amount of money—but they can lengthen transaction times and do not always support participant privacy.

Part of the draw of blockchains is the removal of the third party, though this raises concerns about how to incorporate trust into a transaction. There are four general types of blockchains: public, private, hybrid, and consortium. Trust and access vary among types, following a spectrum of no trust to much trust. In a blockchain, trust is expressed in terms of permission: who receives permission to perform operations on a blockchain and who grants that permission. Permission levels fall along a spectrum as well, from least permissioned (public) to most permissioned (private). Permission levels may vary among hybrid and consortium blockchains depending on the trust and access required.

1. Public, Permissionless

Public blockchains, as the name suggests, are open to anyone who would like to use them. They were developed for use throughout the Internet, allowing two arbitrary parties to safely and securely conduct transactions. Users of public blockchains are able to access current and past records, verify transactions, and conduct mining activities. This type of blockchain is decentralized and uses a consensus model (see Section 2.E) to validate the information exchanged among network nodes and added to the blockchain. Once a record or transaction is deemed valid, it cannot be changed.

Public blockchains are permissionless: No one needs permission to perform an operation. Everyone may mine, everyone may add blocks, and everyone may conduct a transaction. No one grants authority to another participant, and no one needs to ask for authority. Users operate anonymously or pseudonymously. It is almost impossible to verify a user's identity and thus gain some semblance of trust in the user prior to conducting a transaction [60].

This is a zero-trust model. No one assumes anyone else on the network is honest or will deal with them fairly. Use and observation of the blockchain itself must guarantee that transactions are proper—that is, that resources are transferred from the seller to the buyer, that the transfer is recorded and visible, and that only the transfer is recorded. All users have a copy of the ledger and can see each transaction. In this model, trust is replaced with

transparency, which leads to confidence in the system (the blockchain itself) rather than in its users.

Bitcoin uses a zero-trust model. Its objective was to create a financial network completely independent of central banks, and it has generally succeeded. Government organizations could also adopt zero-trust models in certain situations to create publicly available records. This type of blockchain application would give the public access to a government system and a view into government activities, which may in turn increase trust in both the institution and the information it provides.[16]

There are, however, some drawbacks to public, permissionless blockchains. Permissionless blockchains are slower and resource intensive. Large networks require more time and energy to verify transactions, and they do not scale well. A Bitcoin transaction, for example, takes about 10 minutes to fully execute.[17] Also, the lack of restrictions makes permissionless blockchains susceptible to 51% attacks, in which more than 50% of miners in a blockchain join forces to control it (see Section 3.C for more information).

2. Private, Permissioned

Private blockchains are more restrictive and permissioned, specifying roles users may have and granting authority (often to a single entity) to assign, modify, or revoke those roles. Permissioned blockchains reintroduce the concept of trust: trust in the entity administering the blockchain and granting permissions and trust in the users, as they are known.

The first permission granted is the ability to join a blockchain. Some blockchains have a security infrastructure that only accepts vetted users, such as those willing to reveal their identities. Ripple is one example.[18] Another permission is the right to create blocks. All users can conduct transactions, but users who cannot create blocks must use an intermediary participant. In a blockchain like this, users generally retain the ability to read blocks and verify that their transactions were recorded in the leger. However, the right to read blocks may, in some circumstances, be limited. Users may be prohibited from knowing anything about transactions other than their own, or they may be limited to reading particular blocks.

---

[16] Estonia started using blockchains to support government systems in 2012. One example is its use of blockchain for land title registries, which provides Estonian citizens with an accessible and verifiable record system (see https://www.pwc.com/gx/en/services/legal/tech/assets/estonia-the-digital-republic-secured-by-blockchain.pdf for more information).

[17] See https://coinmarketcap.com/alexandria/article/how-long-does-a-bitcoin-transaction-take. During periods of frenetic transactions, a single transaction has been known to take over 24 hours.

[18] Ripple is a corporation that runs a blockchain that supports transactions for the XRP cryptocurrency. See Appendix A and https://ripple.com.

On the other hand, a permissioned blockchain may grant users extra rights. It may reveal all user identities, for example, or the identities of specific users. It may also have open and closed transaction periods in which transactions are visible during open periods and restricted during closed periods.

These characteristics make permissioned blockchains well suited for enterprises where a single entity is responsible for determining access, authorizations, permission levels, and security. This model can assume some degree of trust. But that does not mean permissioned blockchains must be limited to enterprises. Ripple's business model, for example, aims to appeal to anyone with an Internet connection, though users must decide whether they are willing to engage with Ripple, knowing that their identity will not be kept private.

3. Hybrid Permissions

Hybrid blockchains comprise elements of both public and private blockchains. A hybrid blockchain often uses permissionless and permissioned systems in parallel, with the permissionless system allowing transparency among nodes—every node has a copy of the ledger—and the permissioned system controlling access to the ledger (e.g., the ability to add or validate blocks). Transactions and data are kept private, but they can still be verified when needed through means such as a smart contract (see Section 2.H). Users are granted full access to the network, and they are able to remain anonymous until they engage in a transaction. At that point, their identity is revealed to the other party in the transaction.

4. Consortium

Consortium blockchains, also called federated blockchains, are similar to hybrid blockchains, but whereas hybrid blockchains are controlled by a single entity, consortium blockchains are run by a group. Group members control access to the network and user permissions. Consensus procedures on consortium blockchains are controlled by preset nodes. A validator node starts, receives, and validates transactions. Member nodes start and receive transactions. All nodes have a copy of the full ledger.

D. Privacy

Blockchain-based transactions create a new model of privacy. If two individuals wish to conduct a traditional transaction, they might employ a trusted intermediary such as a bank. In this case, each must make their identity known to the bank, if not to each other. (Some organizations, such as auction houses, allow transaction parties to remain anonymous.) The individuals and the bank have a record of the transaction, and the individuals trust the bank's security infrastructure to conceal the transaction from anyone

without a need to know, though applicable laws may require the bank to surrender information.[19]

If individuals want to keep their transactions private, they can conduct them using cash or barter, which in theory leaves no record, but in practice is less straightforward. Except for trivial transactions, both parties would likely want to record entries in their private ledgers. This shifts the burden of maintaining security to the individuals, who may be less well equipped to establish a sophisticated security architecture than a large organization like a bank, even if a bank is a juicier target. (Willie Sutton robbed banks "because that's where the money is.")[20]

Bitcoin's use of blockchain opts for a different model. Every transaction is public, but the identities of the individuals who conduct the transaction are not. It is as if everyone can see all existing cash and observe every transfer between accounts, but without any knowledge of who owns the accounts. There is no such thing as an anonymous transaction, because everyone knows the identity of the coin that was transferred, even if they do not know the sender or recipient. Only the individuals who conduct a transaction are aware of the identities, and each individual only knows their own. This information can be used to construct a comprehensive ledger of one's own transactions.

The act of saving data by pruning a Merkel Tree (Section 2.B.3) is also closely tied to Bitcoin's privacy model. Each Bitcoin transaction records two parties exchanging bitcoin, and coin ownership changes as a result of the transaction, just like with physical coins. And, just as a dime that changes hands cannot be traced to its previous holder,[21] once the Merkel Tree recording a Bitcoin transaction is pruned, there is no way to determine the coin's previous owner. This lack of transparency is a deliberate design feature to provide privacy; it is not an inherent property of blockchains. Neither is the use of a Merkel Tree, though using them means the hashing operation always has a fixed-size input, which increases the predictability of the time required to complete a transaction.

As such, there are some limitations to privacy in blockchains, particularly those that lack pruning methods similar to what Bitcoin employs.

---

[19] The Currency and Foreign Transactions Reporting Act, also known as the Bank Secrecy Act, requires banks to report every deposit or withdrawal over $10,000. (31 U.S.C. §5311 et seq. See https://www.fincen.gov/resources/statutes-regulations.)

[20] See https://www.fbi.gov/history/famous-cases/willie-sutton#:~:text=When%20asked%20why%20he%20robbed,bank%20early%20in%20the%20morning for more information.

[21] A somewhat dubious claim in this age of DNA fingerprinting, but useful for the sake of the analogy.

1. Zero-Knowledge Proofs

Bitcoin achieves its desired level of privacy by hiding identities behind public/private key pairs.[22] The only record of a participant's identity is a public key. Moreover, Nakamoto [4] recommends that each participant generates a new public/private key pair for each transaction. This guards against associating a set of transactions with a single user, which in turn might lead to identifying that user through patterns of behavior, though this approach has proven insufficient [34] (see Section 3.F).

Zero-knowledge proofs have been proposed to address the problem. A zero-knowledge proof allows a party to demonstrate they possess some knowledge without revealing what the knowledge is. Imagine a sentry guarding a gate, challenging all comers to state a password before they enter. The problem with this time-honored method is that the act of stating the password makes it public. The knowledge is available to all within hearing range. What if, instead, the sentry could simply ask "Do you know the password?" and the visitor could reply "Yes, I know the password," and the sentry would know whether or not the visitor was telling the truth? ("Yes, I know the password" is not the password.) This eliminates the risk of revealing the password to an eavesdropper. This is the motivation behind zero-knowledge proofs.

There are immediate privacy benefits to using zero-knowledge proofs in blockchains. Transaction participants can authenticate themselves without revealing even their public keys. Zero-knowledge proofs could allow organizations or individuals to use private datasets to implement smart contracts (Section 2.H) without revealing the actual data. More directly, examining a blockchain should reveal the validity of all transactions on the chain. A zero-knowledge proof can be used to avoid revealing details about the transactions. The information proves the transactions have occurred and can be placed in a block.

E. Conflicts and Consensus Models

A successful blockchain is predicated on the assumption that new blocks are constantly added at regular intervals. Each block records one or more transactions. New transactions cannot be added to a block until a new block is ready. Every blockchain must have protocols and procedures for adding blocks.

When a node discovers a nonce while mining, it immediately creates a block based on that nonce and broadcasts the block to all other nodes. What happens if two nodes simultaneously discover a nonce? In a large distributed network, the time between when a node broadcasts a block and when all nodes receive that block can be significant. Even if

[22] Public and private keys are cryptographic keys used to encrypt data and create and verify digital signatures. Public keys, as the name suggests, are generally widely available and can be used to encrypt data for a specific recipient, who then decrypts the data using their private key (see https://csrc.nist.gov/glossary/term/public_key).

it is only a matter of seconds, some overlap can be expected in the course of generating 666,044 blocks (per Figure 2-4) in an 11,000-node network.

In that circumstance, the network would contain two (or more) blockchains with different blocks after some node. Nodes update their chains as soon as they receive blocks, so nodes "close" to one broadcasting node will receive one block, and nodes "close" to another broadcasting node will receive another block. The rules for resolving these conflicts are as follows:

1. A node extends its chain with the first block it receives.
2. Once a node receives a block, it immediately starts mining the next block.
3. If a node receives another block before it or another node mines the next block, it stores the block in its blockchain as another branch.
4. Nodes consider the longest branch the official branch. When a node detects that one branch has grown longer than other, it drops the shorter branch.

It is possible, but highly unlikely, that the two branches could grow simultaneously for several iterations. As Figure 2-4 shows, mining time variance is much larger than broadcast time. According to one study, conducted when the Bitcoin blockchain had 556,400 nodes, conflicts in the Bitcoin blockchain occurred in just 0.0012% of block broadcasts [5]. If the probability of a single conflict is so small, the probability of a blockchain in which conflicts cause two branches of the same length is vanishingly so.

(This is absolutely not permission for blockchain implementors to ignore the possibility of blockchains with branches. The previous paragraph referred to branches caused by mining conflicts. Branches can arise in other ways, in particular deliberate attempts to defraud a blockchain network; see Section 2.F.)

As such, there is always a possibility in a distributed network that two (or more) subnetworks will become disconnected and modify the blockchain independently for some period of time. How, then, to resolve the conflict when the subnetworks reconnect?

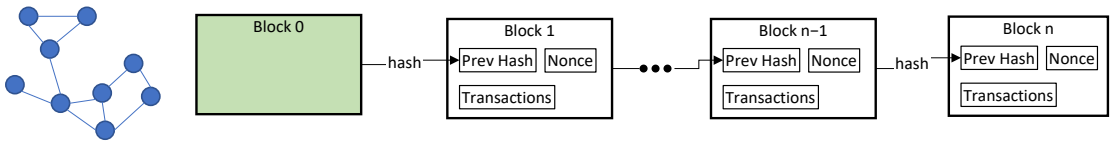

Figure 2-6. Blockchain After Adding *n* Blocks

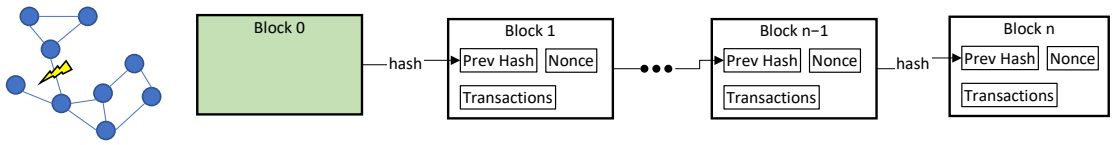

Figure 2-7. Network Split

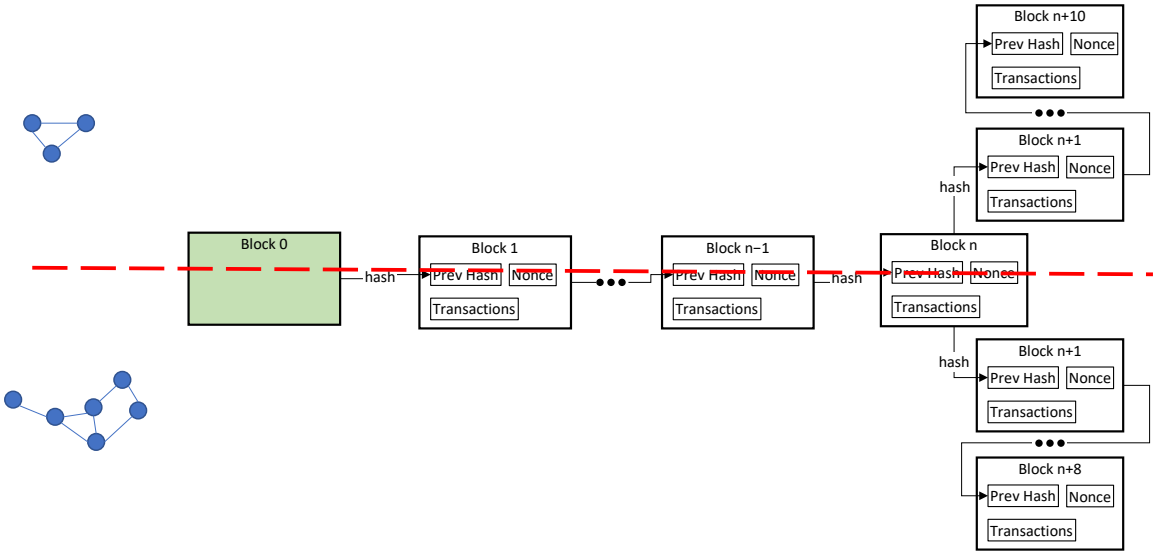

Figure 2-8. Blockchain Evolving Independently After Split

The figures above illustrate the situation: The blockchain grows to *n* blocks (Figure 2-6), the network is then partitioned into two subnetworks (Figure 2-7), and the blockchain evolves independently on each subnetwork (Figure 2-8). Keep in mind that every node has an independent copy of the blockchain. In Figure 2-8, the three nodes above the dashed line each have a copy of the first *n* blocks plus 10 others, whereas the six nodes below the line each have a copy of the first *n* blocks plus 8 others, and the set of 10 and 8 are distinct, with different nonces and transactions. When the subnetworks reconnect, what blocks will be created to record the transactions on both branches? And, not inconsequentially, will the miners on both branches be rewarded? Potential conflicts can arise for a simpler reason than disconnected subnetworks. Suppose two nodes mine blocks at about the same time. Notifying every node in the network is not instantaneous, and it is probable that some nodes will receive notice of both blocks before all nodes receive notice of one block—the latter being a signal that a new block can be unambiguously added to a chain.

