




MATS
UNIVERSITY

NAAC
GRADE **A+**
ACCREDITED UNIVERSITY

MATS CENTRE FOR OPEN & DISTANCE EDUCATION

Blockchain Technology

Bachelor of Computer Applications (BCA)
Semester - 4



SELF LEARNING MATERIAL



MATS UNIVERSITY

www.matsuniversity.ac.in



Bachelor of Computer Applications

ODL BCA DSE 01

BLOCKCHAIN TECHNOLOGY

Course Introduction	
MODULE 1 BLOCKCHAIN TECHNOLOGY	3
Unit 1: Introduction to Blockchain and Its Architecture	4
Unit 2: Blockchain Consensus and Security	11
Unit 3: Blockchain Applications and Development	21
MODULE 2 BASIC CRYPTOGRAPHIC PRIMITIVES BEHIND THE BLOCKCHAIN	39
Unit 4: Hash Functions and Cryptographic Foundations	40
Unit 5: Digital Signatures and Cryptographic Techniques	51
Unit 6: RSA Algorithm and Key Management	59
MODULE 3 BITCOIN BASICS	93
Unit 7: Introduction to Bitcoin and Its Working	94
Unit 8: Bitcoin Creation, Transactions, and Governance	99
Unit 9: Bitcoin Security, Price Volatility, and Future Trends	108
MODULE 4 CONSENSUS	127
Unit 10: Fundamentals of Consensus in Distributed Systems	128
Unit 11: Consensus Protocols in Distributed Systems	134
Unit 12: Consensus in Open Systems and Blockchain	150
MODULE 5 PERMISSIONED BLOCKCHAIN	169
Unit 13: Permissioned Blockchain and Its Use Cases	170
Unit 14: Smart Contracts and Their Limitations	181
Unit 15: State Machine Replication in Blockchain	192
References	205

COURSE DEVELOPMENT EXPERT COMMITTEE

Prof. (Dr.) K. P. Yadav, Vice Chancellor, MATS University, Raipur, Chhattisgarh

Prof. (Dr.) Omprakash Chandrakar, Professor and Head, School of Information Technology, MATS University, Raipur, Chhattisgarh

Prof. (Dr.) Sanjay Kumar, Professor and Dean, Pt. Ravishankar Shukla University, Raipur, Chhattisgarh

Prof. (Dr.) Jatinderkumar R. Saini, Professor and Director, Symbiosis Institute of Computer Studies and Research, Pune

Dr. Ronak Panchal, Senior Data Scientist, Cognizant, Mumbai

Mr. Saurabh Chandrakar, Senior Software Engineer, Oracle Corporation, Hyderabad

COURSE COORDINATOR

Dr. Vaibhav Sharma, Assistant Professor, MATS University, Raipur, Chhattisgarh

COURSE PREPARATION

Dr. Vaibhav Sharma, Assistant Professor and Mr. Digvijay Singh Thakur, Assistant Professor, MATS University, Raipur, Chhattisgarh

March, 2025

ISBN: 978-93-49916-28-9

@MATS Centre for Distance and Online Education, MATS University, Village-Gullu, Aarang, Raipur-(Chhattisgarh)

All rights reserved. No part of this work may be reproduced or transmitted or utilized or stored in any form, by mimeograph or any other means, without permission in writing from MATS University, Village- Gullu, Aarang, Raipur-(Chhattisgarh)

Printed & Published on behalf of MATS University, Village-Gullu, Aarang, Raipur by Mr. Meghanadhu Katabathuni, Facilities & Operations, MATS University, Raipur (C.G.)

Disclaimer-Publisher of this printing material is not responsible for any error or dispute from contents of this course material, this is completely dependent on AUTHOR'S MANUSCRIPT.

Printed at: The Digital Press, Krishna Complex, Raipur-492001 (Chhattisgarh)

Acknowledgement

The material (pictures and passages) we have used is purely for educational purposes. Every effort has been made to trace the copyright holders of material reproduced in this book. Should any infringement have occurred, the publishers and editors apologize and will be pleased to make the necessary corrections in future editions of this book.

COURSE INTRODUCTION

This **Blockchain Technology** course is designed to provide students with a deep understanding of distributed ledger technologies, cryptographic principles, Bitcoin fundamentals, consensus mechanisms, and permissioned blockchain models. The course aims to equip students with the knowledge required to understand blockchain systems, secure transactions, and develop blockchain-based applications. The course is structured into five modules, each covering key aspects of blockchain technology and its applications.

Module 1: Blockchain Technology

This module introduces students to the fundamentals of blockchain technology. It covers the core principles of distributed ledgers, blockchain architecture, design, and various consensus protocols used to maintain security and integrity. The module also explores privacy concerns, real-world use cases in finance, supply chain, and government sectors, and provides an overview of **Hyperledger Fabric**, a permissioned blockchain framework for enterprise applications.

Module 2: Basic Cryptographic Primitives Behind Blockchain

Cryptography is the backbone of blockchain technology, ensuring data security and integrity. This module delves into cryptographic techniques, including **hash functions, digital signatures, and public-key encryption**, which form the basis of blockchain security. Students will explore **RSA encryption** and how public and private keys are used for securing blockchain transactions.

Module 3: Bitcoin Basics

Bitcoin was the first practical implementation of blockchain technology, and this module provides an in-depth study of how Bitcoin works. Students will learn about the creation and governance of Bitcoin, how transactions occur, and security concerns related to its use. The module also examines Bitcoin's price volatility, market trends, and its future prospects in the global financial landscape.

Module 4: Consensus Mechanisms

Consensus mechanisms are essential for maintaining blockchain integrity and preventing fraud. This module covers various consensus protocols, including **Proof-of-Work (PoW)** and **Proof-of-Stake (PoS)**, and their impact on blockchain security. The course also introduces students to the differences between **synchronous and asynchronous systems** and explores how consensus is achieved in both open and permissioned blockchain networks.

Module 5: Permissioned Blockchain Models

This module focuses on **permissioned blockchains**, which differ from public blockchains by allowing only authorized participants to interact with the network. Students will explore how **smart contracts** enable automated transactions, as well as **state machine replication** and its role in ensuring blockchain consistency. The module also examines real-world use cases, design challenges, and limitations of permissioned blockchain systems.

MODULE 1

BLOCKCHAIN TECHNOLOGY

LEARNING OUTCOMES

By the end of this Module, students will be able to:

- Understand the fundamentals of blockchain and its key components.
- Explain blockchain architecture, design, and protocols.
- Understand blockchain consensus mechanisms and their significance.
- Identify security and privacy challenges in blockchain technology.
- Explore real-world blockchain applications in finance, supply chain, and government.
- Learn about Hyperledger Fabric as a blockchain development

Unit 1: Introduction to Blockchain and Its Architecture

1.1 Introduction to Blockchain

The technology behind blockchain was born in 2008, with the release of the Bitcoin whitepaper authored by the pseudonymous Satoshi Nakamoto. Essentially, the blockchain is a new way of storing data and establishing trust in digital systems. In contrast to centralized systems where one entity is in control and validates information, blockchain is a decentralized mechanism where the responsibility of establishing authority is over a network of participants. This paradigm shift introduces a mechanism of trust based on mathematical algorithms and decentralized consensus instead of centralized intermediaries. That's what blockchain is all about, immutability regarding transactions that is transparent, secure and cannot be altered. Each transaction is bundled together with others into a "block" that is then cryptographically linked to the previous blocks, forming a "chain" of transaction history. This structure is what gives blockchain its name and its potent attribute of immutability—once information is entered and confirmed by the network, it is virtually impossible to alter or delete. Blockchain is not a mere record-keeping innovation. It also introduces the notion of trustless systems where parties with no knowledge or trust of each other can interact and transact securely without intermediaries. It is because of the fusion of distributed consensus mechanisms, cryptographic verification with economic incentivization that guides participants to coalesce around and maintain the integrity of the network. Blockchain technology came into the limelight with cryptocurrencies such as bitcoin itself, showing its ability to create secure digital transmission of value without reliance on third parties. But blockchain is capable of much more than digital currencies. Enterprise blockchain solutions now solve problems for sectors as diverse as supply chains, healthcare record systems, digital identity verification and intellectual property rights management. Never another tool yet still, we are working up until now with such devices without most extreme ace in the behind thus do yet without a doubt debate, they are never the best nor well in agreement. Public blockchains, including Bitcoin and Ethereum, are open for participation, as anyone who wishes to join the network can do so, validate transactions, and maintain a copy of the ledger. By comparison, private or permissioned blockchains limit participation to authorized



Notes

participants, classically sacrificing full decentralization for improved efficiency and regulatory compliance. On the other hand, hybrid models adopt characteristics from both classes to provide a middle ground between openness and controlled access.

Blockchain technology has evolved through several generations. Blockchains of the first generation were mostly concerned with facilitating the transfer of cryptocurrency. Second-generation platforms such as Ethereum brought programmability through smart contracts— self-executing agreements where the terms are written directly into code. This set a new paradigm, allowing on-chain transactions to be part of complex decentralized applications (dApps), containing entire bots auto-executing business logic. Referred to as the third generation of blockchain, these projects aim to solve the scalability, energy consumption and interoperability challenges, helping the tech become suitable for mainstream and enterprise adoption. That said, blockchain does come with some considerable challenges. Many platforms still face limitations in, throughput per se, and are still struggling with scalability in comparison to centralized systems. Such energy consumption, especially in proof-of-work consensus mechanisms, raises questions about environmental sustainability. Wide adoption is also limited by regulatory uncertainty, security vulnerabilities, and technical complexity of implementation. However blockchain technology is still evolving quickly, as many researchers and developers are addressing these issues. The technology has evolved toward more efficient, usable, and sustainable solutions with the introduction of new consensus mechanisms, layer-2 scaling solutions, cross-chain interoperability protocols, and improved governance models. With organizations and governments alike beginning to appreciate the transformative potential of blockchain for digital dealings, we are seeing increased investment in both public and enterprise blockchain projects. That is, the underlying value proposition of the technology — providing the means of secure, transparent, and efficient digital transactions without centralized control — means it can be a foundational piece of the digital economy and next-generation internet infrastructure of the future.