Different blockchains have addressed the deconfliction problem by defining and implementing different consensus models. A consensus model lets users work together to agree on who has the right to add the next block to a blockchain and how to resolve conflicts (competing branches) in a blockchain. This section presents an overview of common consensus models.

1. Proof of Work

The proof-of-work consensus model is based on the idea that the right to add a block should be earned through demonstrating that effort has been made. This effort—known as mining (see Section 2.B.2)—requires the use of computational power to solve mathematical puzzles to verify a transaction. Bitcoin uses proof of work. Considering that Bitcoin is for financial transactions, and earning money requires work, Bitcoin's choice of consensus model is unsurprising. In a coin-based economy, minting coins ultimately requires mining metals.

Proof-of-work consensus resolves conflicts by having nodes keep copies of conflicting blocks and diverging chains until such time as one chain is longer than the other. This is a probabilistic approach. Given the difficulty of mining, it can be proven that the likelihood of divergent chains continually increasing in length at the same rate decreases exponentially with each new conflict.[23]

Once a conflict is resolved, nodes drop the shorter chain, merging its transactions. This approach requires each node to maintain a pool of transactions for all conflicted blocks. The node that mines the next block includes these transactions in the block.

2. Proof of Stake

The proof-of-stake model was developed in response to the heavy computational demands the proof-of-work model places on blockchain infrastructure [13,14], as well as the amount of electricity required to support mining operations. (See Section 3.B for more on energy consumption related to proof of work.) Briefly, the approach works as follows. A single node is randomly chosen from all nodes in the network. This node creates a new block, which it digitally signs; the signature, being hard to forge, is the guarantee of the block author's authenticity. One thousand more nodes are then selected and are responsible for verifying the new block's integrity. Once these nodes are satisfied, they sign the block, and the block becomes the end of the blockchain. Signing a block is not computationally difficult, nor is choosing nodes and propagating the block to those nodes.

The probability that a node will be selected is proportional to the node's investment. The nature of "investment" depends on the purpose of the underlying blockchain. If a proof-of-stake-based network is being used for cryptocurrency, investment is synonymous with number of coins. The more cryptocurrency a node possesses, the more often it is selected to create or verify blocks. This proof of stake in a blockchain earns the right to create blocks.

There are several proposed approaches for keeping a single node from having a high probability of selection. For example, in the Peercoin network [14], coins age if they are

---

[23] See [3], Section 11.

unspent, and the longer they are unspent, the higher the probability the node possessing them will be used to sign the next block. A single node possessing a sizeable percentage of coins would have to withhold them from circulation, negating the point of having wealth, which is to spend it. In practical terms, there is a growing disadvantage to being a majority stakeholder.

In September 2022, the Ethereum blockchain network switched from proof of work to proof of stake in an event known as "The Merge."[24] Users have to deposit 32 ether (ETH) to activate the software needed become a validator on Ethereum's network. This transition reduced Ethereum's direct energy consumption by 99% [70].

3. Proof of Capacity

In the proof-of-capacity consensus model, the right to mine cryptocurrency does not come from showing that work has been done, but instead from demonstrating a willingness to devote a portion of a disk drive to mining—that is, proving that one has storage capacity rather than central processing unit (CPU) cycles. (Proof of capacity was originally termed proof of space [15], but the former term seems to be more common.) Both express commitment to the growth of a blockchain network. Proof of capacity, however, was devised with the intent of requiring significantly less energy than proof of work—according to one source, it is 30% more efficient.[25]

Proof of capacity assumes nodes are divided into two categories: provers and verifiers. A prover sends proof to a verifier that the prover has the claimed capacity. Constructing the proof requires minimal computational resources, as does verifying it.

If this seems more like a client-server architecture than a fully distributed peer-to-peer network, that is explained by the background research. Proof of work originated from attempts to combat email-based spam [16]. An email server, it was thought, should accept requests from clients who want to send a small number of messages to a small number of recipients, but it should reject client requests to send large numbers of messages to thousands or millions of recipients. The scheme required more work for each recipient that a normal user would scarcely notice but a spammer would consider an unreasonable up-front cost. Analogously, in proof of capacity, one node acquires additional rights (mining or sending email) with respect to another by proving it has more disk space. Disk space is inexpensive but not free and, if the per-right amount required is large, gaining inordinate leverage with respect to other nodes becomes impractical.[26]

---

[24] See https://ethereum.org/en/roadmap/merge/.

[25] See https://www.investopedia.com/terms/p/proof-capacity-cryptocurrency.asp.

[26] Dziembowski et al. propose 100GB per email address for the right to send messages, although they never justify the figure. [14]

4. Proof of Elapsed Time

Proof of elapsed time was developed by the Intel Corporation to support the Hyperledger Sawtooth blockchain.[27] It uses Intel's Software Guard Extensions (SGX), a set of security-related instructions built into some modern Intel processors. SGX allows a process to execute in an *enclave*, a private region of memory that cannot be read or shared outside that process. An enclave is cryptographically encoded during its existence to prevent it from being observed.

Proof of elapsed time executes on a permissioned blockchain network with a central server. On start-up, a node, called a validator, registers itself with the server. For each registered validator, the server generates a random wait time and broadcasts that time to the node. The validator sleeps for that amount of time, wakes up, and notifies the server. The first validator to notify the server wins the right to add a block to the blockchain. At this point, the cycle starts again, with the server generating a new set of random wait times.

The server must use an algorithm that gives each registered validator an equal chance of waking up first. The server must also verify the validator is legitimate (i.e., that it did not simply sleep for a very short time). Any knowledge of how the server is operating, such as the seed it uses for random number generation, could be used to corrupt the mining process. The server executes in an enclave to minimize the chance of network corruption due to process observation.

Validator nodes need not actually sleep before notifying the server. They can switch to other tasks, and this is the claimed advantage of proof of elapsed time. No validator resources, CPU, or storage are 100% (or even close to 100%) devoted to mining.

5. Hybrid Consensus

Instead of using a single consensus model, block mining can use some combination of the examples listed above. This hybrid approach is not uncommon in consensus models other than proof of work. In the simplest case, proof of work is used to mine the genesis block, then some other approach is used to mine subsequent blocks. The genesis block should be hard to forge, because verifying a blockchain ultimately entails verifying the genesis block. Proof of work, as Section 2.E.1 shows, guarantees that property.

Hybrid consensus models have also been introduced into a mature blockchain for the purpose of switching consensus models. If a blockchain network's participants decide proof-of-work mining is consuming too many resources, they might switch consensus models by declaring that all blocks mined prior to some date are verified according to a proof-of-work algorithm, whereas all blocks mined after that date are mined and verified using some other algorithm. This is typically done by creating a fork (see Section 2.F) and

---

[27] See Appendix A and https://sawtooth.hyperledger.org/.

allowing transactions on the original blockchain to accumulate until the branch can be safely terminated.

The Peercoin cryptocurrency [14] uses both proof of work and proof of stake. Originally, nodes mined Peercoin using proof of work. This was done to build up a sufficient distribution of coins to make a 51% attack unlikely (see Section 3.C). Peercoin then switched to a hybrid model, in which both proof of work and proof of stake could be used for mining. Peercoin's designers assumed that proof of stake, being less resource-intensive, would predominate and slowly replace proof of work. Exact numbers could not be located, but a page on Peercoin's website says:

> Today the majority of blocks in Peercoin are created through proof-of-stake while a small minority are created through proof-of-work.[28]

F. Forks

Several sections have stated that a blockchain aims to be a single chain, not a tree. This is not strictly true. In some circumstances, blockchain participants may desire to have the chain diverge into separate paths with no intention to merge them. This is known as a *fork*.

(The previous paragraph is also not strictly true. Blockchains are still too new to have unambiguous, widely accepted terminology. Some definitions of "fork" include any split, including those eventually resolved by consensus [17]. However, this viewpoint appears to be in the minority.)

Blockchain forks come to be for several reasons:

1. ***Changes to blockchain software.*** Not every change requires a fork. However, if someone discovered a bug that could somehow compromise a blockchain—improper implementation of a protocol, for example—all subsequent transactions must be considered suspect.
2. ***Changes to blockchain protocol.*** This occurs when someone discovers a flaw in the protocol itself (as opposed to software implementing the protocol) that bad actors could exploit. Alternately, a new protocol might make the blockchain operate more efficiently. For example, a Bitcoin fork in 2014 increased transaction efficiency from seven transactions per second to 24 transactions per second.[29]

[28] See https://university.peercoin.net/#/11-economics-of-peercoin.
[29] See https://www.investopedia.com/tech/history-bitcoin-hard-forks/.

3. ***Changes to blockchain capability.*** Blockchain participants may determine they want future blocks to record different kinds of transactions or things other than transactions.

All three cases require changing blockchain software. And in all three cases, one does not simply cease adding blocks until all nodes are running the new software. For one thing, ceasing operation is impractical: in Bitcoin, where new blocks are added about every 10 minutes, waiting for a fix requires shutting down an important service for an unspecified period of time. For another thing, in a permissionless blockchain, no node can dictate actions to others. Finally, some nodes may not want to incorporate software changes. They may have a commitment to a sequence of blocks by the time they receive an update. They may not want to adopt a new protocol or provide a new capability.

Whatever the reasons, the consequence is a situation in which the software used by the subset of nodes that choose not to switch is incompatible with the software used by the subset of nodes that do (Figure 2-9). The former nodes cannot process any "new-style" blocks added after the fork, and the latter nodes cannot, or at least choose not to, process any "old-style" added after the fork. If a change is to the structure of blocks, the latter nodes must still be able to recognize the old structure (for history and verification), but they no longer need to know the old protocol (if that changed) as they will not be adding blocks under that protocol any longer.

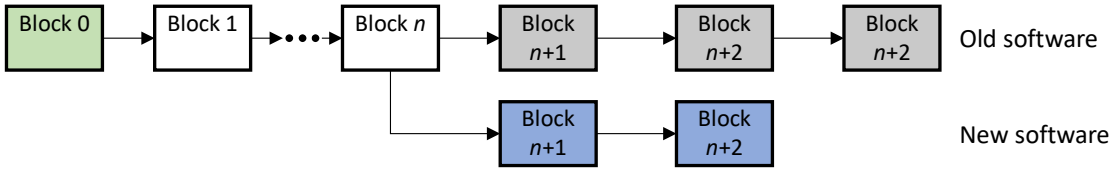

Figure 2-9. Blockchain Fork

In Figure 2-9, there is a shared history up to block *n*. After that, the two chains diverge, and the history of one is not the history of the other.

There are two kinds of forks: hard forks and soft forks. A hard fork is oriented toward switching completely to the new branch. The new software deliberately does not recognize transactions on the old branch or attempts to add blocks to that branch.

By contrast, a soft fork yields software that recognizes both. However, a soft fork is introduced with the expectation that all nodes will soon switch to the new branch. Whereas a hard fork fails to recognize all transactions on the old branch, a soft fork is written to achieve the same capability of the consensus model—except, of course, that the new branch is accepted as official irrespective of its length relative to the old branch. A soft fork preserves transactions.

One reason to introduce a hard fork is in response to fraud. In June 2016, hackers attacked the Ethereum blockchain, making off with approximately $50 million worth of cryptocurrency.[30] Ethereum blockchain users subsequently voted to create a hard fork just before the point at which the fraud occurred. They rolled back all transactions on the blockchain and created a new cryptocurrency at the point of the fork. Ethereum users not associated with the fraud could exchange their old tokens. After that, the rules for the blockchain changed, and new software was distributed. Users were able to either move to the new Ethereum blockchain (known as Ethereum 2.0) or stick with the old branch, which was renamed Ethereum Classic.

G. Incentives

1. The Necessity of Incentives for a Blockchain

Incentives provide something of benefit for users, but they serve an equally holistic role as well: to keep the blockchain pure. Proof of work was devised in response to spam. It benefits users, whose mailboxes are less cluttered. More precisely, it benefits "honest" users who want to use email for the purpose for which it was intended, and therefore have incentive to perform the little amount of work necessary for that purpose. It penalizes "corrupt" users by forcing them to work exorbitantly hard to abuse email.

Proof of work equally benefits the email network as a whole. Exact statistics are hard to come by: entering "how many emails are spam" in Google returns pages that report wildly different results, ranging from 45–85% in 2023. One source, Statista, reports a significant downward trend in the past decade [45]. Still, if over a quarter of email traffic is spam, that is a large drain on network resources.

An email network is not a blockchain, of course. The previous two paragraphs motivate incentives to keep corrupt users out of a blockchain. In a cryptocurrency-based blockchain, a corrupt user presents a dangerous possibility: double-spending. Double-spending is akin to counterfeiting physical currency. A perfect copy of a paper bill is accepted as real until such time as a bank notices two bills with the same serial number—an event unlikely to occur until long after the counterfeit bill has been spent. Fortunately, counterfeiting a paper bill is very difficult.

Creating a copy of a crypto coin is ridiculously easy. A crypto coin is nothing more than a string of bits. A blockchain needs a protocol that prevents a corrupt user from simultaneously presenting this string of bits to two other users.

This is where incentive plays a role. A blockchain is ultimately a single sequence and, if split into two, the longer chain is accepted as the valid one. A corrupt user cannot double-

[30] See [17] for a good overview of this attack.

spend a coin on the same chain: Each transaction records transfer of coin ownership, and it would be easy to determine that the user no longer possesses the coin. In Figure 2-10, some user U_c spends the red coin on block n+1; the transaction records the new owner U_{h1}. U_c's attempt to transfer the coin to U_{h2} in block n+2 fails because the chain records that U_c no longer owns it.

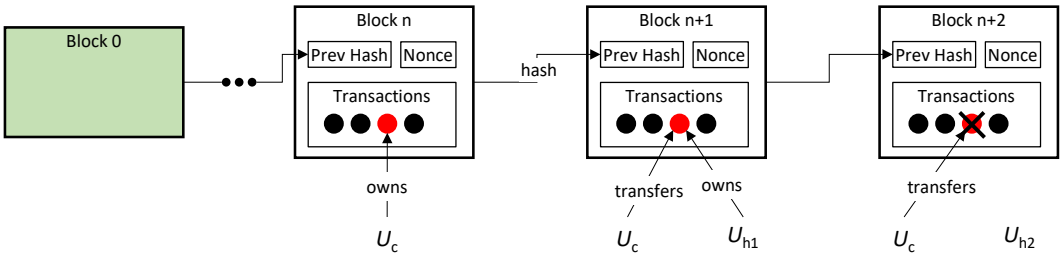

Figure 2-10. Double Spending on the Same Chain

The corrupt user might try to create a new branch of the chain, double spending the same coin on that branch. But ultimately, one branch or the other (but not both) will be accepted as part of the chain, and merging the transactions will reveal the double-spending attempt.

Incentive makes this possible. Only the corrupt user has incentive to maintain the chain with the double-spending attempt—that is, to add nodes onto that part of the blockchain. The rest of the users have incentive to add nodes to the other part of the chain, ensuring it grows quicker than the corrupt chain and is accepted. It does not matter if a large number of users are corrupt, as long as they are not colluding in their corruption. Their incentive is still to work on the longest chain, foiling any single individual's attempts. (This is not to say colluding corrupt users are not an issue: See Section 3.C.)

2. Incentives in DoD-Operated Blockchains

If incentives are driven by a desire to earn wealth, where does that leave DoD? The typical combat system is not designed to earn wealth for its users. One might appeal to the common weal: practically speaking, that implies a mandate for all DoD-operated nodes to engage in mining. But mining hardware is expensive and energy intensive. The logistics of setting up and running a crypto-mining operation in some remote post in the middle of a desert does not add to the attractiveness of using blockchain during combat operations.

Consensus based on proof of work is probably a nonstarter for combat operations, then. As Section 2.E notes, some consensus models were developed with the explicit goal of being less resource-intensive than proof of work. DoD should investigate these models for use in systems where resources are scarce. These other models often require trusting central nodes. That may be a price to pay for using blockchains.

Proof-of-stake consensus models may be a viable approach to fielded systems. If a blockchain provides value to a mission, every partner has an incentive to ensuring the continuation of that value for the duration of the mission. Nodes prepopulated with stakes at the start of a mission have incentive to perform (low-cost) coin generation.

Not all DoD systems are combat systems. In situations in which a contractor operates a system for DoD, DoD might consider monetary incentives for the contractor to perform mining. As part of the cost of running the system, the contractor would be expected to mine a certain number of coins per some predetermined time period. The contractor is incentivized to continue mining to be paid, and DoD is incentivized to keep the contractor mining to ensure the system receives whatever benefits accrue from having a blockchain (see Section 2.G for more information).

Some blockchains offer users an intangible benefit: trust. This is not to be underestimated. Nodes in networked systems devote significant resources to ensuring communications received are trusted (that the sending node is trustworthy, that the message was not corrupted in transit, that messages can be expected at regular intervals) and communications sent are transmitted. Establishing this trust is not free. If a blockchain can be the mechanism and can be implemented and operated at costs competitive with other approaches, it is logical to use it. In this sense, blockchains are one tool in a system designer's bag and should be subject to the same cost-benefit analysis as every other tool.

3. Incentives for Building Wealth

Section 2.B.2 describes the difficulty of mining. Difficulty translates to resources: CPU cycles that must be devoted to the purpose and electricity to power those cycles. One may ask, then, why a node should bother? Why not let others do the mining and use the resources for some money-making operation?

In the case of Bitcoin, the answer is that mining is a money-making operation. Each discovery of a new nonce and publication of a block earns bitcoins for the miner.[31] The miner may then use these bitcoins in transactions. The miner may also use certain institutions to exchange them for traditional currencies. When the exchange rate is high enough, it is easy to see that mining is a worthwhile and potentially profitable activity. Bitcoin is desirable and obtained through mining, so people mine it. Any community that wants to use a blockchain must provide an incentive, financial or otherwise.

[31] When Bitcoin launched in 2009, the miner received 50 bitcoins. The reward reduces by half about once every four years. As of May 25, 2020, the miner is rewarded with 6.25 bitcoins. See https://www.investopedia.com/tech/how-does-bitcoin-mining-work/#how-much-a-miner-earns for a timeline.

H. Smart Contracts

Much of the discussion in this paper assumes a transaction is a transfer of cybercurrency ownership, mimicking the real-world exchange of a physical coin. This perspective is a convenience. It simplifies the concept of a transaction at times when the nature of what is recorded in a block is secondary to some larger conceptual point.

Blockchains derive from ledgers, which record transactions, but this historical antecedent does not limit what a block can store. One interesting application of blockchains is to record "smart contracts." A smart contract is a computer program that, when executed, automatically carries out some sequence of events. An event could be a transfer of cryptocurrency, but more generally, it is an execution of the terms of a contract. Here, "contract" is meant in the most general sense of the word:

> A mutual agreement between two or more parties that something shall be done or forborne by one or both.[32]

A contract is "smart" if some technological means exists to enforce it. The original proposal for smart contracts, which dates from around 1994, offers vending machines as a canonical example.[33] We consumers know, when we approach a vending machine, that its maker has constructed it to dispense merchandise when we feed it coins and press the right buttons or pull the right knobs; we and the maker agree on this mode of operation for our mutual benefits. The vending machine is constructed such that operating it outside the contractually accepted mode (feeding it counterfeit coins, disassembling it, etc.) is sufficiently difficult and risky to disincentivize breaking the contract. (We consumers are also aware that we bear most of the contractual risk. Who has not encountered a broken vending machine? And how many persons have taken the time to call the number printed on the machine for the sake of recouping the cost of a candy bar?)

Because executable code embodies a smart contract, the contract's execution is subject to any rules and regulations implemented by this code and is also subject to penalties. Consider a smart contract governing long-term rental of a property. The contract might require:

- An initial up-front cryptocurrency payment from the lessee.
- A monthly cryptocurrency payment from the lessee.
- The return of a deposit to the lessee upon termination of the rental.

---

[32] This definition is from the *Oxford English Dictionary*. See https://www.oed.com/view/Entry/40328?rskey=4qd8qt&result=1#eid.

[33] Smart contracts date from the early 1990s. The first references to smart contracts appear in unpublished manuscripts (https://scholar.google.com/scholar?q=nick+szabo+smart+contracts). The earliest published papers date from around 1997 [18].

Once the lessee makes initial payments, the smart contract transfers the payments to the lessor and grants to the lessee a "key" that allows access to the property. Subsequently, the smart contract can bill the lessee each month and enforce penalties should the lessee fail to pay. Just as in traditional rental agreements, these penalties can be monetary (extra cryptocurrency) or restrictive (revocation of the key).

(The explanation makes analogies to legally enforceable contracts, but it is important to state that the legal status of smart contracts is still indeterminate. In 2017, Belarus became the first country to legalize smart contracts.[34] Several U.S. states, including Arizona, Nevada, and Tennessee, recognized smart contracts in 2018.[35] One student devoted a master's thesis to the validity of smart contracts under Dutch law [20]. There is no federal legislation regulating or governing smart contracts.)

Just like a regular contract, a smart contract can involve any number of individuals in whatever roles are appropriate to the contract. What it need not involve is intermediaries. Bitcoin was implemented to eliminate third parties in monetary transactions, so a smart contract can in theory eliminate intermediate lawyers, certified public accountants, witnesses, delivery messengers, and anyone else besides the parties affected by the contract. The intent is to cost less and be faster. Of course, if the law requires the participation of a third party, the smart contract is invalid without involving that party. A smart contract's block entry and hash can provide the same proof as a witness's signature, but until the law says a witness is superfluous in a blockchain-based contract, "witness" is a necessary role. Certain potential uses of smart contracts, therefore, depend on the legal system keeping up with technology.

Smart contracts are written in purpose-designed languages. As is to be expected from such new technology, these languages vary widely in syntax and features, though they do seem to possess certain commonalities:

- ***They conceptualize entities that can be involved in a contract.*** The languages recognize that a smart contract involves some thing, or things, and participants with an interest in those things (ownership, responsibility, etc.). The entities are represented using identifiers. More precisely, the entities are identifiers, and it is up to the framework in which the contract is executed to resolve those identifiers into other identifiers or physical entities (persons, property, etc.).[36]
- ***They conceptualize actions and events.*** The languages allow for the expression of steps participants need to take and events that need to occur. Rather like

[34] See https://pipiwiki.com/wiki/Decree_on_Development_of_Digital_Economy.

[35] See https://pipiwiki.com/wiki/Agoric_computing.

[36] Or virtual entities. In March 2021, Mike Winkelmann (a.k.a. Beeple) sold a purely digital artwork at auction for $69.3 million. The provenance of the artwork was recorded on a blockchain. See [33].

activity models such as IDEF0 [21], they account for action inputs, outputs, controls and constraints, and participants. Events can serve as triggers: For example, the end of one activity can trigger the start of another.

- ***They include sequencing and timing.*** A home sales contract is usually contingent on assessments by a financing organization and involves a final walkthrough that must occur within a preset time period prior to closing day. Smart contracts assume electronic transactions will require similar activities with similar types of constraints. The smart contract can also extend to the real world. A home sale using a smart contract might require a (physical) person to electronically sign that an activity has been completed.
- ***They recognize the possibility of failure.*** Just as in the real world, there is always the prospect that a buyer will not be able to pay or that the merchandise will prove to be shoddy. A smart contract will be written such that it can be terminated without the originally expected objectives having been achieved.

There are interesting DoD applications for smart contracts. Consider how they might be used during training, such as within a simulation environment. As personnel successfully master simulation components, they are granted the rights to operate the physical equipment to which their simulations correspond. A soldier might have the ability to practice with a rifle but not to fire it outside the simulation environment until they have demonstrated proficiency at hitting simulated targets. A tank driver might have the right to drive only at slow speeds until they have demonstrated proficiency at maneuvering a simulated tank across challenging terrain at high speeds. This kind of permission-based behavior could easily and automatically be enforced using a smart contract.

3. Limitations of Blockchain

Despite its potential, blockchain technology has limitations. Some of these are inherent. Others are a consequence of the practical realities of implementing blockchains.

This section provides an overview of the major limitations. Its aim is not to be comprehensive, listing each and every flaw, but to give an idea of why blockchain may be inappropriate in some situations.

A. Transaction Processing Speed

Blockchains based on proof of work are inherently inefficient. This is a design decision. Conflicts are not prevented; they are possible but are avoided by establishing a work threshold. In Bitcoin this threshold averages 10 minutes. Simple probability predicts a low likelihood that two computers will simultaneously mine coins, the square of that probability that the next two mining operations would be simultaneous and on the different blocks, and the cube of the original probability that the subsequent two mining operations would again be simultaneous and on the different chains. One source estimated there were about 1,000,000 Bitcoin miners as of February 16, 2023 [66]. It would be a gross simplification to say the probability of any two miners repeatedly simultaneously mining coins is directly related to this value—the quality and amount of hardware devoting to mining varies from miner to miner—but the large magnitude hints at the low chance. Two Bitcoin blocks have been mined simultaneously, but the Bitcoin blockchain has always accepted one of the two blocks before the next mining operation succeeded.

This protection against conflicts comes at a cost. It arbitrarily limits the number of transactions per unit time. In Bitcoin, a block is one megabyte and a transaction is about 250 bytes, giving an upper limit of around 4,000 transactions per block. In practice, including network latency, there are around 2,760 transactions per block [24] or about 4.6 transactions per second. The Visa credit card network processes 1,700 transactions per second [24].

Another way to measure speed is the average time to confirm a transaction. The graph in Figure 3-1, from data taken on April 12, 2023 [25], shows how wildly this value has fluctuated over the course of a year. Variance aside, on no day was the average time less than 6.23 minutes (on April 16, 2022). On March 29, 2023, confirmations averaged 589.98 minutes. Bitcoin is not exhibiting real-time or even near-real-time behavior.

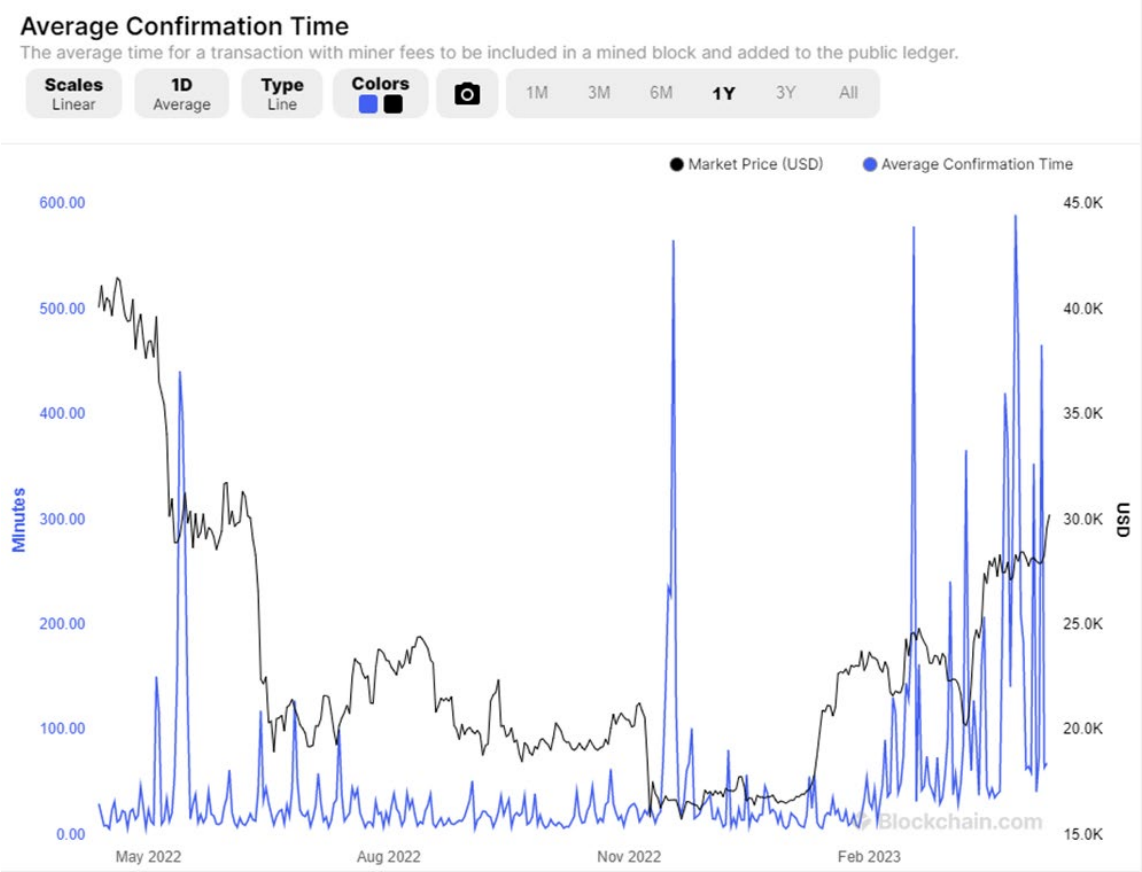

Note. "Average Confirmation Time", blockchain.com, accessed April 12, 2023, https://www.blockchain.com/charts/avg-confirmation-time.