1.2 Blockchain – Architecture, Design and Protocol

This is how blockchain systems are designed to work: It is an interesting architecture of various technology components integrated together for a distributed ledger that is secure. Looking at this architecture entails looking not only at the structural building blocks, but also at the processes underlying its unique capabilities. A distributed ledger — a synchronized database replicated on each node in the network — lies at the very core of blockchain architecture. This ledger is maintained collectively by network participants as opposed to traditional database controlled by central administrators. The nodes all have a full copy of the blockchain in a stored database, so if any particular node goes offline, the remaining nodes still preserve the data and provide huge redundancy, thus removing single points of failure. Blocks are the fundamental structural elements of the blockchain. Each block has three main components: the cryptographic hash of the last block (which forms the connection, or the "chain"), a timestamp, and the transaction data. This makes it so that once information is encoded into a block, it is virtually impossible to change, as doing so would require recalculating the hashes for all blocks after that point, which would lead to a broken chain that all other nodes in the network would recognize immediately. By contrast, transactions are the basic units of work recorded onto the blockchain.

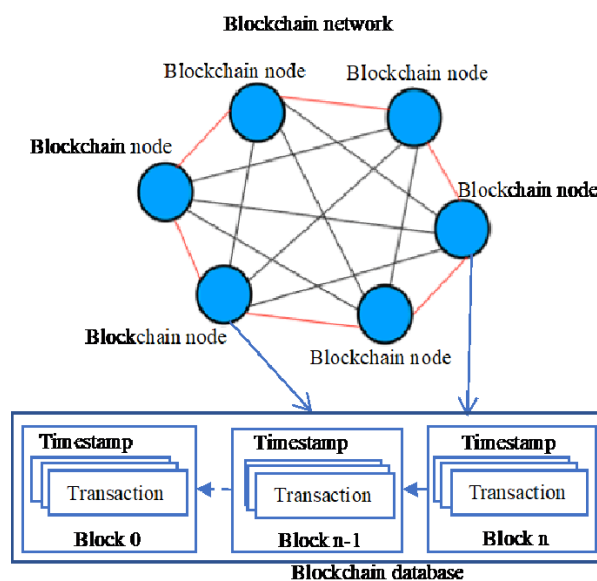


Figure 1: Architecture of Blockchain
[source – www.researchgate.net]



Notes

A transaction typically consists of a sender and a recipient (usually as cryptographic addresses), a value or data that is being transferred, and a digital signature that proves the sender has authorized the transfer. Transactions are verified by validating nodes according to the rules of the network's protocol before being included in a block, ensuring they adhere to requirements around format, digital signatures, and state transitions. Encryption is an important aspect of Blockchain security and integrity. Public-key cryptography provides the foundation for secure digital identities based on key pairs: private key—secret, public key—publicly known to others—for signing and verifying transactions, respectively. Cryptographic hash functions, which provide fixed-length, unique representations of arbitrary data, enable efficient verification of data integrity. A minor alteration in the material results in a substantially diverse hash output, so it becomes clear if material has been manipulated. The P2P network architecture is the communication layer of the blockchain system. Instead, nodes connect with each other to share the latest transactions and blocks without the need for central servers. By operating in a decentralized fashion, there are no single points of control, or failure, making it more resilient towards attacks and censorship. The various protocols used by the P2P network are for data propagation, node discovery, and maintenance of the network. It is a large advancement of the architecture, and it uses by second and third-generation blockchains. Executed through self-executing programs, smart contracts run on the blockchain, automatically carrying out functions when certain conditions are met. Smart contracts create user defined functions and application behaviors that extend the blockchain from simple value transfer to complex business logic, open the door for decentralized applications, automated agreements and programmable digital assets. State is the current snapshot of all accounts, contracts, and data on the blockchain. Each transaction changes the state according to rules of the protocol. Different implementations of blockchains choose different means of keeping track of their state—Bitcoin has an Unspent Transaction Output (UTXO) model to keep track of which outputs to which previous transactions have not yet been spent, while Ethereum chooses an account-

based model, keeping track of the balances and state information for each address.

There are certainly real-world differences between blockchain protocols, and they generally have diverse mechanisms regarding network governance and protocol upgrades. On-chain governance allows stakeholders to vote on proposed changes using the blockchain itself, while off-chain governance derives social consensus among developers, miners, and community members. Fork mechanics enable modifications of the protocols via soft forks (backward compatible changes) or hard forks (breaking changes that require all participants to upgrade). The architectural design of blockchains consists of inherent trade-offs between three important properties commonly composed in to the "blockchain trilemma" : decentralization, security and scalability. Most blockchains are optimized for two of these properties, at the expense of the other. Bitcoin, for instance, prioritizes security and decentralization but at some cost to the transaction throughput obtainable, while more decentralized blockchains can also achieve higher levels of scalability but only at potential cost to decentralization. There are several different types of blockchain architectures, which are distinguished by their access models. Public blockchains enable any user to validate transactions and maintain the ledger, which maximizes transparency and censorship resistance. Private blockchains limit these roles to approved participants, allowing for more control and privacy at the expense of decentralization. In consortium blockchains, control is distributed to a permissioned-set of organizations, providing a mild trade-off between decentralization and requirements on privacy and efficiency. Layer-2 scaling solutions are architectural add-ons, providing new capabilities to the base blockchain tier. There are payment channels, which allow for off-chain transactions with periodic settlements to the main chain, sidechains, which digests transact in a specialized chain but has tether to the main chain, and rollups, which batch multiple transactions into a single memoir that is then submitted to the main chain. These approaches seek to address scalability constraints while benefiting from the security of the underlying blockchain. The problem of communication between blockchains was solved by interoperability protocols. Cross-chain bridges allow for the transfer of assets between two or more separate blockchains; interoperability standards define common formats or interfaces for different blockchains to interact with each other.

As the blockchain ecosystem becomes more diverse and the delineation between blockchains becomes less, these sets of native components become more and more important for blockchains to be able to leverage each other. Different Blockchain implementations have varying approaches to data storage. Blockchains all keep transaction records, but their response to extra data diverges. Others bring data into the chain directly, maximizing availability, but necessitating larger storage requirements for the nodes. Others turn to off-chain storage solutions such as InterPlanetary File System (IPFS) or decentralized storage networks, writing only cryptographic

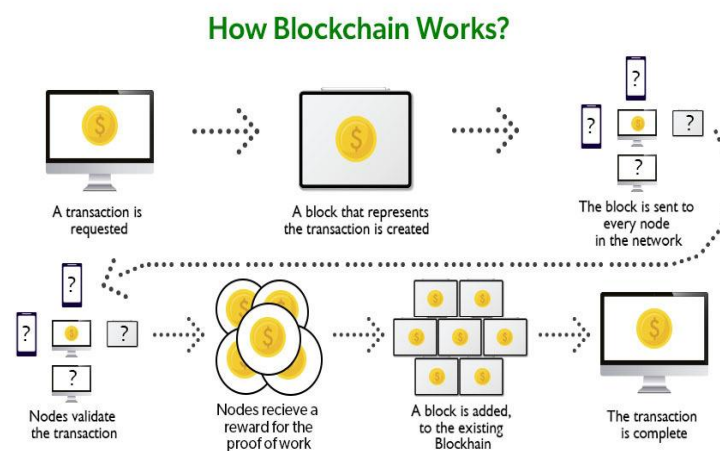


Figure 2: Working Methodology of Blockchain
[source - <https://komodoplatform.com>]

references on-chain to verify the integrity of externally stored information. This is an essential architectural consideration for enterprise applications, and privacy-enhancing technologies build on top of this concept.

Zero knowledge proofs make it possible to prove information without disclosing the original data. With ring signatures, several signatures are combined to shroud all transaction participants. Changes the nature of the transaction: Monies become part of transactions and transaction values are confidential, but validation is still possible. They solve the problem of conflicting transparency requirements of the public blockchain and privacy requirements of the business transactions. The architecture of a blockchain system can be divided into multiple software layers, including the consensus layer for consensus on the blockchain state, networking layer for peer communication, data layer for storage and retrieval, application layer for



Notes

user and application interfaces, and the security layer for cryptographic operations and access controls. The key innovations in blockchain architectures aim to eliminate these limitations while retaining the technology to continue exploiting its benefits. Composable, modular architectures enable different blockchain components to specialize and interoperate. Sharding techniques split the blockchain state across a subset of nodes to enhance sharding scalability. Another feature is that zero-knowledge architectures improve privacy and efficiency. Such developments derive from the continuous refinement of blockchain architecture towards a more scalable, flexible, and useful technology for multiple use cases.

Unit 2: Blockchain Consensus and Security

At the heart of blockchain technology are consensus protocols, a mechanism that gives the decentralized participants of the network a way to achieve consensus on what the current valid state of the distributed ledger is without having to rely on centralized coordination. These mechanisms address one of the most central obstacles to reliable consensus: that participants in a distributed system can be anonymous, geographically distributed, and potentially malicious. Consensus protocols are among the most active areas of blockchain research, with a variety of approaches that offer various trade-offs between security, decentralization, performance, and energy efficiency. The original blockchain was Proof of Work (PoW), which was offered by Bitcoin. It involves participants (miners) solving resource-heavy cryptographic puzzles to validate transactions and add new blocks to the chain. The challenge of these puzzles adjusts automatically, ensuring that the block time remains constant, even when we add more computing power to the network. When a miner solves the puzzle, they broadcast their solution to the rest of the network, and other nodes can easily verify for themselves that the solution is correct. This mechanism discourages attacks as it demands immense computational power to take over 51% of the mining power to be able to execute the attack.