Figure 3-1. Transaction Confirmation Time

With these numbers, Bitcoin is not ready to challenge legacy transaction-processing systems. The spike in November 2022 may have been triggered by the collapse of FTX, a cryptocurrency exchange that handled huge amounts of Bitcoin transactions [62]. Spikes in February and March 2023 may have resulted from the cascading collapses of Silvergate Capital, Silicon Valley Bank, and Signature Bank, which sent companies and individuals searching for alternative options [63]. Even discounting these peaks, Bitcoin's transaction processing speed is inadequate for commercial needs as well as for the DoD's military operations.

This especially slow transaction speed is a Bitcoin problem. No doubt its creator did not foresee its explosive growth. There have been many proposals to improve Bitcoin's efficiency; Figure 3-1 shows they either have not been adopted or have not been sufficient. The data in this section does not present an argument against blockchain per se; it only shows that the Bitcoin blockchain is too slow for potential DoD use. Other blockchain-based approaches can be considered viable:

1. Create a new, DoD-specific proof-of-work blockchain.

2. Use proof of stake or some other consensus model (see Section 2.E).
3. Employ layer 2 solutions (see Section 3.G).

B. Energy Consumption

Transaction processing time is related to the amount of energy devoted to mining. One million Bitcoin mining operators consume vast amounts of energy. Exact numbers are difficult to discern—no one knows how much mining is done by malware[37]—but it is clear that the effect is huge, and consumption keeps growing. Statistics often compare the amount of energy used by Bitcoin mining to countries: As of January 2023, Bitcoin was estimated to consume 127 terawatt hours of electricity per year, which surpasses Norway [64]. A 2023 investigation by *The New York Times* into 34 U.S.-based Bitcoin mining operations found that those companies used at least 30,000 times as much power as the average U.S. home, around 3,900 megawatts of electricity. For those mines, between 75–99% of the electricity they consume is generated by fossil fuel plants [65]. There are plans to shift mining to sites that use green energy [27], but green-energy mining does not address Bitcoin's staggering and ever-growing energy needs, nor does it address the fact that green energy will remain only part of the energy generation picture for the foreseeable future.

Energy needs for proof-of-work consensus argue against DoD using blockchains in military operations, where energy sources may be scarce and the energy grid unreliable. Bitcoin has proven itself practical in environments where the grid is stable and energy relatively cheap. Bitcoin mining has shifted toward using more renewable energy sources, and there is an opportunity for mining to use excess energy produced by renewable sources, such as wind and solar generation plants. In the long term, using Bitcoin mining to absorb some of the excess energy may help stabilize energy prices.[38]

C. 51% Attacks

A blockchain grows block by block, and if two miners simultaneously discover a new hash, a state of conflict exists. That this state should continue for any significant period of time is, as discussed in Section 2.E, extremely unlikely. The longer chain will soon be established, and transactions will be recorded on it.

This approach assumes the independence of miners. Suppose a group of miners cooperate and manage to add a sequence of nodes. Each of these miners has control over the transactions that are placed in the blocks they create. They can refuse transactions from

---

[37] Inserting mining malware on a computer is known as cryptojacking. See https://www.avg.com/en/signal/bitcoin-miner-malware.

[38] See https://www.coindesk.com/consensus-magazine/2023/07/27/enhancing-profitability-of-wind-and-solar-through-bitcoin-mining/ for more information on powering Bitcoin mining with renewable energy sources.

other miners, effectively stalling the blockchain for all other users. If one of these miners creates a block as part of a conflict, they can reverse transactions on the conflicting block and double-spend coins. The original Bitcoin paper foresaw this, and pointed out that miners do not have incentive to cooperate. It is highly unlikely that a small group of miners can work together to hijack a blockchain. Their corrupt efforts would quickly be detected (other nodes would notice their transactions are not being processed), and the miners would be ejected from the network. The profit they would make from hijacking one or two nodes would probably be small compared to the benefits they could reap from continuing to participate honestly.

The situation changes if a large number of miners decide to cooperate. If more than 50% of miners work together, the probability one of them will be the next to mine a coin becomes greater than 50%. If that happens, they control the blockchain. This is known as a 51% attack. The group members now reap all the coins mined and exert control over all transactions placed in blocks.

51% attacks are unlikely, but not unknown. Bitcoin itself suffered a 51% attack in July 2014 when Bitcoin mining pool[39] ghash.io briefly exceeded 51% of mining resources [36]. Bitcoin was still relatively novel, and the miners were not looking to subvert the technology. One miner opted to sell 50% of his own Bitcoin rather than compromise the cryptocurrency's integrity [37]. Ghash.io subsequently limited its participation to 40% [38].

Other chains have not been so fortunate. A 51% attack on the Bitcoin Gold blockchain in July 2018 resulted in a loss of $18 million [39]. The Ethereum Classic blockchain was hit by three 51% attacks in August 2020 [42].

These and other recent attacks have occurred on smaller blockchains. It is less likely that the larger blockchain networks could be attacked: recruiting and coordinating that many actors is too great a logistical challenge. Then again, in 2019, two mining pools acted together to carry out a 51% attack on the Bitcoin Cash blockchain [41]. This was not malicious: They were attempting to stop a bad actor during a network fork. However, that two mining pools can act together to control 51% of a large blockchain shows how far blockchains have grown from the original vision of small, independent miners.

D. Quantum Computing

Quantum computing is the use of quantum phenomena to perform computations. It is based on the quantum-mechanics principle that a physical system may simultaneously be in many states. In the quantum circuit model of quantum computing, a quantum computer comprises a collection of qubits, each of which can be in one of three states: 0, 1, or the

[39] A mining pool is a group of miners who agree to pool their resources.

superposition of the 0 and 1 states. The ability to be in each of these states simultaneously means a quantum computer with n qubits can simultaneously be in 3^n states. In other words, a quantum computer can explore an exponential number of possible solutions compared to a traditional computer. The implications for U.S. government networks, and especially DoD, goes well beyond the risks to blockchain alone.

A quantum computer can perform integer factorization in polynomial time.[40] Because blockchain security assumes that integer factorization is intractable—indeed, this assumption is fundamental to every blockchain implementation—the existence of a quantum computer with a sufficiently large number of qubits would compromise blockchain integrity. Different blockchains use different hashing algorithms, so no single number describes the tipping point for all blockchains, but it has been estimated that a quantum computer with 4,000 qubits would compromise Bitcoin [46].

In these days when random access memory (RAM) is measured in gigabytes and performance in gigaflops,[41] 4,000 may seem almost insignificant. Quantum computers have, however, proven notoriously difficult to design, build, and operate. Consider the following four factors:[42]

1. A quantum computer works with quantum-level phenomena (i.e., very, very small things) and requires extraordinarily precise fabrication technology.
2. Quantum computing is based on probabilities, so an operation does not give an exact answer of 0 or 1; it gives a probability.
3. Quantum computers need more error correction than classical computers. One approach is to use additional qubits. The number of additional qubits can be orders of magnitude higher than the number of qubits performing the computation.
4. Much quantum computing technology operates in supercooled environments. Although these environments are not rare, they are costly.

These factors explain why it is likely to be some time before anyone produces a quantum computer with 4,000 qubits. Figure 3-2 shows growth in qubits, starting from systems produced in 1998 up to 2017 [48]. Statistics since then seem hazy. In 2018, Rigetti announced plans to deliver a 128-qubit computer by 2019 but missed its deadline. On May 11, 2021, it was announced that Chinese researchers had created a quantum computer prototype with "the largest number of functional qubits in the world – 62" [49]. IBM

---

[40] Shor's Algorithm, developed in 1994 by Peter Schor, is the original algorithm for factoring numbers on quantum computers in polynomial time. It has complexity $O\big((\log N)^2(\log\log N)(\log\log\log N)\big)$.

[41] A gigaflop is "a unit of measure for the calculating speed of a computer equal to one billion floating-point operations per second" (https://www.merriam-webster.com/dictionary/gigaflop).

[42] These factors are a summary of material in [45].

created a 433-qubit quantum computer in 2022 and, in May 2023, announced plans to build a 100,000-qubit system by 2033 [71]. The current numbers are well below 4,000, and unlike traditional microprocessors, the growth rate in Figure 3-2 looks closer to linear than exponential. A 2019 National Academies of Science report makes the following prediction [50, p. 157]:

> Given the current state of quantum computing and recent rates of progress, it is highly unexpected that a quantum computer that can compromise RSA 2048 or comparable discrete logarithm-based public key cryptosystems will be built within the next decade.

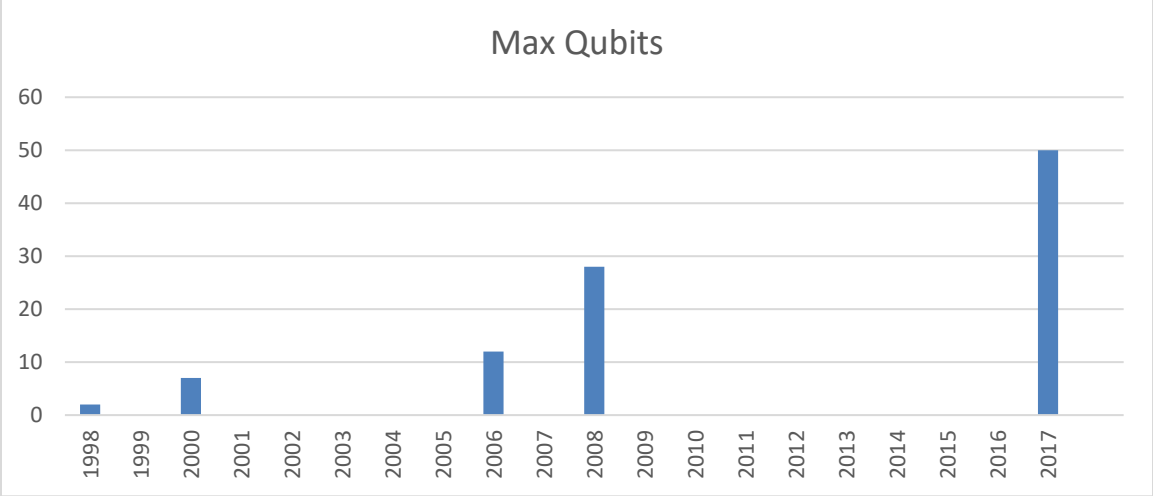

Figure 3-2. Quantum Computing Progress

In the absence of significant technological breakthroughs, quantum computing is unlikely to comprise blockchain in the near term. If DoD choses to use blockchain, it should accept the longer-term risk and plan to adopt technologies that quantum computing cannot compromise (see [46] for examples). Considering the widespread use of technologies that assume factoring is difficult, compromised blockchains may prove to be the least of DoD's problems if quantum computing ever becomes viable for this class of problems, a situation known as *quantum supremacy* [52]. DoD should be investigating how to migrate away from factoring technologies for reasons independent of blockchain.

E. Centralization

Bitcoin's creator envisioned a peer-to-peer network. Every participant possesses an identical copy of the ledger. No one's copy is more important than anyone else's, and no node has special rights. A transaction, although recorded on every ledger, involves the purchaser and the seller, and no one else. Centralized institutions, such as banks, are excluded.

It has not worked out that way. People do not exchange fiat currencies for the thrill of the exchange. They make payment for goods and for services rendered. This applies to both Bitcoin and traditional currencies. The buyer and seller must agree on the amount of currency suited to the goods or services. Considering the fluctuating value of cryptocurrencies, it is unwise to set a cryptocurrency price. By the time the transaction settles—which can take a very long time, as discussed in Section 3.A—an original cryptocurrency investment may be worth much less.

Compared to traditional currencies, cryptocurrencies are still used rarely and by a small fraction of the world's population. If someone restricts themselves to cryptocurrencies, they cannot purchase goods at the majority of brick-and-mortar stores or even online. Unless and until such time as cryptocurrencies are nearly universal, cryptocurrency users will need the ability to convert their holdings into traditional currencies and vice versa.

Persons, organizations, and governments with an interest in cryptocurrencies have addressed these issues in several ways:

- Cryptocurrency exchanges, which let users buy and sell cryptocurrencies—that is, exchange dollars, euros, or some other traditional currency for a set amount of cryptocurrency.
- The stablecoin, a cryptocurrency that is backed by some reserve asset and therefore less likely to experience price fluctuations.
- Government cryptocurrencies, which are issued by some central government that usually exerts authority over the cryptocurrency in some way.

There is overlap between government cryptocurrencies and stablecoins, although not every stablecoin is issued by a government, and government cryptocurrencies are not necessarily stable. In the U.S., one of the most popular stablecoins is Tether, which is pegged to the U.S. dollar.[43] Tether is a collaterized stablecoin, meaning that it is backed by reserves. There are a number of stablecoins pegged to the U.S. dollar, including USD Coin, a token that operates on the Ethereum blockchain, and Binance USD, which is regulated by the New York State Department of Financial Services [75].

There is also overlap between government cryptocurrencies and central bank digital currencies (CBDC), which are issued and backed by a bank. China's digital yuan, also referred to as e-CNY, is a prime example. China issued e-CNY in 2021 as a stablecoin, so it is as stable as China's reserves—in other words, very stable and likely to remain so. China's implementation is the anthesis of Bitcoin: The Chinese government is able to track

[43] This means that one Tether token is always equivalent to the value of one dollar. Token values fluctuate as the dollar does.

every transaction made with e-CNY. This raises a question as to the degree to which the e-CNY will rival other digital cryptocurrencies. As of May 2023, all government staff in Changshu, which is in the province of Jiangsu, are fully paid in e-CNY.[44] Many users outside China may be loath to accept the invasion of privacy that comes with e-CNY use. However, China is economically powerful enough to force some of its international partners to trade in e-CNY if it chooses.

CBDCs have gotten more popular since China initiated the e-CNY. As of December 2022, 11 countries have launched CBDCs, 18 countries are working toward CBDC pilot programs, and 32 have CBDCs in development, including the U.S.[45]

The March 9, 2022, Executive Order (EO) on Ensuring Responsible Development of Digital Assets identifies six priorities related to the development and regulation of digital assets: protecting consumers and investors, promoting financial stability, countering illicit finance, reinforcing U.S. leadership in the global financial system and economic competitiveness, promoting access to safe and affordable financial services, and supporting responsible innovation.[46] The EO also requests an exploration of a potential U.S. CBDC, specifically tasking the Federal Reserve to continue its research into a CBDC. The Federal Reserve has not yet made a recommendation regarding CBDC development.[47] In response to the EO, the White House Office of Science and Technology Policy published a report in September 2022 that identifies policy objectives and examines technical design choices for a CBDC system. (It does not address whether the U.S. should pursue a CBDC.) The report also examines some potential effects a CBDC may have on federal processes and systems. Though there would be benefits to adopting a CBDC, including increased ability to make payments regardless of the available infrastructure, CBDCs also present cybersecurity and privacy risks. A cyberattack on a CBDC could be used to compromise additional federal infrastructure. The need to collect and store information to verify payments—including sensitive business information and personally identifiable information—may open opportunities for identity theft or fraud [72].

The U.S. Congress is also pursuing federal regulation and oversight for stablecoins and other digital assets. The House Financial Services Subcommittee on Digital Assets, Financial Technology and Inclusion held a hearing on April 13, 2023, titled

---

[44] See https://www.cnn.com/2023/04/24/economy/china-digital-yuan-government-salary-intl-hnk/index.html.