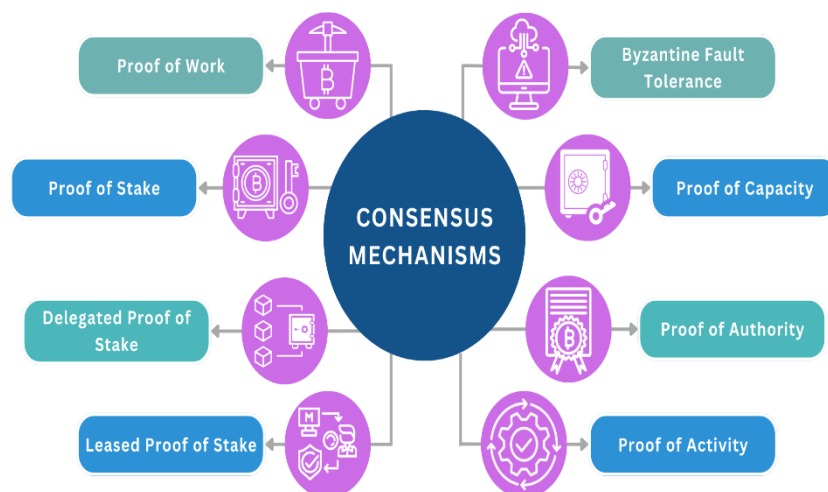


Figure 3: Consensus in Blockchain

Nonetheless, PoW is heavily criticized because the competitive mining process expends a lot of electricity, causing significant environmental concerns. You refer to it as proof-of-work, and from what I know one of the disadvantages of proof-of-work is the centralization of the mining process due to the economy of scale that is not decentralised. In contrast, Proof of Stake (PoS) emerged as an energy-efficient alternative to PoW. Instead of competing with computational work, validators are chosen to write blocks based on how much cryptocurrency they “stake” as collateral. This method removes the processor-intensive mining process, significantly decreasing energy consumption. Validators can be slashed (lost money) for malicious behavior or operational downtime, offering an economic security model like PoW without the requisite energy use. Many major blockchain platforms such as Ethereum (after its "Merge" upgrade), Cardano and Polkadot use variants of PoS consensus. Critics cite the risk of power being concentrated in the hands of wealthy stakeholders, as well as the so-called "nothing-at-stake" problem, where validators may be incentivized to validate both conflicting chains (although modern PoS implementations have proven solutions to these issues). Delegated Proof of Stake (DPoS) is an approach based on PoS that integrates a democratic representation system. An example of this mechanism is that token holders vote one-time to elect a few delegates (typically 21-100 delegates) that will verify a transaction and create a block. This voting mechanism provides accountability; on-site delegates who aren't pulling their weight can be voted out. By limiting the block production to a select group of trusted delegates, DPoS can deliver significantly higher transaction throughput than traditional PoS or PoW systems. Variations of DPoS are used by blockchains such as EOS, TRON, and Lisk. The main argument against DPoS is its relative lack of decentralization, as only a pale of the majority of those can produce blocks, this creates points of failure and censorship. PBFT and its variants provide a solution to the Byzantine Generals Problem -- several parties must agree. The PBFT consensus relies on one primary validator to suggest blocks and other validators to exchange messages to affirm that suggestion. The system can tolerate up to a third of the nodes are Byzantine (arbitrarily malicious) and still reach consensus. PBFT guarantees high transaction throughput and immediate finality (once a transaction is included in a block, it gets



Notes

confirmed without the need to wait for any confirmations). But it requires known validator identities and scales poorly as connecting cost increases quadratically with validator count. Adaptations of PBFT exist in permissioned blockchains such as Hyperledger Fabric and some public-private hybrid systems. Proof of Authority is a reputation-based consensus mechanism in which block validators are chosen based on their identities. These trusted validators are staking their reputation rather than crypto; engaging in malicious behavior threatens their professional and reputational capital. PoA provides high throughput with low computational overhead, and succeeds in permissioned systems with pre-identified participants. Various networks such as VeChain, POA, and many enterprise blockchain implementations use PoA due to its efficiency and clear accountability system. Its main drawback is its centralized approach — optimization comes at the cost of decentralization, hence, it is better suited for consortium blockchains or private networks than fully public ones. For example, Proof of Elapsed Time (PoET), an Intel project for the Hyperledger Sawtooth platform, achieves a lottery-based consensus that has each validator request a random wait time from a trusted execution environment (Intel provides the Software Guard Extensions). The validator that waited the least time receives the permission to create the next block. This approach retains PoW's fairness, without its energy consumption — the puzzle solving portion is replaced by a lottery and, therefore, saving energy. However, PoET's reliance on proprietary hardware and trust in the manufacturer of that hardware does present potential centralization concerns. In Proof of Burn (PoB), validators must "burn" the cryptocurrency by sending it to unspendable addresses in order to reach consensus. This entails sending coins to a verifiable address where they are permanently destroyed, leading to a commitment to the network and granting mining rights in proportion to the amount destroyed. PoB exchanges the continual energy expenditure cost of proof of work for a one-time resource commitment which makes PoB a significantly more sustainable choice. Implementation examples shall include Slimcoin and certain token distribution mechanisms. Others have criticized the permanent removal of coins from circulation, and questioned whether these coins have any long-term economic sustainability.

The Proof of Space (PoSpace) or Proof of Capacity approach requires validators to reserve large amounts of storage space to earn the right to create blocks. Now they commit large amounts of storage on the assumption that they can publish millions of possible proofs of the cryptographic puzzle and, that the probability of getting a valid block, will be larger and larger. This method uses available storage rather than using energy to compute continuously. Variations of this consensus method are used in cryptocurrencies such as Chia and Spacemesh. PoSpace is indeed more efficient than PoW, but still consumes resources, with the need for competitive participants to have low-cost storage capacity. Some hybrid consensus mechanisms are formed from different mechanisms to take advantages of the strength of different consensus mechanisms while compensating for its weaknesses. As an example of this approach, Algorand combines a lottery-based selection of validators using Verifiable Random Functions (VRFs) with Byzantine agreement among the chosen validators in two phases. Casper FFG, which was an early effort that aimed to complement PoW block production with PoS finality confirmation as Ethereum transitioned to PoS at the protocol level. Such hybrid approaches provide better security-performance trade-offs than their single-method counterparts. DAG based consensus is a structural paradigm shift from traditional blockchain architecture. Instead of batching transactions into chronological blocks, DAG systems create a larger graph structure between transactions, enabling new transactions to verify older transactions in parallel. Such structure allows for concurrent processing of transactions yielding a higher throughput than traditional blockchains. Several projects, like IOTA's Tangle, Hedera Hashgraph, or Nano, take a DAG-based paradigm with various consensus procedures. Challenges are maintaining global state across the network and transaction finality. Federated Byzantine Agreement (FBA), which is used in networks such as Stellar and Ripple, is a mechanism whereby nodes choose their own “quorum slices”—sets of trusted peers they care about the consensus of. When overlapping quorum slices produce agreement across the entire network, the system reaches global consensus. It is a more streamlined, lower capacity solution, while still providing more versatility than traditional BFT approach. Critics point out that FBA security is based on the correct configuration of the quorum slice, and can potentially introduce a centralization risk as the largest participants may have an outsized influence on quorum formation. The



Notes

Avalanche consensus, implemented by the Avalanche platform, uses a probabilistic approach where nodes repeatedly sample random subsets of peers to validate transactions. This metastable voting procedure rapidly converges to network-wide consensus with sub-second finality and large throughput. Avalanche enables very strong resistance to 51% attacks and requires minimal energy use (similar to PoS systems). As the mechanism is relatively new, the long-term security properties are still being evaluated in the wild. A lot of research has been done to tackle challenges regarding consensus mechanisms, which are still evolving. In fact, there are sharded consensus approaches that partition node validators into node subsets responsible for reaching consensus, which are already being tried in networks like Ethereum 2.0(Serenity), Zilliqa and so on. VRF Based Mechanisms: They use cryptographic functions to select validators in a way that is unpredictable but verifiable to others with a guarantee of fairness and security. Essentially, multiple parties can collectively compute a signature while maintaining security guarantees through these cryptographic constructs, thus leading to reduced communication overhead during consensus by combining signature generation into one threshold signature scheme. Choosing the right consensus mechanism is a matter of individual use case and priority for the blockchain. For public blockchains where achieving maximum decentralization and censorship resistance is the goal, PoW or PoS is often selected regardless of performance limitations. Enterprise or consortium blockchains usually tend to use some BFT-based approach or PoA for higher transaction throughput and clear accountability. As an example, financial applications may prefer mechanisms that give transaction finality without risk of reorganization, or content distribution networks might be content with some eventual consistency to achieve greater throughput. With the maturity of blockchain technology, consensus mechanisms are placing more and more emphasis on sustainability, scalability, and practical utility. Both academic and applied research are ongoing at the intersection of decentralization and other principles such as security and performance, with recent innovations being announced in practice. The number of consensus approaches reflects the change in the perception of blockchains from a one-way-system with cryptocurrency only to a distributed computing framework for various requirements from industries and applications.

1.4 Security and Privacy Aspects of Blockchain

The core design principles of blockchain technology have transformed how we perceive digital security and privacy. Essentially, a blockchain is a tamper-proof, distributed ledger which is utilised in trustless environments by employing cryptographic methods. Unlike traditional centralized systems where trust is vested in a single entity, the blockchain security model is a radical departure which built trust in a distributed manner across a network. At the core of blockchain security is the concept of cryptography. Every transaction is encrypted using asymmetric cryptography, which means the users have the public and private keys. The public key is kind of like an address that other people can use to send you some assets and the private key is kind of a secret that is used to sign off on transactions. This digital signature algorithm confirms the legitimate owner to transfer his/her assets by avoiding double expenditure through the fingerprint algorithms. Blockchain uses cryptographic hash functions (like SHA-256 in Bitcoin) to generate unique fingerprints for data blocks, connecting them in such a way that any alteration is immediately detectable. The nature of blockchain as a distributed system greatly improves its security profile. In contrast to centralized networks that can go down with a single point of failure, a blockchain's ledger is distributed across multiple nodes. This makes it necessary for an attacker to simultaneously compromise a majority of the nodes to successfully change the blockchain — which becomes exponentially harder as the network increases in size. This distributed design really removes single points of failure that have traditionally been an issue in centralized systems, which makes blockchain almost immune to the traditional vectors of attack such as denial-of-service attacks. Another major security feature is the consensus mechanisms used by blockchain networks. These protocols — Proof of Work (PoW), Proof of Stake (PoS), and so on — are what allow nodes across the network to come to an agreement about the validity of transactions and the state of the ledger. Proof of Work (PoW) is the first consensus mechanism introduced by Bitcoin, in which participants (miners) compete to solve complex mathematical challenges, using computational power to validate blocks. This resource need is why it is economically unlikely for any bad actor to try to attack the network as