[45] See the Atlantic Council's CBDC tracker for more information (https://www.atlanticcouncil.org/cbdctracker/).

[46] The White House, Executive Order on Ensuring Responsible Development of Digital Assets, March 9, 2022.

[47] See https://www.federalreserve.gov/central-bank-digital-currency.htm for more information, including a 2022 discussion paper examining the pros and cons of a U.S. CBDC and other publications and testimony.

"Understanding Stablecoins' Role in Payments and the Need for Legislation."[48] This hearing was in support of the development of yet unnamed draft legislation "to provide for the regulation of payment stablecoins, and for other purposes"[49] and "to provide requirements for payment stablecoin issuers, research on a digital dollar, and for other purposes."[50]

F. Pseudonymous Identities

In theory, the identities of blockchain users are anonymous. The Bitcoin paper states the following [4, p. 6]:

> … privacy can still be maintained by breaking the flow of information in another place: by keeping public keys anonymous. The public can see that someone is sending an amount to someone else, but without information linking the transaction to anyone.

The reality is more complex. Bitcoin (and other blockchain) users are not anonymous; instead, they are pseudonymous. The public key is a "name" open for all to see. Moreover, one traditionally chooses a pseudonym with an intent to deceive or mislead[51] or to suggest.[52] And in these circumstances, there is nothing to trace the pseudonym back to an individual.

Cryptocurrency is different. Unless all someone wants to do is exchange a cryptocurrency coin and never exchange that coin for some good, service, or other cryptocurrency coin, they must link their public key to another identity that is known outside of the cryptocurrency's blockchain.

Criminal networks, who early on were some of Bitcoin's most eager users, have learned this the hard way. Bitcoin is decentralized and largely outside of law enforcement's control. Cryptocurrency exchanges are not. Network traffic analysis can sometimes link blockchain transactions to outside sources.[53] In most cases, this kind of analysis only

---

[48] Press Release, "Hill Delivers Remarks at Hearing on Stablecoins' Role in Payments and the Need for Legislation, April 19, 2023, https://financialservices.house.gov/news/documentsingle.aspx?DocumentID=408714.

[49] The full text of the draft bill is available here: https://docs.house.gov/meetings/BA/BA21/20230518/115973/BILLS-118pih-Thediscussiondraftdefinesp.pdf.

[50] The full text of the draft bill is available here: https://www.congress.gov/118/meeting/house/115753/documents/BILLS-118pih-Toproviderequirementsforpaymentstablecoinissuersresearchonadigitaldollarandforotherpurposes.pdf.

[51] One author's parents once bought an 1835 letter signed by Andrew Jackson and were subsequently disappointed to learn 1830's writers often signed letters to newspapers with "Andrew Jackson."

[52] For example, Alexander Hamilton, James Madison, and John Jay signed *The Federalist Papers* "Publius" to suggest the classical wisdom underlying the proposed United States Constitution.

[53] The Chainalysis company specializes in this kind of investigation. See https://www.chainalysis.com/.

begins after investigators have a lead from some other source, such as arresting a drug offender and discovering a Bitcoin wallet on their phone or computer. And criminals, alerted to this technique, adopt additional security measures to hide their tracks (zero-knowledge proofs are one approach—see Section 2.D.1), forcing law enforcement to adapt. Nevertheless, it is unwise to think that any blockchain truly guarantees pseudonymity, let alone anonymity.[54] Law enforcement organizations, with support from the public sector, are actively involved in tracing identities [78].

Furthermore, blockchain implementations have existed for just over a decade and are not as mature as comparable distributed-data technologies, in particular database management systems. Experience suggests that, compared to mature technologies, they are more likely to have faults and are less likely to be fault-tolerant. Some research backs up this point. A 2017 paper describing the Blockbench framework for comparing private blockchain implementations concludes that the private blockchains were simply not as robust as database management systems [53]. There is no consensus on blockchain architectures, and each new blockchain introduces some variation intended to address a problem with previous ones. The immediate problem may be solved, but often at the cost of overlooking some hitherto unconsidered attack.

G. The Blockchain Trilemma

The blockchain trilemma is the idea that public blockchains must compromise either security, scalability, or decentralization to work. Bitcoin is a great example. Proof of work, Bitcoin's consensus mechanism, enables decentralization and security but hinders scalability, which refers to the number of transactions a network can process per second (an average of 4.6 transactions per second). There are multiple ways to address the trilemma, though none are perfect solutions. Switching to proof of stake removes some of the scalability issues inherent in proof of work, particularly the time and resources needed to mine.

Sharding, a method to boost scalability for layer 1 blockchains,[55] divides the blockchain network into smaller pieces, known as shards. Rather than rely on the entire network to process and validate transactions, each shard can process transactions in parallel, boosting overall capacity. Each shard may have its own smart contracts as well. Sharding is growing in popularity—the Zilliqa[56] and Cardano[57] public blockchain

[54] See [32] for more discussion.

[55] Layer 1 refers to a base network and its underlying infrastructure; the main network within a particular ecosystem. Layer-1 blockchains, such as Bitcoin and Ethereum, process transactions on their own networks. See https://academy.binance.com/en/articles/what-is-layer-1-in-blockchain for more information.

[56] See https://www.zilliqa.com/ for more information.

[57] See https://cardano.org/discover-cardano/ for more information.

platforms use sharing, among others—but it does present some challenges. Compared to the entire blockchain network, shards are more susceptible to hackers, who could overtake a shard and then disseminate malicious transactions across the entire blockchain. It can also be difficult to implement network sharding, since it requires splitting the network and reassigning node states.

Another option is to employ layer 2 solutions to ease the burden on the main chain. Layer 2 solutions, as the name suggests, add a second layer on top of the main blockchain network (also known as the parent or main chain). Layer 2 solutions generally comprise two parts: a network to process transactions and a smart contract that links the protocol to the main chain.[58] Transactions can then offload from the main chain to the layer 2 protocol, which lessens the main chain's computation burden and facilitates faster processing.[59] Some layer 2 protocols can process thousands of transactions per second; Lightning Network, a layer 2 solution for Bitcoin, can process over 1 million transactions per second.[60]

One type of layer 2 option is a side chain, which is a secondary blockchain built on top of the main chain. A side chain relies on the main chain for security and dispute resolution processes, but it processes transactions separately. Side chains can use different consensus mechanisms from the main chain: If the main chain uses proof of work, for example, a side chain can use proof of stake or another method.

A plasma chain is a version of a side chain that is anchored to the Ethereum Mainnet, the main, public Ethereum blockchain.[61] Like side chains, plasma chains have their own consensus mechanisms to validate transactions and add blocks. But unlike side chains, the root of each transaction on a plasma chain block is published to Ethereum. This lends extra security to plasma chains: Plasma chains inherit security from the main chain's consensus mechanism, and plasma chain users can point to block roots on the main chain as proof of transaction if something disrupts or corrupts the plasma chain.[62]

---

[58] See https://chain.link/education-hub/what-is-layer-2 for more information.

[59] See https://crypto.com/university/blockchain-scalability.

[60] https://timesofindia.indiatimes.com/business/india-business/exploring-the-top-5-layer-2-crypto-projects-of-2023/articleshow/102840191.cms?from=mdr

[61] See https://ethereum.org/en/enterprise/ for more information about plasma chains.

[62] https://docs.plasma.group/en/latest/src/plasma/sidechains.html#:~:text=Simply%20stated%2C%20a%20plasma%20chain,!)%20and%20still%20be%20safe.

4. DoD Applications for Blockchain

This section presents some examples of how DoD might use blockchains. The section covers applications for both day-to-day operations and missions. The emphasis is explicitly on blockchains, not cryptocurrencies.

A. Materiel Ledgers

The first example of how DoD could use a blockchain is the direct application of the technology a blockchain models: ledgers. DoD is an enormous organization with a huge materiel inventory. Would recording transactions on a permissionless blockchain improve DoD's ability to manage materiel movement?

It does not seem likely. Permissionless blockchains promote recording transactions between autonomous peer nodes. DoD is a hierarchy, not a collection of peers. Actions requiring materiel transfer are initiated from the top. Each organization has a budget approved by a higher-level authority (ultimately the U.S. Congress). General needs are known in advance, and transactions can be predicted. It might be useful to have a non-repudiable blockchain record of transactions as they are carried out, but only if DoD had a need to block fraudulent transactions in materiel transfer, which does not seem to be the case. There are instances of problems accounting for materiel,[63] but they are thefts, not frauds. Blockchains do not address theft of materiel.

Furthermore, each service has devised extensive and complex processes, practices, and policies to support materiel transfer and has considerable organizational support devoted to implementing them. The Army, for example, centralizes control and dispersal of materiel at Army Materiel Command, an organization employing upwards of 190,000 military, civilian, and contractor personnel.[64] For the past two decades, the Army has been migrating its logistics support systems to Army Global Combat Support System (GCSS), which centralizes logistics information technology (IT) operations [58]. The Army-GCSS supports and implements the many Army logistics processes. The decision to tinker with that implementation requires detailed cost-benefit trade-offs that are outside the scope of this paper.

---

[63] For example, https://ktla.com/news/nationworld/top-general-shocked-by-ap-report-on-missing-military-guns-mulls-systematic-fix/.

[64] See https://www.amc.army.mil/ for more about the Army Materiel Command.

B. Supply Chains

It is no secret that DoD faces enormous challenges establishing and running supply chains. Long before the DoD, the maintenance and operation of fighting forces has been plagued by quality and timing issues, as well as shortages. To these old problems, we can add the modern problem of security. Adversaries, given access to supply chains, could always attempt to slip in shoddy equipment in hopes of causing premature failures. Now they can substitute equipment that deliberately fails on command (an example is the Stuxnet worm's devastating attack on Iran's nuclear production capability) or, unknown to the user, transmits information to the adversary.

The typical product DoD acquires is a complex system composed of several interacting parts, many of which are complex subsystems in their own right. The bill of materials for a system can easily include the company that manufactures each part. Those parts that are subsystems, however, perhaps contain parts manufactured by other companies, which in turn may assemble parts made by yet more companies. Keeping track of all manufacturers in a supply chain, with an eye to vetting each and every one, can become an enormous, costly task.

The problem has become especially acute in software supply chains. Today's applications run on top of layer upon layer of software, in whose development the application developer has neither participation nor visibility. Applications also rely on millions of lines of code in external libraries. Except in very unusual circumstances, the application developer has no choice but to trust that all these externally developed components and systems have no malicious functionality.

Blockchains offer an approach to reducing supply chain risk management. They have already been employed in diverse supply chains:

- Pacific Northwest National Laboratory has begun two pilot programs that use blockchains, one of which emphasizes how blockchain can improve asset management and supply chain security [54].
- The IEEE has formed a working group to investigate blockchains for food security and propose standards for their use. When an outbreak of a foodborne illness such as *e. coli* occurs, it is vital that investigators be able to trace food as far back as necessary to the source of contamination, even to the farm that produced it. In this age of international food shipments, different countries' laws make quick tracing especially problematic. Blockchain offers a fairly straightforward solution [55].
- Deloitte published a report concentrating on the financial aspects of using a blockchain in a supply chain [56]. Aside from its emphasis on pricing, the report makes the same points as others: blockchain can improve supply chain quality,

help streamline and automate administrative processes, and reduce time and effort necessary to assess supply chain risks.

- U.S. Customs and Border Protection (CBP) has requested information on commercial blockchain/distributed ledger technologies for processing import and export trade data. CBP wants to explore using blockchain to improve data sharing, increase supply chain transparency, and make it easier to find and flag suspicious transactions [77].

The advantage of using a blockchain ledger for a supply chain is straightforward. Suppose a block in the chain represents a manufactured part, including the manufacturer, date of acquisition, and the parts it comprises. There must exist a previous block in the chain for each constituent part, all the way back to elementary components. For hardware, this might stretch back to the operation that produced the raw materials. For software, the stopping point would be the repository containing the source code. If that source code requires external libraries, each of those libraries must exist in a previous block in the chain, their provenance recorded and verifiable.

With this blockchain in place, determining the full provenance of a complex system is straightforward. One traces back through the blockchain, finding all manufacturers and looking for any associated with adversaries or known to have been compromised.

To ensure the supply chain is not tampered with, the block can include a hash of the code (for software) or a hash of all part identifiers (for hardware). That would facilitate reasonably quick review and verification of all components. For added security, a block for software could include a hash of the tools used to compile and deploy a given component. This would help mitigate the kind of attack used on SolarWinds. An attacker's modification of a compiler would be detectable.

Putting this kind of supply chain in place is no trivial undertaking. It requires participation from every manufacturer possibly involved in a supply chain. Conversely, it requires every manufacturer to use only products whose manufacturer participates in the blockchain. For hardware, one can imagine DoD creating incentives to participate. Companies manufacture hardware with an eye to selling it. If DoD regulations prohibit buying parts not recorded on a blockchain-based supply chain ledger, and likewise prohibits buying parts composed of parts not recorded on that same ledger, companies have a financial motivation to record their manufacturing activities on a blockchain.

Could this also be feasible for a software-based product? A large percentage of software is developed and distributed through open-source channels. Much of the rest is proprietary, the source code closely held by its developer. A full supply chain requires vetting the underlying operating system. Imagine the effort required to review the Linux kernel and its standard tool suite or the difficulty of vetting an operating system released by Microsoft. Some amount of trust will probably be necessary. DoD would have to accept

and allow selected Linux or Microsoft software releases. As for tools and libraries that run on top of those systems, restricting components to software factories, such as DoD's Platform One,[65] could increase confidence that only trusted components were being used to build applications.

C. C2 and C3 Systems

DoD units use command and control (C2) systems and command, control, and communications (C3) systems for planning, executing, and reviewing missions. Systems that implement the NATO-sponsored Joint Consultation, Command and Control Information Exchange Data Model (JC3IEDM) are an example.[66] These systems are used to record, view, and query information about entities and actions of interest to military planners. All systems access a database whose schema is the same data model (the JC3IEDM), and the database is currently implemented using a database management system. Each of the NATO countries that participate in the Multilateral Interoperability Program (MIP), which is responsible for publishing and maintaining the JC3IEDM, implement their own version of the system. In coalition operations, each system has its own database management system and database for reasons of resilience and national security. Select information is shared between databases using database replication features of the database management system.

C2 and C3 systems are typically write-once. Information, once entered in a database, cannot be deleted or modified. If someone estimates an enemy force has 1,000 combatants and later changes that estimate to 500, the database preserves the initial estimate. The intent is to preserve all reporting, ensuring that after-action reviews have access to all information gathered, accurate or not. Inaccurate information may prove helpful in identifying and understanding failures.

Modern database management systems support this level of access protection. Table permissions can be set to prohibit delete and update operations for the average user. However, modern database management systems also have, and indeed require, the concept of a user with unlimited powers. This person, the database administrator, can establish special permissions for themselves. An adversary who gains administrative rights can subvert the normal permissions and make arbitrary modifications to database contents. In the JC3IEDM, the adversary could modify records without harming database consistency; that is, it would be difficult to detect that the constraints on modification had been violated. Log files can be examined after the fact, although logging must be properly

---

[65] Platform One provides products and services for creating container-based infrastructure. See https://software.af.mil/dsop/services/.