Notes

they'd have to own greater than 51% of the computing power of the total network. PoS networks, by contrast, choose validators based on the amount of cryptocurrency they own and are willing to “stake” as collateral, providing financial disincentives for dishonesty. The unchangeability of blockchain records is among its most powerful security features. A transaction is considered confirmed once it is in the blockchain, and changing or deleting it then would be quite challenging. Each block includes a cryptographic hash of its predecessor, forming a chain such that altering any given block would require redoing all subsequent blocks — a process that becomes ever more resource-demanding as the chain extends. The immutability of data stored on a blockchain creates a reliable audit trail, which is why blockchain is especially useful for use cases needing tamper-evident storage of records. Even considering these strong security features blockchain systems have some vulnerabilities. Smart contracts – self-executing contracts with the agreed terms of the arrangement directly written into code – have opened up new attack vectors. This risk is exemplified by the infamous DAO hack of 2016, when attackers took advantage of a vulnerability in Ethereum’s smart contract code to funnel millions in cryptocurrency into their own pockets. On the trajectory of blockchain development follows, of course, the need for best practices where code auditing and formal verification mechanisms can be implemented to the code due to horrific security breaches that can occur throughout implementation errors in smart contract code. A similar threat exists in the form of quantum computing. As such, many blockchain cryptographic algorithms, specifically elliptic curve cryptography applied in digital signatures, might be broken by a future quantum attack. This has catalyzed investigations into quantum-resistant cryptographic algorithms that might be able to safeguard blockchains in a world of quantum computing. Then again, the blockchain community is aware of this potential vulnerability and is already addressing the need for quantum resistant solutions before quantum computers become powerful enough to be a genuine threat. 51% Attack 51% attack is still a theoretical risk for blockchains based on Proof of Work consensus. If a single actor, or a group of colluding attackers, controlled in excess of 50% of the mining power of the network, they would be capable of manipulating the blockchain by either double spending coins or preventing some transactions from being

confirmed. Although you'd need a massive amount of computational cycles to attack a large network like bitcoin for the it to be economically viable, smaller blockchain networks can still be attacked. This points to the fact that security in blockchain systems frequently depends on the size and distribution of the network. Another major security issue in blockchain systems is private key management. In contrast to traditional banking systems where lost passwords can simply be reset, loss of private keys in blockchain systems usually results in irreversible loss of access to the stored assets. HR, CS, and mult signature cryptographic approaches have emerged as best practices for key management, but they require users to keep up a higher level of security awareness than in traditional finance. In addition to technical security concerns, governance mechanisms are critical to ensuring the security of a blockchain. Hard forks — when the community splits over alterations to the protocol — can divide the network and potentially diminish its security. Governance Frameworks Balancing Innovation With Stability Are Key To Long-term Security Of Blockchain Systems. For many blockchain projects, the challenge of updating protocols in response to emerging threats while maintaining consensus across the network has prov. The blockchain is a mixed bag from a privacy point of view. This pseudonymous characteristic of most public blockchains provides some level of privacy, with users interacting through cryptographic addresses rather than directly with personal data. Yet this pseudonymity does not mean that they are completely anonymous. Using advanced analysis techniques like transaction graph analysis, researchers can identify behavioral patterns that might tie an address to a person in the physical world. This has resulted in emergence of diverse privacy-preserving technologies in the blockchain ecosystem. One of the most promising privacy technologies available on blockchain is zero-knowledge proofs. These cryptographic techniques enable one party to demonstrate to another that a statement is true, without revealing any information apart from the truthfulness of the statement. For example, zk-SNARKs (Zero-Knowledge Succinct Non-Interactive Arguments of Knowledge) make it possible to have transactions where the sender, receiver, and amount are kept private even if the blockchain itself is still secure. Zcash was the first public blockchain to implement zk-SNARKs, providing users with the option of shielded transactions that offer enhanced privacy. Ring signatures are a more nuanced yet common alternative way to implement blockchain privacy seen



Notes

in privacy coins such as Monero. This technique enables a user to sign a message for a group, such that it is computationally infeasible to identify the key that generated such a signature. By using stealth addresses (used to generate one-time addresses for each transaction) and confidential transactions (which keep the amounts of transactions hidden), so-called ring signatures create an airtight privacy solution that makes tracing financial flows on the blockchain really impractical. Mixers or tumblers are a more centralized solution to the problem of improving blockchain privacy. In summary, such services combine and redistribute the funds of several users, obfuscating the transaction trail and making it difficult to follow the flow of funds. But these services also exhibit a degree of centralization and require users to trust the mixing service not to run away with their funds or keep logs that could later be used to breach their privacy. CoinJoin and other decentralized mixing protocols provide these concerns by separating the deposit & withdraw phase and thus allowing multiple users to deposit funds anonymously into a single transaction without any trusted 3rd party being required.

COLUMN ENCODING

Off-chain solutions like Lightning Network (in short, LN) are designed to complete trading tasks off the main chain that provides more transaction privacy. These payment channels are established and closed on the blockchain but the chain of transactions that can occur throughout while the channels are open remain private and off the blockchain. Not only does this approach enhance privacy, but it also alleviates scalability issues by offloading credit information from the primary blockchain. The intrinsic privacy characteristics of blockchain technology have drawn the investigation of regulatory organizations worrying about possible abuse for illicit activities. This has resulted in the development of regulatory regimes like Know Your Customer (KYC) and Anti-Money Laundering (AML) obligations on cryptocurrency exchanges and other blockchain businesses. They have introduced tension across the blockchain ecosystem, requiring that many entities know your identity and verify it, essentially taking away the privacy gain that brought many to use blockchain to begin with. Privacy vs Compliance — One of the Biggest Hurdles to Mainstream Blockchain Adoption Privacy-enhancing technologies present with challenges as it

needs to meet both user confidentiality and compliance in terms of transparency and accountability. Such solutions include selective disclosure, whereby users can disclose only some (e.g., their age) instead of all (e.g., their name) of their data to trusted parties, leaving others in the dark and thereby striking a balance between these competing data access rights. Another way to improve privacy in blockchain systems is through confidential computing. Confidential computing enables data privacy during processing by executing transactions and smart contracts in secure enclaves secured with hardware-level encryption. Intel's Software Guard Extensions (SGX) and ARM's Trust Zone are examples of technologies that establish a trusted execution environment to keep sensitive data secure from unauthorized observers (and even the operating system admins of the machines executing the processes). As Blockchain technology progresses, marketplace of this tension between transparency and privacy continues to be crucial area of consideration. Public blockchains are based on principles of openness and verifiability, but users are now demanding stronger privacy guarantees. Public systems provide some advantages over the trustlessness of other blockchains, but the control over data visibility limits them. Moving forward, hybrid solutions that integrate aspects of both sectors will likely prove to be the sweet spot for many use cases. However, the security and privacy features of blockchain remain a rapidly evolving field, shaped by papers on new cryptographic techniques, regulatory developments, and the increasing capabilities of potential adversaries. The transformative potential of research in post-quantum cryptography, mixnets for confidential smart contracts, and scalable privacy solutions will impact the safety of blockchain technology in the years to come. So as blockchain systems become more entwined with existing digital infrastructure and financial systems, addressing these security and privacy challenges will be key for unlocking the technology's full potential.



Unit 3: Blockchain Applications and Development

1.5 Various Use Cases – Finance, Supply Chain, Government

The barrel has alluded to by the coming of Blockchain technology has already taken a step beyond the mere ways as the main framework of cryptocurrencies to become a responsibility changing mechanism in all approach. Why this is happening? Its decentralized, immutable, transparent, and secure nature has made it positioned as a solution to long-standing problems in different industries. Some of the most promising and further developed application domains of blockchain are finance, supply chain management, and government services. In finance, blockchain has ushered a revolution that goes well beyond its original use cases in cryptocurrency. Called decentralized finance or DeFi, it is arguably the most ambitious rethinking of financial services. To do so, DeFi platforms use smart contracts to generate programmable regulations and actions, omitting traditional intermediaries such as banks or brokerages. Such platforms allow for lending, borrowing, trading, and investing through protocols that operate autonomously on blockchain networks — predominantly on Ethereum. With access to the internet, DeFi services can be operated by anyone, eliminating the restrictions based on geographical locations or social status that exists in bank-controlled systems, which offers the potential of financial inclusion to the 1.7 billion adults without access to bank accounts worldwide. Cross-border payments are yet another financial sector where blockchain exhibits clear advantages over existing infrastructure. Traditional banking channels to send money internationally include multiple intermediaries, exorbitant fees, and settlement times measured in days. Blockchain-based payment systems can process these transfers in minutes, or even seconds, and for transaction fees that are a fraction of what exists today. Targeting this use case is Ripple's XRP Ledger and Stellar's payment network, which partners with financial institutions to optimize cross-border payment flows. For migrant workers who send remittances back home to their families, these blockchain-based systems could result in decent savings and faster access to funds that people desperately need. The security trading and settlement industry has taken to blockchain and distributed ledger technology to eliminate existing market infrastructural inefficiencies. Traditional securities transactions can take a complicated post-trade process

that can take T+2 (trade date plus two business days) for settlement because of multiple reconciliation processes amongst intermediaries. Block chain platforms can provide near-instantaneous settlement by creating a single and shared record of all transactions that removes the need for reconciliation. Projects to implement blockchain-based clearing and settlement systems have been undertaken by major financial institutions and stock exchanges, such as the Australian Securities Exchange (ASX) and the Swiss Stock Exchange (SIX). These implementations look to diminish counterparty risk, release on capital that is presently occupied during the settlement period and decrease administrative costs. Trade finance, an inherently paper-heavy and fraudulent process, has turned out to be a natural candidate for modernization using blockchain. Facilitating international trade generally requires a handwritten document (such as bills of lading, letters of credit, or insurance certificates) that is verified by multiple participants, usually through manual processes. Blockchain platforms such as IBM's Trade Lens and we. trade provide mutually agreed-upon, tamper-proof archives for these documents, that anyone with permission can access and validate in real time. This minimizes the possibility of document fraud, accelerates the verification process, and adds transparency for all players in the trade ecosystem. Through streamlined processes, blockchain can close the global trade finance gap estimated to be at \$1.5 trillion by the Asian Development Bank — where small and medium-sized enterprises are especially impacted in developing economies. Blockchain implementation is another financial subsector that benefits, with smart contracts automating the process. Smart contracts can streamline claims processing if the pre-defined conditions are satisfied, by eliminating the need for human intermediaries, thereby reducing processing time and administrative costs. Particularly well suited to blockchain are insurance technologies like parametric insurance that pay out upon the occurrence of triggering events, rather than assessed damages. An example of that is flight delay insurance on the blockchain which can automatically issue payments if flight data confirms an eligible delay without a need for filing claims and manual processing. A classic example of consortium models is B3i (Blockchain Insurance Industry Initiative), which unites the biggest insurers to create standardized blockchain solutions for areas such as reinsurance and other complex insurance processes.