[66] The JC3IEDM is a NATO Standard Agreement (STANAG) for information exchange in coalition operations. See https://standards.globalspec.com/std/1031445/STANAG%205525.

configured,[67] and log files must be regularly monitored for unexpected activity. Relying on logging is a reactive strategy.

Blockchain, by contrast, offers a proactive strategy. Imagine an implementation in which every addition to a C3 database is recorded on a blockchain. This is equivalent to cryptocurrency, in which every entry represents an addition to the "database" of cryptocurrency transactions. A new hash value is computed from the previous hash value for the database combined with the new record. Any attempt to delete or modify an existing record will yield an improper hash value. The software enforces the write-once policy at the time records are written, catching illegal modifications immediately.

D. Real-Time Distributed Systems

Suppose a force is poised to attack a target. The force is split into multiple units, each positioned at a different location with respect to the target. The force's overall commander believes an attack will succeed if all units attack simultaneously. However, should one or more units not participate in the attack, the defenders will be able to marshal resources, concentrating their own forces against individual units before shifting to oppose other units. The more units that do not attack, the more likely the entire plan will fail.

This is known as the Byzantine Generals Problem [57]. In the classic formulation, the unit commanders (Byzantine generals) have surrounded a city, and their commander needs to send them the attack time. However, the commander believes some of the generals may be traitors. Furthermore, the generals communicate by sending messages via their lieutenants, some of whom may be traitors, and others of whom may be captured en route. Assuming there are not too many traitorous generals and not too many traitorous lieutenants, how can the loyal generals be confident they know the correct time to attack?

If one thinks of the generals as systems, and of the lieutenants as communication channels, then the problem encapsulates ensuring the resiliency of a real-time distributed system:

1. How does a working system determine whether other systems in the network are working correctly?
2. How does a working system determine that the content of the messages it sends and receives are transmitted properly?

Researchers soon realized that a blockchain can be used to solve the Byzantine Generals Problem. Traitorous and captured lieutenants correspond to network faults, which can be categorized, respectively, as message corruption and network link failure. A

[67] For example, the MySQL DBMS does not enable logging by default. See https://dev.mysql.com/doc/refman/8.0/en/server-logs.html.

blockchain addresses message corruption through cryptographic security: a block's hash verifies that it correctly encodes the information placed in it. A distributed ledger addresses network link failure: the absence of an expected transaction demonstrates its incomplete transmission throughout a network.

Traitorous generals are, roughly speaking, dealt with by participation incentives as discussed in Section 2.G. A traitorous general is analogous to a blockchain participant trying to form a new chain off the main chain. As Section 3.C discusses, this is unlikely to succeed unless the participant controls more than 50% of mining operations. The scenario can tolerate a single traitorous general (faulty system) and in fact may be able to tolerate multiple traitorous generals. In any real-life operation, the planner needs to know how systems can fail before the entire operation is compromised.

This may seem imprecise—it is. Fault tolerance is a complex, probabilistic field, and blockchains are not a magic solution. A blockchain cannot guarantee a distributed system will continue operating in the presence of node or network failures. It cannot even guarantee that failures can be detected. A blockchain can, however, guarantee that the probability of failures occurring without detection is exceedingly low. If each node's ledger does not record all expected communications in a period consonant with the expected time to propagate messages across the network, each working node can conclude the probability of a fault is high and act accordingly.

This confidence comes at a cost. In traditional distributed C2 networks, commands can be broadcast as soon as they are composed, and responses are transmitted as quickly as they can be formulated. In a blockchain, transactions cannot be placed in a block until a nonce is generated. Depending on the blockchain implementation, this may take several minutes. This kind of distributed system is not real-time. It may be suited to operational, strategic, and tactical planning, though its suitability for operational use is questionable. In combat situations, soldiers need to obey orders instantly and cannot wait until another block is generated. The risk of having to delay operations while awaiting a nonce is too high. Even for networks of automated systems, the use of blockchains must be carefully considered based on the possibility of delay. Recent research on drone swarms illustrates this point. DARPA's Offensive Swarm-Enabled Tactics (OFFSET) program is studying swarms of 250 or more drones.[68] The challenge of coordinating that many drones (for example, having each maintain a complex flight pattern to discourage targeting while simultaneously avoiding collisions) requires a large amount of network communications. On the one hand, a blockchain is an excellent technology for maintaining a fault-tolerant network and tracking which drones are operational. On the other hand, drones are

[68] See https://www.darpa.mil/work-with-us/offensive-swarm-enabled-tactics.

constantly moving and need to communicate their new positions very frequently. They cannot await a miner.

In operational settings, there is also the issue of which nodes can serve as miners. Mining is computationally and energy intensive, so it may not be realistic to assume that mining can occur in theater, where resources may be scarce. Mining is, of course, not the only way to generate new blocks, but 51% attacks are a threat if block generation is too easy.

These considerations suggest that blockchain is better suited to planning military operations than carrying them out. The Byzantine Generals Problem is a planning problem, which is one reason blockchains can be used to implement solutions. Blockchain implementations operate in near-real time, not real time.

E. Automated Workflows

DoD may be able to use blockchain-based systems and smart contracts to automate certain workflows. Section 2.H enumerates an example: Using smart contracts to track training certifications and grant permissions to use tools and systems once the required certification(s) have been achieved. Blockchains could also streamline existing systems to track and record approvals for documents that need to be coordinated among multiple offices and authorities.

Smart contracts facilitate a "first A, then B" system that benefits workflows that require specific criteria to be met before proceeding to the next phase. One example is workflows supporting acquisition and procurement. Generally, acquisition processes are rife with reviews and decision points that only occur once certain criteria have been met. Smart contracts could facilitate certain workflows by triggering a notice for review once the required criteria have been achieved and recorded. Doing so may require specific, considered applications. For example, in a program developing a cyber capability, an initial requirements review may require multiple documents listing requirements related to how the capability should function when used operationally, characteristics of the intended operational environment that need to be considered during development, and any constraints or limitations that may affect the initial capability design, among others. These documents may come from different stakeholders. Adding the documents to the blockchain would serve as the impetus for the smart contract to move the workflow to the requirements review stage.

Another example is workflows supporting contracting processes. A smart contract may generate a payment for services rendered once proof of completion is received and verified. For phased contracts, submitting a deliverable may start the funding process for the next phase, to the extent of a sponsor automatically transferring money to a contractor. These applications provide additional functionality not found in conventional databases.

They may also reduce the burden on contracting officers by automating some of the steps involved with closing out a contract or initiating a funding action.

Blockchain-based systems can also support document coordination and approvals among multiple offices and authorities. The blockchain would keep a clear record of who reviewed and approved the document and ease the burden of managing version control—almost always a challenge, especially when multiple parties review and make changes.

These examples assume that the activities needed to verify the requirements of a smart contract occur on chain or within the boundaries of a particular blockchain, but this is not always the case. Smart contracts cannot operate outside of a blockchain, but they often require data that resides outside the blockchain boundary (also known as off-chain data). Using smart contracts to track training and grant permission to use certain tools or systems is one example: The smart contract needs to be able to contact the network controlling a particular tool. This can occur through the use of oracles. Blockchain oracles serve as bridges between a particular blockchain and the outside world, as well as between different blockchains. Oracles can collect and aggregate data for use within a blockchain and allow smart contracts to send commands to off-chain systems to trigger specific actions (e.g., unlock a system, make a payment, or store data).

Decentralized oracles echo blockchain decentralization and allow oracles to pull data from multiple sources while maintaining the core premise of a blockchain. The DoD could use distributed oracles to support blockchains gathering verifiable, open-source data to support financial processes (e.g., contracting), open-source intelligence gathering, or other missions.

F. Electronic Health Records

Ideally, electronic health records (EHR) provide secure storage for patient data that is easily accessible by both the patient and the healthcare provider, no matter the provider's location or organization. Actual implementation has yet to reach this goal. Though EHRs have been effectively implemented for individual practices or healthcare groups, when it comes to interoperability, they are more akin to digital versions of paper-based medical records than anything else. Sharing information with the patient or other providers, particularly providers outside a particular practice or group, often requires printing documents to then be faxed or hand carried. EHRs are designed to comply with HIPAA rules for privacy and security,[69] though they are also prime targets for exploitation, as

[69] HIPAA stands for the Health Insurance Portability and Accountability Act of 1996 (Pub. L. No. 104-191, §264, 110 Stat. 1936). HIPAA's privacy rule protects individuals' medical records and protected health information (PHI), which includes information about an individual's past, present, and future physical or mental health; provision of health care; past, present, or future payment for provision of health care; and other identifying information (e.g., name, birthdate, Social Security Number) (see https://www.hhs.gov/hipaa/for-professionals/privacy/laws-regulations/index.html for more information).

medical records contain the information needed for identity theft and command high prices on the dark web.

The DoD has seen success with its EHR modernization program, but interoperability outside the Department is still an issue. In April 2023, the Department of Veterans Affairs (VA) paused its years-long EHR modernization effort to address issues and complaints from veterans and VA clinicians [67]. The VA's effort is meant to replace its current EHR with the same system DoD implemented to improve overall capability, effectiveness, and interoperability, particularly with DoD. But moving patient records from DoD to the VA when military personnel retire is difficult, and a January 2023 Government Accountability Office (GAO) report identified issues with the initial implementation [68].

Blockchain-based EHRs can support interoperability among different providers, practices, and healthcare groups. Rather than a centralized EHR, a blockchain-based EHR provides a decentralized system that would allow access to data across multiple independent systems. This eases the administrative burden of sharing records among different providers and makes it easier for patients to access their information. A hybrid permission model would allow healthcare providers to add blocks and enable patients to see and access their data when needed.

A challenge for this type of system is the amount of data that would need to be stored. EHRs contain large amounts of data, including images, and storing large amounts of data on a blockchain may make the system slower and less efficient. One potential solution is to use a hybrid data storage model, which stores metadata (e.g., access controls, a record of transactions, etc.) on the blockchain and uses a HIPAA-compliant cloud to store the actual EHR data [69].

Blockchain EHRs should support HIPAA compliance, and several current companies claim that they do. Standard features—cryptographic hashes, public and private keys assigned to users—maintain data privacy and immutability and help protect against tampering and theft. However, blockchain-based EHRs are a nascent category, and regulation has not yet caught up to the technology.

G. Monitoring Media Use

Though misinformation and disinformation have been around for almost as long as humans have, the threat of disinformation to national security and social cohesion came into prominence with the revelation of Russia's sophisticated disinformation campaign to influence the results of the 2016 and 2020 U.S. presidential elections. Effective disinformation campaigns such as Russia's exploit existing divisions and hot-button issues, making use of existing arguments and media to bolster a desired narrative. It is far easier to manipulate existing media than create new media.

One emerging threat is deepfakes, which are modified images, audio, or video that make it seem like someone did or said something they never did or said. A notable example is from March 2022, when a manipulated video showing Ukrainian President Volodymyr Zelensky ordering Ukrainian soldiers to surrender to Russian troops circulated on social media and appeared on a Ukrainian television station's website and live feed. Though this particular video was not particularly convincing—the consensus was that Zelensky's accent was incorrect and his voice did not sound accurate—it illustrated the potential threat deepfakes could pose to national security [73]. In May 2023, a deepfake image showing an explosion at the Pentagon went viral on Twitter. The Pentagon Force Protection Agency and the Arlington County Fire Department issued statements on Twitter stating that there was no explosion or incident at the Pentagon, but the hoax still caused a brief dip in major stock market indices [74].

Blockchains record and track data, creating a system that ensures data is preserved and remains unchanged. A blockchain's distributed ledger serves as a permanent, immutable archive that can be searched and used to prove data authenticity and integrity. A blockchain system tracking images, audio and video could help verify media provenance and make it easier to determine when media has been manipulated. This would give DoD greater control over the media it creates and support the dissemination and maintenance of narratives that support U.S. interests. For example, if an adversary manipulates an image that DoD created and placed on a blockchain, DoD could present the original and altered images side by side to highlight the changes and reinforce the meaning and context of the original. This would build trust in the DoD as the source of the original media and help expose attempted disinformation campaigns. Generally, it is very difficult to regain control of the narrative once it has been hijacked by a disinformation campaign. Disinformation is designed to maximize engagement by eliciting an emotional response, and it tends to spread faster and farther than true information. (That facts are often less sensationalized than fiction does not help.) Instead of fact checking, providing verifiable proof of a disinformation campaign can help debunk attempted disinformation and offer a compelling alternative narrative.

Media copyright sites such as Pixsy[70] use blockchain to monitor image use online. Pixsy is targeted toward artists and photographers as a tool to protect their copyright by identifying when and how images are posted online. Pixsy also provides legal services to support copyright infringement cases. Though DoD may not be interested in Pixsy's legal support, this type of monitoring application would help DoD identify when and how its media are posted and used, providing insight into both DoD's scope and reach online and how its media may be adopted and/or manipulated to support misleading narratives.

[70] See https://www.pixsy.com/.

A tool like Pixsy would serve DoD best as part of a suite of solutions aimed at countering disinformation and promoting accurate narratives. People tend to dig their heels in when presented with information that contradicts their views, even if there is irrefutable evidence proving that information to be true. It may be most effective to employ both proactive and reactive approaches to countering disinformation. Using a blockchain to provide proof of media authenticity is a reactive approach. A proactive approach would be teaching or disseminating core tenets of media literacy, which helps people think more critically about the information they consume. In fact, targeted media literacy campaigns have the ability to temporarily inoculate people against the persuasive powers of a disinformation campaign.[71] Another proactive approach is to use existing platforms, such as official statements, press releases, social media, or others, to establish factual and compelling narratives that can then be supported using reactive tools when needed. Building trust—both in a narrative and in the source of that narrative—takes time and requires a multi-pronged approach.

H. Blockchain for Intel Gathering, Operations Support, and Humanitarian Work

Though blockchains seem better suited to support planning activities than full military operations (see Section 2.G.2 for more information), there are other uses that can support operations, intelligence gathering, and humanitarian work.

One of the most straightforward uses is to transfer funds. The DoD can use blockchains to securely transfer funds to military units in theater. Back-office and management functions are often the first thing lost during operations in theater, especially if there are scant resources available to support financial record keeping or funding disbursement. Rather than dedicate needed resources toward record keeping, using a blockchain to provide and disseminate funds removes the immediate need for traditional financial management processes and oversight. Blockchains automatically provide a secure, practically immutable record of transactions and support remote auditing. They also provide safeguards against the theft.[72]

The DoD can also use blockchains to transfer funds externally. The DoD has a robust humanitarian mission. Non-governmental organizations, non-profits, and other aid groups already use cryptocurrencies—usually stablecoins, as they are less subject to market

---

[71] Inoculation theory posits that it is possible to confer resistance against malicious persuasion attempts before they happen, much like a vaccine confers protection against a virus. Exposure to the online game *Bad News*, in which players assume the role of a fake news creator and learn about common misinformation techniques, was found to be effective at reducing susceptibility to online misinformation. See [76] for more information.