Notes

One of the most important institutional responses to innovations based on blockchain are Central Bank Digital Currencies (CBDCs). These are national currencies in digital forms, that are issued and regulated by the central banks, and that take advantage of the blockchain or the distributed ledger technology. China has made significant strides with its Digital Currency Electronic Payment (DCEP) system, and several other countries, such as Sweden, the Bahamas, and the European Central Bank, are at various states of CBDC development. CBDCs are designed to provide the efficiency and programmability of blockchain-based digital currencies while retaining the stability and regulatory oversight of traditional fiat currencies. They could provide governments with better monetary policy transmission and lower cash handling costs, as well as improved tracking of financial flows to curtail illicit activities. Outside of finance, supply chain management remains one of the most attractive application areas for blockchain technology. Global supply chains are growing increasingly complex with many supply chains involving dozens of parties across multiple countries posing major challenges for transparency, traceability, and efficiency. Enter blockchain, providing a shared and immutable record of a product's journey from its origin to consumer. This feature truly shines in sectors where the provenance and ethical sourcing of goods is a primary concern. Food Safety and Traceability: Blockchain Supply Chain Implementation Outbreaks of foodborne illness typically take days or weeks to trace to their source, historically because of fragmented record-keeping across the supply chain. IBM's Food Trust platform, which is used by some of the world's largest retailers, including Walmart and Carrefour, allows food products to be tracked almost instantly from field to store. For instance, when Walmart performed a trace test in 2018 to determine the origin of mangoes, it took them 2.2 seconds to do so with blockchain and almost seven days with the traditional method. This allows for quicker, and more precise recalls, which can greatly limit the effects of contamination events. The pharmaceutical drugs supply chain has also looked to blockchain to address this counterfeiting issue — a global issue that the World Health Organization says afflicts more than 10% of pharmaceutical products in developing countries. The Drug Supply Chain Security Act (DSCSA) requires end-to-end traceability of prescription drugs in the United States, and the ability of blockchain technology to fill this need is being widely

debated. MediLedger is a specialized blockchain network that has been designed specifically for the pharmaceutical industry to verify the authenticity of products at every point of transfer of ownership, thus ensuring an unbroken chain of custody from the manufacturer to the pharmacy. This capability not only assists in ensuring patient safety, but also helps protect the pharmaceutical company's brand and meet increasingly stringent regulatory requirements. Another supply chain use case where blockchain immutability adds tangible value is luxury goods authentication. This could be achieved through blockchain solutions that allow high-end fashion brands such as LVMH to go one step further — enabling consumers to not only entrust their purchases to the brands but also to check the authenticity of their purchases on an immutable record of the item's production and ownership history. Each item undergoes this process is assigned a unique identifier stored on the blockchain, resulting in a digital certificate for each item that cannot be forged. Not only does this assist in tackling the major issue of luxury goods counterfeiting, it also increases the consumer experience by showcasing the craftsmanship and materials used in their purchases. Consumers and regulators alike increasingly demand verification of ethical sourcing and sustainability. The technology aids in creating verifiable records that are a testament to adherence to labor standards, environmental regulations, and fair trade practices. The diamond trade has traditionally been marred by issues of conflict stones, making it a prime candidate for blockchain adoption for this reason. Another example is De Beers' Tracr platform, which uses blockchain technology to follow diamonds from mine to retail, providing a tamper-proof record of each stone's journey—making sure only ethically sourced diamonds ever enter the supply chain. Similar initiatives are in place for minerals used in electronics, cocoa, coffee and other commodities where there are concerns over ethical sourcing.

The fusion of blockchain and traditional trade finance processes has turned the paradigm of supply chain financing upside down. Blockchain helps such financial institutions extend more competitive financing options due to lower risk and increased visibility by creating verified records of shipments, deliveries, and ownership transfers. Small suppliers who have less favorable payment terms with large customers can leverage verified blockchain records of proof of acceptance of the shipment to obtain immediate



Notes

financing instead of waiting for payment for 30, 60, or even 90 days. Marco Polo, a cooperation system between R3, and TradeIX specifically addresses this nexus of supply chain management and trade finance. Blockchain applications in the governmental field have been progressively developing in various public services, tackle improvements in transparency and efficiency and security. Voting systems are one of the most often talked about government applications. With immutable and verifiable voting records, blockchain-based elections could allay some fears about election integrity that arise in anonymous voting mechanisms. Estonia was a pioneering country for digital government services, and it has experimented with blockchain for its i-Voting system; in the U.S., West Virginia piloted a blockchain application for voting for overseas military personnel in the 2018 midterm elections. But challenges persist, especially concerning the accessibility of platforms and the safety of endpoint devices used for voting — blockchain demonstrates immense potential for modernizing electoral systems. Another enticing government blockchain use case is identity management. Existing identity systems are also often fragmented between multiple government departments, hard to access and attractive to fraudsters. An access solution to the above problem is that government services can access digital identity data based on blockchain making it highly secure and on permission where the user can allow certain services to verify their identity. Switzerland's city Zug has launched a blockchain-based digital ID system for its citizens, and Georgia's government records land titles on blockchain, which reduces land disputes, since it facilitates to prove ownership. The applications illustrate how blockchain can fundamentally change the way citizens interact with government services in a more secure, portable digital identity (DID) manner.

Blockchain in Tax Collection and Compliance Tax calculations and payments can now be automated on the basis of the verified transaction data using smart contracts, this may result to a faster process, lower rates of tax evasion and easy compliance for businesses. Public procurement processes in many jurisdictions are susceptible to corruption and inefficiency. In procurement, blockchain can improve transparency through creating immutable records of bids, contract awards, and later changes. Recently, Mexico's state-owned petroleum company Pemex piloted the use of

blockchain for procurement in a bid to reduce corruption within its contracting processes. Blockchain can promote the integrity of the bidding process by storing bid information in a transparent and immutable manner, ensuring contracts are awarded on the basis of merit, not on the basis of back-door influence. Such transparency enables better resource utilization of taxpayer money and increasing public confidence in government purchases. Another governmental area where blockchain can boost efficiency and accountability is grant distribution and tracking. Traditional grant management is a convoluted process with cumbersome reporting requirements, and little clarity on how funds are used in the end. By acting as middlemen, Blockchain platforms can facilitate the application process, automate disbursement of funds based on performance milestones, and maintain transparent records of fund utilization. The European Union is exploring the use of blockchain to track research grants and the United States Agency for International Development (USAID) piloted blockchain to track foreign aid. These implementations are intended to enhance administrative efficiency while increasing accountability in public money utilization. Healthcare records management links governmental and private interests, especially in countries with national healthcare systems. How easy is it for you to change providers, and what challenges does anyone face in terms of getting a complete picture of all sides of your health? Blockchain health record systems can enable patients to own their medical data, whilst providing on-demand access to a non-modifiable complete pool of records for genuine and authorized healthcare providers. Estonia's e-Health system uses a blockchain-based approach to secure patient records, while Singapore's Ministry of Health has investigated the application of blockchain-based technologies for patient consent management, as well as secure data sharing across healthcare providers. Such implementations should weigh privacy considerations against the potential advantages of data accessibility for enhanced care coordination and research. While the three sectors above are widely considered the most important ones for blockchain application, there are still many “niche” use cases in other industries. In energy, blockchain peer-to-peer trading platforms allow homeowners with solar panels to sell their excess electricity directly to neighbors. Blockchain technology can deliver verifiable, tamper-proof academic credentials and reduce fraud while making it easier for employers to verify credentials. Blockchain technology has vast potential applications

across various industries, including real estate, where it can simplify property transactions and create immutable records of ownership. In media and entertainment, blockchain facilitates new models of content monetization that allow creators to have direct interactions with their audiences, circumventing traditional intermediaries. The continued evolution of blockchain use cases across financial services, supply chain, government, and other sectors will likely be influenced by a number of trends — regulatory developments that create clarity without stifling innovation; technology advancements that better serve scalability and performance requirements; governance models that steer toward a balanced scorecard between decentralization and pragmatism; and integration strategies that will help blockchain fit as an enabler rather than a replacement for working systems. With these and other developments well underway, blockchain technology stands to deliver on its promise as a transformative leap across many areas of human endeavor.

1.6 Hyperledger Fabric – A Platform for Blockchain Development

Hyperledger Fabric: The Most Prominent Enterprise Blockchain Technology Unveiled in 2015, under the auspices of the Hyperledger Project – an open-source collaborative project hosted by the Linux Foundation – Fabric emerges as a radical departure from experimental, early-stage distributed ledger implementations, focusing on the explicit needs of enterprise environments: privacy, scalability, performance, and fine-grained access control.

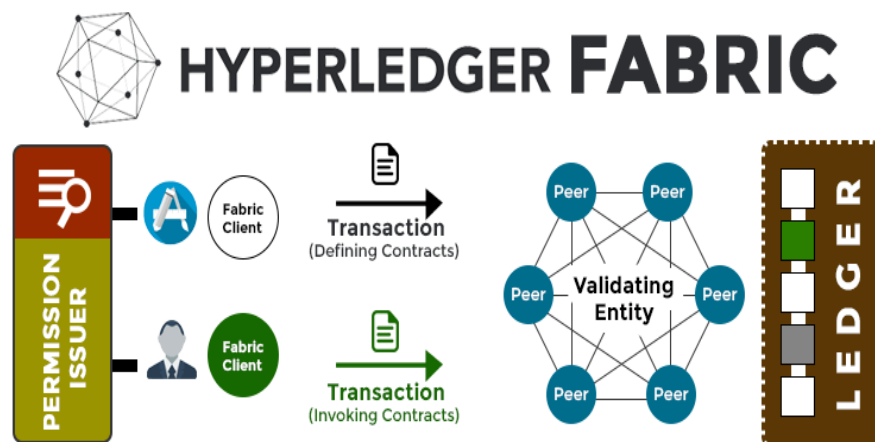


Figure 4: Hyperledger Fabric

Originally released in 2017, Hyperledger Fabric has quickly become one of the most popular platforms for building and executing blockchain applications in a wide range of sectors, including financial services, health care, supply chain management, and government services. The design principles behind Hyperledger Fabric demonstrate types of block chain architecture to be largely rethought for enterprise settings. In contrast to public blockchains, such as Bitcoin and Ethereum, which function on permission-less networks where anyone can join, Fabric maintains a permissioned model whereby participants are identified by known identities (e.g., organization entities). This design choice solves important enterprise needs around data privacy and regulatory compliance. It is the modular architecture of the Fabric that helps it to stand out; with consensus mechanisms and membership services designed as plug-and-play components helps organizations customize blockchain networks for their distinct requirements. This modular approach widely contrasts with the monolithic design that many blockchain platforms follow and has been especially beneficial for enterprises with complex and specific requirements. Hyperledger Fabric's execute-order-validate architecture is also unique and differs from the order-execute model of many blockchain systems at the infrastructure level. In classical blockchains, transactions are executed sequentially by all nodes but ordered first (through mining or other consensus mechanisms). This causes performance bottlenecks and constrains throughput. Fabric's execute-order-validate (EOV) paradigm alters this flow: first, transactions are executed in parallel, in a dedicated set of nodes known as endorsers, then they are ordered through a consensus service, and finally validated before being committed to the ledger. This architecture allows them to successfully hit scale without compromising transaction throughput or slowing it down with bottlenecks as it allows the processing to be done in parallel, which solves the fundamental performance limitations that have stymied blockchain adoption in high-volume enterprise applications. Identity management is integral to the Hyperledger Fabric security model. Fabric uses a complex PKI (Public Key Infrastructure) that manages digital certificates of all network participants through its MSP (Membership Service Provider) component. All that together gives rise to an infrastructure that allows for attribute-based access



Notes

control (ABAC), which means permissions can be granted based on any number of attributes on the user (rather than just their identity) and their organizational role. The MSP can support multiple certificate authorities, thus enabling different organizations to continue to maintain control over their own identity management while still being part of a single shared network. The requirement of C, which gives huge flexibility of using users of different kinds with no interdependent pressure on the outside remains very precious in consortium blockchains, as multiple organizations need to collaborate with each other while managing their existing identity management systems and organizational hierarchies.

Hyperledger Fabric's channel architecture provides a means for strong data segregation and confidentiality. Channels act as individual ledgers that are shared among selected sets of network participants, where private messages and transactions are conducted privately and are not 'visible' to other parties on the network. This element specifically targets the privacy issues that have hindered widespread adoption of blockchains in industries like financial services and healthcare, where confidentiality is of utmost importance. Firms can create channels for various business relations, building essentially "virtual blockchains" making up the overarching chain. This architectural model allows for advanced multi-party workflows, where some information can be broadcast to a large audience, whereas other data can be processed by only certain stakeholders, replicating the diverse information sharing requirements for real-world business relationships. In Hyperledger Fabric, smart contracts are referred to as chaincode, where chaincode is programmatic code that governs the business logic of assets on the blockchain. Fabric supports various programming languages for chaincode, such as Go, Node.js, and Java, making it easier for developers to hop onboard them with familiar languages instead of the specialized blockchain-specific languages that so many others require. In Fabric, chaincode is deployed in a secure container environment that ensures isolation and prevents any malicious code from impacting the entire network. This technique secures the business logic while supporting the independent deployment and upgrading of business logic without interfering with the network. Chaincode can be written in mainstream programming languages, which has greatly expedited development cycles for blockchain applications

while increasing the number of developers who can build on the platform. The consensus view in Hyperledger Fabric is aligned with the modular nature of the architecture as it decouples transaction validation from ordering and offers pluggable consensus implementations. This decoupling enables network architects to choose consensus mechanisms suitable for their unique trust models and performance needs. Fabric's default ordering service, which works on a particular version of Apache Kafka or uses a variant of Raft, only provides crash fault tolerance as opposed to Byzantine fault tolerance, and thus optimizes for performance in environments where the identities of the participants are known. To provide stronger byzantine fault tolerance where needed, Fabric's modular architecture allows alternative consensus implementations to be plugged in. Since specific consensus middleware is abstracted from the application logic, Fabric networks can upgrade consensus mechanisms as needs change, or as new, better algorithms surface, without having to redesign the application. Hyperledger Fabric ledger structure consists of two separate components: blockchain and world state. A blockchain is an append-only data construct guaranteeing an immutable history for data, in this case, transactions, stored as a linked sequence of blocks. Implementing a world state as a database (for example LevelDB or CouchDB), it stores the current value of all the assets, allowing very efficient queries without having to go through the entire transaction history. CouchDB integration is especially notable as it enables deep queries over JSON data, giving way to complex data model and advanced query functionality that are not found in typical blockchain systems. The dual-component nature of this ledger allows developers to balance common blockchain features, such as tamper-proofing and immutability, with the performance needs of enterprise applications, wherein query performance is often different than the performance typically associated with the blockchain.

In Hyperledger Fabric, private data collections is an additional privacy method which allows a subset of organizations in a channel to endorse, commit and query for private data without the need to have separate channels. This feature helps cater to complex privacy requirements in scenarios where a subset of channel participants needs to keep certain data secret from other participants of the channel. The local database of each peer stores the actual private data and only a hash of it is committed to the



Notes

ledger, ensuring privacy but anchoring the private content to the shared ledger for integrity. The mechanism enables use cases including price negotiations, where competing providers may need to share confidential pricing data, but not with each other with the buyer – a requirement that is common in multi-party business processes, and difficult to address on previous blockchain platforms. Network governance in Hyperledger Fabric is based on policies that govern all aspects of network operation, from chaincode instantiation to changes to the channels configuration. Such policies typically need to be signed by several different entities — a sort of distributed governance model that is correct for consortium blockchains, where no single organization should be allowed to exercise unilateral control. From a simple majority to more complicated weighted voting based on responsibilities, Fabric's governance model is a policy framework that allows for such flexibility in how to express rules. This governance structure offers the formal mechanisms necessary for enterprise blockchain networks to change and grow over time without losing the consensus of participants, one of the major challenges preventing the long-lived continuation of enterprise blockchain networks across organizational borders. Fabric's transaction flow from execute order to validate is just one of many architectural innovations to their approach to scalability. This allows for the partitioning of the network, enabling different groups of participants to process transactions in parallel, and you can create multiple channels. Dependency on only a subset of peers for validation in endorsement policies can ease computational overheads while processing transactions. Endorsing, committing, and ordering peers are separate peer types, and this allows organizations to provide resources according to the roles that they will perform in the network. In summary, these scalability features work together to allow Hyperledger Fabric to handle the transaction volumes needed for enterprise applications, especially in areas such as supply chain management where a large numbers of organizations may engage through a common network. Hyperledger Fabric networks are generally deployed using container technologies especially Docker and Kubernetes for container orchestration of the componenet. This containerization approach allows for easier deployment across heterogeneously infrastructure environments and enables on-demand scaling of network resources based on workload changes. Fabric with its DevOps tooling compatibility and practices has

made things simpler for enterprises to convert to enterprise IT environments while providing network management through existing operational procedures and tools. Deploying Fabric components across multiple cloud providers or on-premises infrastructure has been especially useful for consortium blockchains where the participants might have different preferences and requirements for the infrastructure. The ability to integrate with existing tools is critical in the growth of enterprise blockchain, as these networks need to plug into existing apps to provide business value. Also, Hyperledger Fabric supports multiple ways of integration, for example: event-driven architecture to trigger things based on changes in blockchain state, or direct API calls. Multi-Cloud Support: Multi-Container Deployment Model

This includes any combination of virtualized and cloud-based container environments, with the intention of being able to deploy the same application at all sites. But Fabric also supports external chaincodes, which provides a method of integrating with existing legacy systems if being rewritten as on-chain logic is unrealistic. These integrations allow Fabric to become an orchestration level across organizations and use enterprise systems instead of forcing their replacement.

Hyperledger Fabric has been around a long time and performance characteristics have improved with each release – recently documented metrics show over 20,000 transactions per second throughput with the right set of configurations. Performance tuning in Fabric is an art, where a single repository has many ways to allocate hardware, configure the network, design the chaincode, and select the endorsement policy. This ability to horizontally scale on the platform when adding channels and the ability to scale vertically on hardware gives multiple opportunities if performance requirements are not being met. Fabric has performance characteristics that have been adequate for many of the enterprise use cases it has been applied to, including high-volume supply chain tracking and financial transaction processing, but as is typical for a general-purpose platform, fine-tuning for specific workloads requires careful configuration and testing. Hyperledger Fabric has proved its versatility in real-world uses across many industries. In finance, HSBC and Société Générale have used Fabric to build underlying networks for trade finance platforms that digitise the documentary aspects of international trade. In supply chain, Walmart has used Fabric for food traceability allowing them to track produce from farm



Notes

to store. For example, in healthcare, organizations have leveraged the use of Fabric to build secure, permissioned networks for the secure collaboration of shared patient data, in a way that is compliant with regulations. Government use cases of Fabric include land registry systems, customs documentation, and contractor management. These different implementations demonstrate Fabric's flexibility to suit multiple business domains and regulatory environments.