[72] Pero, M.C., "Understanding Bitcoin and Its Utility for Special Operations Forces," U.S. Naval Postgraduate School, March 1, 2022.

swings—to send financial aid to governments, humanitarian organizations, and individuals in need. Using a cryptocurrency lets organizations send funds securely and often out of sight of authoritarian regimes. Sometimes, the individuals in need are using cryptocurrencies themselves: Women in Afghanistan, currently subject to Taliban rules barring them from opening and maintaining their own bank accounts, have turned to Bitcoin as a source of income. This allows Afghani women to maintain their own digital wallets, which offers them some financial independence and ensures they have funds to support their families.[73]

The same methods could support special operations and intelligence work. In addition to sending funds for humanitarian purposes, the DoD could use blockchains to anonymously fund resistance groups in certain countries. Generally, providing funding behind enemy lines means resorting to physical cash transfers, which are subject to theft, or bank transactions, which can be traceable. Blockchain removes the cash component, provides anonymity, and adds flexibility, as cryptocurrencies can be converted into local hard currency when needed. The DoD could also use blockchain to fund human intelligence (HUMINT) operations and sources, reducing the need for face-to-face interactions and helping maintain the source privacy and security. Anonymous financial transactions on a blockchain may also provide a means of hiding messages in plain sight, another boon for intelligence activities.[74]

The DoD can also use blockchain to support counterintelligence and counter-terrorism activities. Adversaries and terrorist organizations use cryptocurrencies to raise, transfer, and launder funds to evade sanctions and pay for illicit activities or services (e.g., terrorist activities, assassination plots, drug dealers, or human traffickers).[75] The DoD should continue to develop its understanding of terrorism financing to track illicit activities for intelligence purposes and to develop methods for blocking or disrupting these funding streams.

---

[73] https://nymag.com/intelligencer/2022/09/afghanistans-crypto-lifeline.html

[74] Pero, M.C., "Understanding Bitcoin and Its Utility for Special Operations Forces," U.S. Naval Postgraduate School, March 1, 2022.

[75] Wagman, S., "Cryptocurrencies and National Security: The Case of Money Laundering and Terrorism Financing," *Harvard National Security Journal*, Vol. 14:87, 2022, https://www.hks.harvard.edu/centers/mrcbg/programs/growthpolicy/cryptocurrencies-and-national-security-case-money-laundering.

5. Conclusion

Although blockchain technology has matured, it is still relatively new in terms of application. Blockchain is most commonly used for cryptocurrencies at present, but there is plenty of potential for other applications. The challenge for DoD is to determine the most effective use cases.

Some type of permissioned blockchain seems the likely choice for DoD applications, though there may be circumstances that call for permissionless blockchains. Using a blockchain-based application with partner nations is one example, though that raises the question of administration. One solution may be using a hybrid or consortium permission model. In a consortium model, access is controlled by a group, which could include representatives from DoD and each partner nation involved, though day-to-day administration would likely be delegated to one member of the group. For example, a model like this could apply to an information-sharing agreement between DoD and NATO. DoD could work with NATO members to establish the permission model, and NATO would be responsible for day-to-day administration.

Permissioned blockchains would give DoD necessary control over access to the network for internal blockchain applications. Unlike permissionless blockchains, which promote trust in the system rather than in the users, permissioned blockchains reintroduce the need for an established trust relationship. This makes them a good fit for DoD, as DoD systems already require that the administrative organization trusts the users and vice versa. Also, there is an inherent lack of expectation for privacy when accessing a DoD system. Simply gaining access to a DoD facility requires one's identity to be known and vetted.

As discussed earlier in the paper, blockchains using a proof-of-work consensus model are probably not a good tool for supporting military operations, given that mining is resource and energy intensive. Proof-of-work blockchains, such as Bitcoin, tend to be most practical in environments with stable and reliable access to energy and in contexts that do not rely on real-time decision making. However, this does not preclude the use of proof-of-work models as a whole. Blockchains may be better suited for operational planning than actual operations, as they operate in near-real time.

Using a different consensus model, such as proof-of-stake, would lessen the burden on DoD in terms of energy and resource consumption. However, these models would still be best applied in situations that do not require real-time information.

There are many potential blockchain implementations for DoD. Implementing blockchain-based systems would facilitate workflows and help ensure the integrity of DoD-created information. Doing so would also require buy-in from senior leaders. It would be no small effort to provide DoD organizations and programs with blockchain systems. DoD would have to put sufficient infrastructure in place, including, in some cases, infrastructure to support multiple instances of a system on unclassified and classified networks. Integrating blockchains onto a classified system would likely require specific security configurations and approvals. DoD would also have to provide comprehensive user training. A phased approach to adoption that initially targets a few applications would focus efforts on developing the supporting infrastructure and give DoD the opportunity to identify issues and adjust the implementation to best suit the need.

This initial outlay would reap benefits in efficiency and security. Blockchain-based systems can support automated data collection and application and achieve the CIA triad, which is essential for secure, effective systems. As blockchain capabilities continue to grow beyond cryptocurrencies, there will be many additional opportunities for DoD application.

Appendix A. Open-Source Blockchains

The list below provides a sample of available, open-source blockchain platforms and tools.

Ethereum

Ethereum was the first organization to provide smart contracts on its blockchain. Ethereum currently offers:

- Ether, a cryptocurrency used on Ethereum applications.
- The Ethereum Virtual Machine, the platform on which all Ethereum-based smart contracts execute.
- Solidity, a language influenced by C++ and JavaScript, for writing smart contracts.
- Vyper, another smart-contract language that is influenced by Python.
- Yul, a language close to the instruction set of the Ethereum Virtual Machine, and Yul+ (a Yul extension).

The instruction set of the Ethereum Virtual Machine is Turing equivalent, as are the three languages listed above. The intention is to allow a smart contract to be anything that can be expressed in code. See https://ethereum.org/en/developers/docs/ for more information about Ethereum.

Ripple

Ripple is a currency exchange geared toward businesses, particularly financial institutions. Ripple provides RippleNet, a software infrastructure and payments network that supports sending and exchanging XRP, a "carbon-neutral"[76] cryptocurrency that runs on a public, decentralized blockchain. XRP uses a consensus protocol that relies on validator servers to agree on transaction order and outcome, a faster and less-resource-intensive model than proof of work. See https://ripple.com for more information.

[76] See https://ripple.com/xrp/.

Stellar

Stellar is another currency exchange that uses tokens to stand for actual currency or other items considered valuable, such as gold or units of time. Stellar also has lumen, its native cryptocurrency.

Stellar uses a proof-of-agreement consensus protocol that disseminates messages and voting processes among users in a particular group to confirm transactions. Anyone can set up a node, but users must submit identifying information for public record to join a group and participate in transactions. This is meant to help determine trust among users and groups of users. Stellar also imposes a small fee with every transaction to deter malicious behavior. See https://stellar.org/ for more information.

R3 Corda

Corda is a blockchain in that transactions are cryptographically linked, but whereas "traditional" blockchains batch transactions into blocks, Corda confirms each transaction separately and claims to do so in real time. Corda offers a private and permissioned distributed ledger technology platform. Only users involved in a transaction share data; even the communications protocol is kept hidden from users not involved with a particular transaction. Smart contracts on Corda can be written in any Java virtual machine (JVM)compatible language. See https://r3.com/products/corda/ for more information.

Hyperledger Foundation

Hyperledger Foundation (see https://www.hyperledger.org/) is an open-source community hosted by the Linux Foundation that focuses on developing enterprise blockchain applications. It has multiple applications in development and available for use. One is Hyperledger Sawtooth (see https://sawtooth.hyperledger.org/), for which the Intel Corporation developed the proof-of-elapsed-time consensus protocol. Sawtooth transaction processes can be written in Rust, Python, Go, or JavaScript.

Another application is Hyperledger Fabric, a modular application that offers a flexible approach to data privacy. Data can be isolated in channels or stored in private collections that can then be shared as needed. There is similar flexibility regarding smart contract models. See https://www.hyperledger.org/use/fabric for more information.

Solana

Solana is a public blockchain that supports non-fungible tokens (NFT) and decentralized finance, payment and gaming applications, among others. Solana considers itself "censorship resistant," meaning that it is very difficult for unintentional causes (e.g., computer or infrastructure failure) or intentional causes (e.g., malicious actors) to prevent users from adding blocks.

The Solana blockchain uses both proof of stake and proof of history to validate blocks. Blockchains use timestamps to validate transactions in the order in which they were received. Proof of history builds timestamps into the blockchain by using a verifiable delay function (VDF). Nodes need to execute the VDF to add a block. According to Solana, proof of history enables nodes to verify parts of the blockchain in parallel. See https://solana.com/ for more information.

Everledger

Everledger is a private, permissioned blockchain focused on increasing transparency and security in global supply chains. Everledger is ISO 27001 certified. Data is divided into tiers, which allows users to grant access to particular tiers and data sets. Everledger's blockchain asset tracking function helps ensure asset provenance, traceability, and authenticity for multiple markets, including diamonds and gemstones, art registries, battery repurposing and recycling, critical minerals, and fashion and luxury brands, among others. See https://everledger.io/ for more information.

Storj

Storj is a decentralized cloud storage application that uses blockchain to encrypt data, split encrypted data into several pieces (each file is split into at least 80 pieces), and distribute the pieces to uncorrelated nodes for storage. The Storj network automatically reconstitutes the data for download; Storj claims it needs only 29 of the 80 pieces to fully reconstitute a file. Storj is compatible with Amazon S3 and provides an open-source code library. See https://www.storj.io/ for more information.

Brave

Brave offers a web browser and search function that uses blockchain to protect user privacy and does not track user activity or data. Brave also offers a cryptocurrency wallet built into its web browser. The wallet supports a number of decentralized finance applications and non-fungible tokens, and Brave also offers its own cryptocurrency, Basic Attention Token (BAT). See https://brave.com/ for more information.

References

[1] Gartner.com, "Gartner Hype Cycle". Accessed February 2, 2024. https://www.gartner.com/en/research/methodologies/gartner-hype-cycle.

[2] Parker, L., "Medieval Traders as International Change Agents: A Comparison with Twentieth Century International Accounting Firms." *Accounting Historians Journal* 16 (2): 107–118 (1989).

[3] Merkle, R. C. "A Digital Signature Based on a Conventional Encryption Function." In *Advances in Cryptology — CRYPTO '87*, Pomerance, C. (ed). Lecture Notes in Computer Science, 293, Springer, Berlin pp. 369–378 (1988).

[4] Nakamoto, S., "Bitcoin: A Peer-to-Peer Electronic Cash System." https://bitcoin.org/bitcoin.pdf.

[5] Romiti, M., Judmayer, A., Zamyatin, A., and Haslhofer, B., "A Deep Dive into Bitcoin Mining Pools: An Empirical Analysis of Mining Shares." In *2019 Workshop on the Economics of Information Security*, Boston, MA (2019).

[6] Galbraith, J. K., *Money: Whence It Came, Where It Went.* Princeton University Press, ISBN 978-0-691-17166-1 (2017).

[7] Yaga, D., Mell, P., Roby, N. and Scarfone, K., *Blockchain Technology Overview.* National Institute of Standards and Technology NISTIR 8202, Gaithersburg, MD (2018). Available at https://doi.org/10.6028/NIST.IR.8202.

[8] Sherman, A., Javani, F., Zhang, H. and Golaszewski, E. "On the Origins and Variations of Blockchain Technologies." *IEEE Security & Privacy* 17 (1), pp. 72–77, DOI 10.1109/MSEC.2019.2893730 (January/February 2019).

[9] Chaum, D. *Computer Systems Established, Maintained and Trusted by Mutually Suspicious Groups*, University of California, Berkeley, CA (1982).

[10] Wong, E., "Retrieving dispersed data from SDD-1: A system for distributed databases." *Proc. 2nd Berkeley Workshop Distributed Data Management and Computer Networks*, pp. 217–235 (May 1977).

[11] Lamport, L., "The Part-Time Parliament." *ACM Transactions on Computer Systems* 16 (2), pp. 133–169, DOI 10.1145/279227.279229 (May 1998).

[12] Back, A. *Hashcash – A Denial of Service Counter-Measure*, August 2002. Available at http://www.hashcash.org/papers/hashcash.pdf.

[13] Chen, J and Micali, S., *ALGORAND*, May 2017. Available at https://arxiv.org/abs/1607.01341.

[14] King, S. and Nadal, S. *PPCoin: Peer-to-Peer Crypto-Currency with Proof of Stake*, August 2012. Available at https://www.peercoin.net/whitepapers/peercoin-paper.pdf.

[15] Stefan Dziembowski, S., Faust, S., Kolmogorov, V. and Pietrzak, K. *Proofs of Space.* In Gennaro R., Robshaw M. (eds) Advances in Cryptology – CRYPTO 2015. Lecture Notes in Computer Science, vol 9216. Springer, Berlin, Heidelberg. Available at https://doi.org/10.1007/978-3-662-48000-7_29.

[16] Dwork, C. and Naor, M., *Pricing Via Processing or Combatting Junk Mail.* In Ernest F. Brickell, editor, CRYPTO'92, LNCS 740, pp. 139–147. Springer, August 1993.

[17] Antonopoulos, A., *Mastering Bitcoin: Programming the Open Blockchain* (2 ed.). USA: O' Reilly Media, Inc. p. Glossary. ISBN 978-1491954386 (2017).

[18] Siegel, D., "Understanding the DAO Attack" (June 2016). Available at https://www.coindesk.com/understanding-dao-hack-journalists.

[19] Szabo, N., *The Idea of Smart Contracts* (1997). https://www.fon.hum.uva.nl/rob/Courses/InformationInSpeech/CDROM/Literature/LOTwinterschool2006/szabo.best.vwh.net/idea.html.

[20] Schuplen, R., *Smart Contracts in the Netherlands: A Legal Research Regarding the Use of Smart Contracts Within Dutch Contract Law and Legal Framework.* Tilburg University (2019). http://arno.uvt.nl/show.cgi?fid=146860.

[21] *Integration Definition for Function Modeling (IDEF0)*, Federal Information Processing Standards Publication 183 (December 1993). https://web.archive.org/web/20090227122140/http://www.itl.nist.gov/fipspubs/idef02.doc.

[22] Orcutt, M., "Criminals Thought Bitcoin Was the Perfect Hiding Place, but They Thought Wrong." MIT Technology Review (September 11, 2017). https://www.technologyreview.com/2017/09/11/149211/criminals-thought-bitcoin-was-the-perfect-hiding-place-they-thought-wrong/.

[23] Nagarajan, S., "Bitcoin Miners Raked in More Than $1 Billion in Combined Earnings Last Month. Here's How They Make Money." *Business Insider*, February 20, 2021. https://markets.businessinsider.com/currencies/news/bitcoin-miners-bln-earnings-how-they-make-money-transactions-2021-2-1030103871.

[24] Li., K., "The Blockchain Scalability Problem & the Race for Visa-Like Transaction Speed." January 30, 2019. https://towardsdatascience.com/the-blockchain-scalability-problem-the-race-for-visa-like-transaction-speed-5cce48f9d44.

[25] Blockchain.com, "Average Confirmation Time." Accessed April 12, 2023. https://www.blockchain.com/charts/avg-confirmation-time.

[26] Kovach, S., "Tesla Buys $1.5 Billion in Bitcoin, Plans to Accept It as Payment." *CNBC*, February 8, 2021. https://www.cnbc.com/2021/02/08/tesla-buys-1point5-billion-in-bitcoin.html.