Hyperledger Fabric: The developer experience has come a long way since the first release. None of this was too easy to operate — early implementations had complex setup steps and a profound knowledge of the blockchain, the steep learning curve. Recent releases have added tools for certificate management such as Fabric CA and have made setting up a network easier because of the improved documentation. In addition to this, the Hyperledger Fabric Operations Service was introduced, which simplifies operational concerns with regard to running a Fabric network. Development frameworks such as Hyperledger Composer (now an archived project) treated the establishment of patterns for modeling business networks, the influence of which still shapes our thinking on overall Fabric application design. The new Fabric Smart Client framework aims to reduce the effort of creating even more complicated applications involving many peers. Simply put, this means that teams building on top of Fabric face much less of a learning curve than ever before, and that the tools, framework, and capabilities that have been introduced have combined to reduce the level of specialized knowledge involved in building on the Fabric platform. Some security aspects of Hyperledger Fabric relies on cryptographic principles, and some require good operational security practices when deploying the enterprise solution. Private keys used by network participants can be protected by integrating Hardware Security Modules (HSMs). The platform supports Transport Layer Security (TLS) for network communications, protecting against eavesdropping and man-in-the-middle attacks. Fabric has a defense-in-depth security model that provides multiple layers of protection from the network to the application level. Built-in security auditing capabilities, all transactions are cryptographically signed and immutably recorded, forming a comprehensive auditable trail for compliance. Such security features have allowed Fabric deployments in highly regulated

industries, where protecting data and auditability are of utmost importance.

Interoperability: An emerging area of focus for Hyperledger Fabric development is interoperability with other blockchain networks. Exploring mechanisms for communication between Fabric networks and other blockchain systems (Hyperledger or otherwise) is one of the focus areas of Hyperledger's Interoperability Working Group. Atomic cross-chain transactions, notarization schemes and relay mechanisms that provide connectivity between heterogeneous blockchain architectures are technical approaches for cross-chain solutions. They leverage standards such as the Token Taxonomy Framework to ensure a consistent representation of digital assets across the ecosystem. These initiatives also highlight the idea that the future blockchain landscape is not simply a winner-takes-all contest, as the best approach will likely involve multiple blockchain networks across diverse sectors and the established protocols that allow them to interoperate. The Hyperledger Fabric project itself is governed by open-source principles, and development direction is decided through collaborative processes involving a large number of different stakeholders. The big names involved include tech companies such as IBM and Oracle, banks, supply chain companies, and independent developers. Hyperledger Technical Steering Committee maintains an open process for technical decision making through a system of code contribution, peer reviews and an ongoing commitment to quality. 3. This open governance model has been critical to the adoption of Fabric across industries, providing assurance that the platform is not controlled by any single commercial entity, and that the future of Fabric will be defined by the community of users, not a single vendor.

As we step into future, Hyperledger Fabric is continually evolving, with new releases having the potential to address emerging needs of the enterprise users. These features provide support for external chaincode services, which allow smart contract logic to run outside the peer environment, enhanced private data capabilities and performance optimizations that increase throughput and reduce latency. Improvements to the programming model, streamlined management of networks, and increased interoperability capabilities are also penciled in for the project roadmap. As enterprise blockchain adoption matures from experimental projects to the production systems that many organizations rely on today, Fabric is also evolving to



Notes

address the operational challenges of maintaining long-lived blockchain networks in complex organizational environments. Beyond the core network platform offered by Hyperledger Fabric, its ecosystem also includes a group of complementary tools and frameworks that are targeted to help cover specific use cases in the lifecycle of a blockchain solution. Hyperledger Cactus is focused on blockchain interoperability, and making it easier to connect Fabric networks to other blockchain systems. Hyperledger Caliper is a sandbox environment that provides the benchmarks to measure and compare the performance of various blockchain implementations and configurations. Hyperledger Aries is the framework for decentralized identity management, which can be added on top of Fabric's native identity capabilities. This broader ecosystem will help solve issues that go beyond just basic blockchain functionality and enable the creation of integrated enterprise solutions instead of siloed BAT components. 3:23Despite frequent comparison with other public chains, Hyperledger Fabric is a milestone in the development of enterprise blockchain based on the open-source platform, which provides a flexible connection and modular use model for business applications. Its permissioned design, advanced privacy features, and performance-oriented architecture overcome many of the barriers that have previously stood in the way of enterprise level adoption of blockchain. As the technology matures and the ecosystem of tools and frameworks surrounding it expands, Hyperledger Fabric will drive new forms of cross-organizational collaboration across multiple industry sectors. It is the practical insights gained from Fabric's ongoing evolution and experience in the enterprise market that inform much of the current understanding of how distributed ledger technology can be fundamentally relevant to rethinking business, and how it intersects with the underlying demands of enterprise systems.

Multiple Choice Questions (MCQs)

1. **What is the primary feature of blockchain technology?**
 - a) Centralized control
 - b) Distributed ledger
 - c) Manual data processing
 - d) Single-point failure

2. **Which component of blockchain ensures that all transactions are transparent and immutable?**
 - a) Smart Contracts
 - b) Consensus Protocol
 - c) Cryptographic Hashing
 - d) Centralized Database
3. **Which of the following is NOT a consensus protocol used in blockchain?**
 - a) Proof of Work (PoW)
 - b) Proof of Stake (PoS)
 - c) Proof of Authority (PoA)
 - d) Proof of Trust (PoT)
4. **Which of the following industries has benefited from blockchain technology?**
 - a) Finance
 - b) Supply Chain
 - c) Government
 - d) All of the above
5. **What is Hyperledger Fabric primarily used for?**
 - a) Public cryptocurrency transactions
 - b) Enterprise blockchain solutions
 - c) AI-powered chatbots
 - d) Internet of Things (IoT) applications
6. **Which feature of blockchain prevents unauthorized modifications to past transactions?**
 - a) Hashing and cryptography
 - b) Centralized authority
 - c) Editable ledgers
 - d) Limited scalability
7. **What is a smart contract?**
 - a) A self-executing contract with pre-defined rules
 - b) A legal agreement written on paper
 - c) A document stored on the cloud
 - d) A physical contract signed using digital signatures



Notes

8. **Which consensus algorithm is used by Bitcoin?**
 - a) Proof of Stake (PoS)
 - b) Proof of Work (PoW)
 - c) Delegated Proof of Stake (DPoS)
 - d) Byzantine Fault Tolerance (BFT)
9. **What makes blockchain a decentralized system?**
 - a) The use of a single central authority
 - b) Multiple nodes maintaining a shared ledger
 - c) Its reliance on government regulations
 - d) The ability to edit past transactions
10. **Which of the following best describes blockchain's security feature?**
 - a) It allows unlimited access to stored data
 - b) It uses encryption and hashing to ensure data integrity
 - c) It operates on a single-point failure system
 - d) It has no privacy protection mechanisms

Short Answer Questions

1. What is a blockchain ledger, and how does it work?
2. Explain the difference between public and private blockchains.
3. What is cryptographic hashing, and why is it important in blockchain?
4. Name three common consensus protocols and briefly describe their purpose.
5. How does blockchain ensure data immutability?
6. What are smart contracts, and where are they used?
7. What role does Hyperledger Fabric play in enterprise blockchain solutions?
8. Explain the Byzantine Generals Problem in blockchain consensus.
9. What are some key security risks in blockchain technology?
10. How can blockchain improve supply chain management?

Long Answer Questions

1. Explain the fundamental structure of blockchain. How do blocks, transactions, and consensus mechanisms work together?
2. Compare and contrast Proof of Work (PoW) and Proof of Stake (PoS) consensus protocols. Which one is more efficient?
3. Discuss the challenges and limitations of blockchain technology.
4. How does blockchain enhance security and privacy in digital transactions?
5. Describe how blockchain is transforming the financial sector, with real-world examples.
6. Explain the architecture and key components of Hyperledger Fabric.
7. How does blockchain technology prevent fraud and unauthorized access?
8. What are the ethical considerations of implementing blockchain in government and public sectors?
9. Describe a real-world blockchain use case in the supply chain industry and how it improves transparency.
10. What are the future trends in blockchain technology, and how will they impact different industries?



MODULE 2

BASIC CRYPTOGRAPHIC PRIMITIVES BEHIND THE BLOCKCHAIN

LEARNING OFUTCOMES

By the end of this Module, students will be able to:

1. Understand hash functions and their properties: collision-free, hiding, and puzzle-friendliness.
2. Explain the role of hash pointers in blockchain security.
3. Understand digital signatures and their importance in verifying authenticity.
4. Explore public key cryptography and its applications in blockchain.
5. Analyze public key encryption and how it secures data transmission.
6. Understand and analyze the RSA algorithm, including its working principles.
7. Differentiate between public and private keys in RSA encryption.

Unit 4: Hash Functions and Cryptographic Foundations

2.1 Hash Function – Collision Free, Hiding and Puzzle Friendly

(A hash function is a basic cryptographic primitive that takes data of any size and produces a fixed-sized output, thus allowing indexing of arbitrary-sized data with a fixed size hash.) These math functions are used in a multitude of crypto applications, from digital signatures to blockchains. A hash function is cryptographically secure if it has three core properties: collision resistance, hiding, and puzzle-friendliness.