[27] McDonnell, T., "Can Bitcoin Ever Really Be Green?" Quartz, March 10, 2021. https://qz.com/1982209/how-bitcoin-can-become-more-climate-friendly/.

[28] Gartenberg, C., "Facebook Reportedly Plans to Launch Its Own Cryptocurrency." *The Verge*, May 11, 2018. https://www.theverge.com/2018/5/11/17344318/facebook-cryptocurrency-token-blockchain-report-david-marcus.

[29] Andriotis, A., Hoffman, L., Rudegeair, P., and Horwitz, J., "Facebook Building Cryptocurrency-Based Payments System." *The Wall Street Journal*, May 2, 2019. https://www.wsj.com/articles/facebook-building-cryptocurrency-based-payments-system-11556837547.

[30] Kastrenakes, J., "Libra cryptocurrency project changes name to Diem to Distance Itself From Facebook." *The Verge*, December 1, 2020. https://www.theverge.com/2020/12/1/21755078/libra-diem-name-change-cryptocurrency-facebook.

[31] Associated Press, "Venezuela Plans a Cryptocurrency, Maduro Says." *The New York Times*, December 3, 2017. https://www.nytimes.com/2017/12/03/world/americas/venezuela-cryptocurrency-maduro.html.

[32] Ellsworth, B., "Special Report: In Venezuela, New Cryptocurrency Is Nowhere to Be Found." *Reuters*, August 30, 2018. https://www.reuters.com/article/us-cryptocurrency-venezuela-specialrepor/special-report-in-venezuela-new-cryptocurrency-is-nowhere-to-be-found-idUSKCN1LF15U.

[33] Areddy, J., "China Creates Its Own Digital Currency, a First for Major Economy." *The Wall Street Journal*, April 5, 2021. https://www.wsj.com/articles/china-creates-its-own-digital-currency-a-first-for-major-economy-11617634118?st=wszp4bhdgermp1w&reflink=article_gmail_share.

[34] Orcutt, M., "Criminals Thought Bitcoin Was the Perfect Hiding Place, but They Thought Wrong." *MIT Technology Review*, September 2017. https://www.technologyreview.com/2017/09/11/149211/criminals-thought-bitcoin-was-the-perfect-hiding-place-they-thought-wrong/.

[35] Smee, S. "Beeple's Digital 'Artwork' Sold for More Than Any Painting by Titian or Raphael. But As Art, It's a Great Big Zero." *The Washington Post*, March 16, 2021. https://www.washingtonpost.com/entertainment/museums/beeple-digital-artwork-sale-perspective/2021/03/15/6afc1540-8369-11eb-81db-b02f0398f49a_story.html.

[36] Hern, A., "Bitcoin Currency Could Have Been Destroyed by '51%' Attack." *The Guardian*, June 16, 2014. https://www.theguardian.com/technology/2014/jun/16/bitcoin-currency-destroyed-51-attack-ghash-io.

[37] Canellis, D., "One Mining Pool Controls a Dangerously Dominant Share of BitcoinSV's Hash Rate." *The Next Web*, April 26, 2019. https://thenextweb.com/news/calvin-ayre-coingeek-bitcoin-sv-hash-rate-51-percent-double-spend.

[38] Farivar, C., "Bitcoin Pool GHash.io Commits to 40% Hashrate Limit After Its 51% Breach." *Ars Technica*, July 16, 2014. https://arstechnica.com/information-technology/2014/07/bitcoin-pool-ghash-io-commits-to-40-hashrate-limit-after-its-51-breach/.

[39] Roberts, J., "Bitcoin Spinoff Hacked in Rare '51% Attack'." *Fortune*, May 29, 2018. https://fortune.com/2018/05/29/bitcoin-gold-hack/.

[40] Voell, Z., "Ethereum Classic Hit by Third 51% Attack in a Month." Coindesk, August 29, 2020. https://www.coindesk.com/ethereum-classic-blockchain-subject-to-yet-another-51-attack.

[41] Hertig, A., "Bitcoin Cash Miners Undo Attacker's Transactions With '51% Attack.'" Coindesk, May 24, 2019. https://www.coindesk.com/bitcoin-cash-miners-undo-attackers-transactions-with-51-attack.

[42] Voell, Z., "Bitcoin Node Count Falls to 3-Year Low." May 6, 2020. https://www.coindesk.com/bitcoin-node-count-falls-to-3-year-low-despite-price-surge.

[43] Tan, K-L., "Distributed Database Systems." In *Encyclopedia of Database Systems*, Liu, L. and Özsu, T. (eds), 2009. Springer Link. https://link.springer.com/content/pdf/10.1007%2F978-0-387-39940-9_701.pdf.

[44] Wessler, B., "Man Behind Silk Road Website Is Convicted on All Counts." *New York Times*, Feb. 4, 2015. https://www.nytimes.com/2015/02/05/nyregion/man-behind-silk-road-website-is-convicted-on-all-counts.html.

[45] Johnson, J., "Spam: Share of Global Email Traffic 2007-2019." Statista, January 25, 2021. https://www.statista.com/statistics/420400/spam-email-traffic-share-annual/.

[46] Faridi, O., "Threat of Quantum Computing to Bitcoin Should be Taken Seriously, But there's Enough Time to Upgrade Current Security Systems, Experts Claim." Crowdfund Insider, November 7, 2020. https://www.crowdfundinsider.com/2020/11/168869-threat-of-quantum-computing-to-bitcoin-should-be-taken-seriously-but-theres-enough-time-to-upgrade-current-security-systems-experts-claim/.

[47] Franklin, D. and Chong, F., "Challenges in Reliable Quantum Computing." In Shukla S.K., Bahar R.I. (eds) *Nano, Quantum and Molecular Computing*. Springer, Boston, MA (2004). https://doi.org/10.1007/1-4020-8068-9_8.

[48] Feldman, S., 20 Years of Quantum Computing Growth. Statista, May 6, 2019. https://www.statista.com/chart/17896/quantum-computing-developments/.

[49] Qingqing, C., "Chinese Team Develops Quantum Processor With World's Largest Number of Superconducting Qubits." CGTN, May 11, 2021. https://news.cgtn.com/news/2021-05-10/Chinese-team-makes-new-breakthrough-in-quantum-computing-technology-1091kWpt9tK/index.html.

[50] *Quantum Computing: Progress and Prospects*. The National Academies Press (2019). http://nap.edu/25196.

[51] Cho, A., "IBM Promises 1000-Qubit Quantum Computer—a Milestone—by 2023." *Science*, September 15, 2020. https://www.sciencemag.org/news/2020/09/ibm-promises-1000-qubit-quantum-computer-milestone-2023.

[52] Arute, F., Arya, K., Babbush, R. et al., "Quantum Supremacy Using a Programmable Superconducting Processor." *Nature* 574, 505–510 (2019). https://doi.org/10.1038/s41586-019-1666-5.

[53] Dinh, T. Wang, J et al., *BLOCKBENCH: A Framework for Analyzing Private Blockchains* (2017). https://arxiv.org/pdf/1703.04057.pdf.

[54] Roeder, L, *Building Trust in Blockchain for the Electric Grid.* Pacific Northwest National Laboratory, March 29, 2019. https://www.pnnl.gov/news-media/building-trust-blockchain-electric-grid.

[54] Pretz, K., "How Blockchain Technology Could Track and Trace Food From Farm to Fork." IEEE Spectrum, August 29, 2018. https://spectrum.ieee.org/the-institute/ieee-products-services/how-blockchain-technology-could-track-and-trace-food-from-farm-to-fork.

[55] Laaper, S. et al., *Using Blockchain to Drive Supply Chain Innovation.* September 2017. https://www2.deloitte.com/us/en/pages/operations/articles/blockchain-supply-chain-innovation.html.

[56] Lamport, L, Shostak, R., and Pease, M., "The Byzantine Generals Problem." *ACM Transactions on Programming Languages and Systems* 4(3), pp. 382–401, July 1982.

[57] Perna, Ltg. G., "Integration Starts With Enterprise Resource Planning Systems." *Army Sustainment*, p. 2, November–December 2017. https://alu.army.mil/alog/2017/NOVDEC17/PDF/NOVDEC2017.pdf.

[58] Neumann, A.J., Statland, N., and Webb, R.D., "Post-Processing Audit Tools and Techniques." *Proc. NBS Invitational Workshop,* Miami, Florida, 1977. https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nbsspecialpublication500-19.pdf.

[59] Campbell, C., "What Are the 4 Different Types of Blockchain Technology?" *TechTarget*, March 3, 2023, https://www.techtarget.com/searchcio/feature/What-are-the-4-different-types-of-blockchain-technology.

[60] Sharma, T.K., "What Is Hybrid Blockchain? How Can It Help To Solve Everyday Problems?" *Blockchain Council*, October 12, 2018, https://www.blockchain-council.org/blockchain/what-is-hybrid-blockchain-how-can-it-help-to-solve-everyday-problems/.

[61] Huang, K., "Why Did FTX Collapse? Here's What to Know." *The New York Times*, November 10, 2022, https://www.nytimes.com/2022/11/10/technology/ftx-binance-crypto-explained.html.

[62] Melinek, J., "Top Crypto App Downloads Rise Over 15% Following SVB Collapse," *TechCrunch*, March 16, 2023, https://techcrunch.com/2023/03/16/top-crypto-app-downloads-rise-over-15-following-svb-collapse/.

[63] Huestis, S., "Cryptocurrency's Energy Consumption Problem," *RMI*, January 20, 2023, https://rmi.org/cryptocurrencys-energy-consumption-problem/#:~:text=Bitcoin%20alone%20is%20estimated%20to,fuel%20used%20by%20US%20railroads.

[64] Dance, G.J.X., "The Real-World Costs of the Digital Race for Bitcoin," *The New York Times*, April 9, 2023, https://www.nytimes.com/2023/04/09/business/bitcoin-mining-electricity-pollution.html?campaign_id=9&emc=edit_nn_20230410&instance_id=89828&nl=the-morning&regi_id=58911296&segment_id=130026&te=1&user_id=55bb8cb83cab021a5c8fb8553fcd804a.

[66] Carter, R., "How Many Bitcoins Are There in 2023?" *BanklessTimes*, February 16, 2023, https://www.banklesstimes.com/how-many-bitcoins-are-there/.

[67] Nelson, H., "VA Suspends Further Oracle Cerner EHR Implementations for EHRM 'Reset'," *EHR Intelligence*, April 24, 2023, https://ehrintelligence.com/news/va-suspends-further-oracle-cerner-ehr-implementations-for-ehrm-reset.

[68] Government Accountability Office (GAO), GAO-23-106685, "Electronic Health Record Modernization: VA Needs to Address Change Management Challenges, User Satisfaction, and System Issues," March 15, 2023, https://www.gao.gov/products/gao-23-106685.

[69] Dubovitskaya A, Baig F, Xu Z, et al. ACTION-EHR: Patient-Centric Blockchain-Based Electronic Health Record Data Management for Cancer Care. *J Med Internet Res*. 2020;22(8):e13598. Published 2020 Aug 21. doi:10.2196/13598. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7474412/.

[70] Castor, A., "Ethereum Moved to Proof Of Stake. Why Can't Bitcoin?" *MIT Technology Review*, February 28, 2023, https://www.technologyreview.com/2023/02/28/1069190/ethereum-moved-to-proof-of-stake-why-cant-bitcoin/.

[71] Brooks, M., "IBM Wants to Build a 100,000-Qubit Quantum Computer," *MIT Technology Review*, May 25, 2023, https://www.technologyreview.com/2023/05/25/1073606/ibm-wants-to-build-a-100000-qubit-quantum-computer/.

[72] White House Office of Science and Technology Policy, "Technical Evaluation for a U.S. Central Bank Digital Currency System, September 2022, https://www.whitehouse.gov/wp-content/uploads/2022/09/09-2022-Technical-Evaluation-US-CBDC-System.pdf.

[73] Pearson, J. and Zinets, N., "Deepfake Footage Purports to Show Ukrainian President Capitulating," *Reuters*, March 17, 2022, https://www.reuters.com/world/europe/deepfake-footage-purports-show-ukrainian-president-capitulating-2022-03-16/.

[74] Bond, S., "Fake Viral Images of an Explosion at the Pentagon Were Probably Created by AI," *NPR*, May 22, 2023,

https://www.npr.org/2023/05/22/1177590231/fake-viral-images-of-an-explosion-at-the-pentagon-were-probably-created-by-ai.

[75] Abrol, A., "Top 5 Stablecoins In 2023," *Blockchain Council*, March 29, 2023, https://www.blockchain-council.org/cryptocurrency/list-of-stablecoins/#:~:text=Tether%20(USDT)%20is%20the%20most,pegged%20to%20the%20U.S.%20dollar.

[76] Roozenbeek, J., van der Linden, S., and Nygren, T., "Prebunking Interventions Based on "Inoculation" Theory Can Reduce Susceptibility to Misinformation Across Cultures," *Harvard Kennedy School Misinformation Review*, February 3, 2020, https://misinforeview.hks.harvard.edu/article/global-vaccination-badnews/.

[77] Cordell, C., "CBP Wants Blockchain Capability for Its Trade Portal Modernization," *Nextgov/FCW*, July 24, 2023, https://www.nextgov.com/cxo-briefing/2023/07/ftc-hhs-warn-potential-privacy-and-security-risks-embedded-online-health-sites/388708/.

[78] Grigg, G., "Blockchain Analysis for National Security and Law Enforcement Agencies: A Primer," *Chainanalysis*, July 21, 2022, https://www.chainalysis.com/blog/blockchain-analysis-national-security-law-enforcement/.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE (DD-MM-YY)	2. REPORT TYPE	3. DATES COVERED (From – To)
00-01-24	Final	

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
On Blockchains: The Hype, the Technology, and the Potential Applications for the Department of Defense	HQ0034-19-D-0001
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBERS

6. AUTHOR(S)	5d. PROJECT NUMBER
Steven P. Wartik, Michelle G. Albert	C520824
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESSES	8. PERFORMING ORGANIZATION REPORT NUMBER
Institute for Defense Analyses 730 East Glebe Road Alexandria, VA 22305	3001138

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR'S / MONITOR'S ACRONYM
Institute for Defense Analyses 730 East Glebe Road, Alexandria, VA 22305	IDA
	11. SPONSOR'S / MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT

Distribution Statement A. Approved for public release: distribution is unlimited.

13. SUPPLEMENTARY NOTES

Project Leader: Steven P. Wartik

14. ABSTRACT

Blockchains—particularly, the cryptocurrencies they enable—rank among the most-hyped technologies of the 21st century, but their potential uses range far beyond crypto. Blockchains provide a means of securely recording data that ensures the data's confidentiality, integrity, and availability—essential characteristics of a modern, secure system that can provide immense value to the Department of Defense (DoD). This paper attempts to cut through the hype around blockchains and looks at them as a foundational technology that may affect how DoD operates now and in the future.

15. SUBJECT TERMS

blockchain, confidentiality, availability, integrity, cryptocurrency, cybersecurity, mining, privacy, smart contract

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Institute for Defense Analyses
Unclassified	Unclassified	Unclassified	Unlimited	70	19b. TELEPHONE NUMBER (Include Area Code)