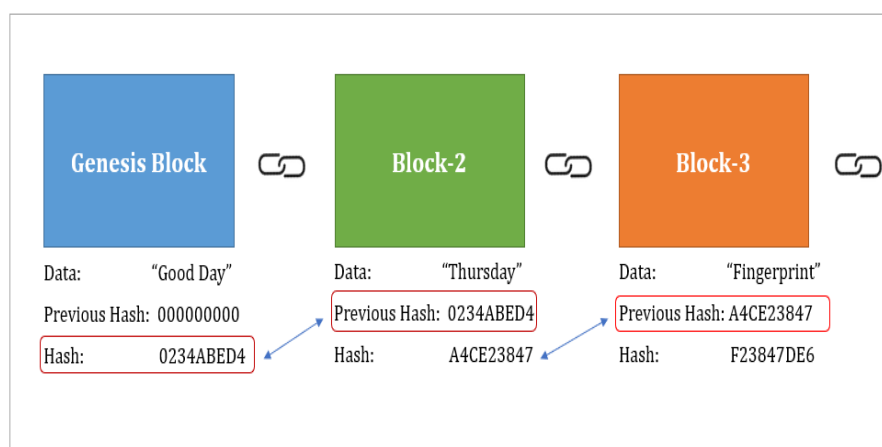


Figure 5: Hash Function in Blockchain

Collision Resistance

Let H be a hash function, we have that H is called collision resistant if it is computationally hard to find two different inputs x, y such that $H(x) = H(y)$. Put simply, though collisions (by the pigeonhole principle, when mapping from larger input space to smaller output space) exist theoretically, finding such collisions should be practically impossible. For example, modern cryptographic hash functions (e.g., SHA-256) have an output size that is large enough (256 bit) that brute-force collision search is infeasible. As such, an approximate 2^{128} operations (by the so-called "birthday paradox", given a 256-bit output) would be necessary to find a collision in practice, an effort that outmatches even the most cutting-edge algorithms by several orders of magnitude. Collision resistance is critical for digital signatures, which would be compromised by an attacker capable of finding collisions.



Notes

In the context of blockchain technology, collision resistance ensures that the hash of a transaction uniquely identifies that transaction, making it very unlikely that a double-spending attack will work. Certain hash functions once widely accepted as collision-resistant have also succumbed to attacks, despite their theoretical hardness. MD5 and SHA-1 are examples of widely used algorithms that were found to be insecure when researchers succeeded in demonstrating practical collision attacks against the algorithms. This underlines the need for continued cryptographic research and transition to more robust hash functions as new vulnerabilities are found

Hiding Property

Hiding property: This is used to prevent an attacker from knowing the original data from the hash value, based on which a hash value is said to be hiding if, given $H(x)$, the original x used to generate the hash is infeasible to be calculated. Informally, for a hash function H and input x sampled from a high min-entropy distribution (i.e.: x is sufficiently random), it should be hard to compute x given $H(x)$. However, salting or other measures to maintain the hiding property should be applied for well-known inputs (for instance common passwords). In blockchain applications, this hiding property enables commitment schemes in which the user commits to a value yet keeps it hidden until revealed. As a very simple example, in the case of a sealed-bid auction, each participant could submit $H(\text{bid} \parallel \text{random_nonce})$ to commit to their bid without revealing the actual amount until the end of the auction. The property of hiding depends on the one-wayness of cryptographic hash functions. Although hash functions are deterministic (a given input will always yield the same output), the transformation function effectively shuffles the input data, rendering reverse-engineering virtually infeasible without prior knowledge of the input.

Puzzle-Friendly Property

A hash function H is said to be puzzle-friendly if for any output value y having high enough min-entropy, it is computationally hard to find input x such that $H(x) = y$ much faster than simply trying random inputs. This property is especially important for proof-of-work systems, such as the ones used in many cryptocurrencies. For Bitcoin, for instance, miners try to find

a block header whose hash is less than a target value. The puzzle-friendly property also guarantees that no single strategy dominates over pure random guessing, again a nice property from the point of view of keeping the playing field level in terms of the computational resources available to those miners. More formally, if H is a puzzle-friendly hash function, and the target value y is chosen from some high min-entropy distribution, and a hint z is provided that is independent of the target (i.e., the target can't be solved faster than brute force), then it will still be computationally infeasible to find x such that $H(x \parallel z) = y$, therefore, no shortcuts can be taken, and solving hash puzzles will require performing the prescribed computational work. The property is also needed to fairly distribute computational tasks and to prove that certain computational effort has been spent. It allows for the design of systems in which members are required to prove some amount of computational effort before being granted benefits, forming the basis backbone of many consensus mechanisms used in distributed systems.

Implementation and Standards

These properties are achieved through Modern cryptographic hash functions that make use of complex algorithms requiring multiple rounds of bit manipulation, substitution and permutation operations. Some of the standard hash functions which are very popular are: SHA-2 family (SHA-256, SHA-512...) Developed by the NSA, published by NIST still secure, used in SSL, TLS, Bitcoin, whole government systems. SHA-3 family: SHA-3 use Keccak algorithm and is part of the first public competition in selecting SHA hash family, SHA-3 internal structure is different from SHA-2, which brings security diversification. BLAKE2 and BLAKE3 — Designed to maximize performance on today's CPUs while providing high security margins, these functions are increasingly common in performance-sensitive applications.

Note that all of these hash functions achieve the same security guarantees at a high level, even though each implements the collision-resistant, hiding, and puzzle-friendly properties through different internal mechanisms.

You are admitted to the world of modern cryptography.



Notes

These three properties of cryptographic hash functions have many applications outside the world of blockchain technology:

1. **Password storage:** The hide property lets systems store hashes of passwords, rather than in plain text, so that user passwords are not disclosed if a database is hacked.
2. **Data integrity verification:** Hash functions produce checksums capable of identifying any alteration in the original data.
3. **Pseudorandom Number Generation:** Hash functions can turn deterministic input sequences into output sequences with statistical properties that are similar (but not identical) to random numbers.
4. **Commitment schemes:** The hiding and binding properties of hash functions let parties commit to a value but not reveal it right away.
5. **PoW systems:** The puzzle-friendly property forms; major computational challenges that are difficult to solve and easy to verify.

Each application benefits from the unique features of hash functions that offers various security assurances. It's important to understand these properties so that you can implement cryptographic protocols correctly and evaluate their security properties..

2.2 Hash Pointer

Its most well-known act is the use of Hash pointers which were combined to create the essence of blockchain itself. A hash pointer is like a normal pointer but with the data it's pointing to, and the cryptographic hash of its contents. Not only can data be located, but it can also performed as provided able to verify its integrity = this dual nature.

Basic Structure and Function

Pointers in with reference to traditional data structures simply point to a memory location, more like to an address where a data type or a data is being stored. This concept is improved by having a hash pointer that contains a cryptographic hash of the referenced data. This adds an incredible layer of verification: when the data changes, it changes the hash, thus ensuring that tampering is evident.

In more formal terms, a hash pointer to a data block X is composed of:

- A reference to the place where we stored X
- $H(X)$ — A cryptographic hash of X 's content

By using a hash pointer to access data, one can recompute the data's hash and verify it with the hash value stored. Any mismatch indicates that the data has been changed after the hash pointer was created.

Post-WI Data Structures (Tamper Evident)

The main benefit of hash pointers is that they enable tamper-evident data structures. We replace regular pointers in the data structure with hash pointers, and thus traditional data structures become tamper-evident. Some key examples include:

Hash Pointer Linked Lists

Thus the way a normal linked list maintains data and pointers only to the next node. In a hash pointer linked list (aka blockchain), each node contains data and a hash pointer to the previous node. This establishes a chain in which each node cryptographically validates the integrity of its predecessors. Suppose an attacker modifies some data in any node, now in order to keep the integrity intact they will also have to modify this hash in all next nodes. This cascading requirement propagates all the way down the chain and finally necessitating the change to the latest hash pointer (often referred to as the “head” or “tip”). In that case, any tampering with the chain can be detected if a trusted party securely stores this head hash pointer.

Merkle Trees

Merkle trees (or hash trees) leverage hash pointers to form a binary tree structure allowing for efficient verification of the integrity of large data sets. In a Merkle tree:

- Hashes of the individual data blocks are at the leaf nodes
- Hash pointers to leaf nodes of non-leaf nodes
- The top node (Merkle root) has a hash that validates indirectly all the data of the tree



Notes

Having this umbrella structure helps take severe efficiency advantages. To ensure a data block is part of the tree, only a logarithmic number of hashes (in respect to all data blocks) needs to be computed — that's a so-called Merkle proof. Therefore, given a dataset, a Merkle Tree (MT) is a data structure that derives from the DAG in which the leaves are the data, which allows you to prove that a dataset is correct without having a single entry in hand and falling into the domain of what is known as "light node". One such use case of Merkle trees in blockchain is the Merkle root who commits for all transactions of the block in Bitcoin, which helps lightweight clients to verify individual transactions without needing to download the complete blockchain.

The fundamental building block of blockchain architectures is hash pointers, and they provide the key security properties of blockchains:

Immutable Transaction History

In blockchain systems, each block is linked to a previous one by a hash pointer, creating a chain of blocks — changing any historical transaction would require recalculating all future blocks. Because creating blocks in these systems generally requires substantial computational effort (in proof-of-work systems) or cryptographic signatures from trusted parties (in permissioned systems), this property makes historical records practically immutable after a sufficient number of subsequent blocks have been appended.

Efficient Verification

The hash pointers, especially organized in Merkle trees, enable the verification of specific pieces of data without the need to process the full dataset. This property allows lightweight blockchain clients (SPV clients) to validate their own relevant transactions without needing to store the entire blockchain.

Timestamping and Ordering

This gives hash pointers an immutable chronological ordering of events. By including a reference to the previous block, each block unambiguously

timestamps all of its contents, which creates a historical record that cannot be backdated or rearranged.

Advanced Applications

In addition to the main use in blockchain systems, hash pointers facilitate several advanced cryptographic applications:

Skip Lists

Skip lists with additional hash pointers can be used to form authenticating data structures with tamper-evidence and searchable properties. These data structures can be leveraged for authenticated dictionaries, certificate transparency logs, etc.

Sparse Merkle Trees

Sparse Merkle trees use hash pointers to produce key-value mappings representations where sparse majority keys are unmapped. The structures allow compact non-membership proofs that can be useful for certificate revocation lists or state verification in certain blockchain architectures.

Incremental Merkle Trees

Incremental Merkle trees that support efficient, progressive updates and canonical views will be beneficial in scenarios where data is pumped in, especially data-heavy application like blockchain or zero-knowledge proof systems.

Implementation Considerations

There are some practical considerations you need to think about when using systems with hash pointers:

Hash Function Selection

Hash pointers are as secure as the underlying hash function. It should be collision-resistant so that attackers cannot find different data that produces same hash. Most current applications use hash functions such as SHA-256, SHA-3, or BLAKE2/3, which offer strong collision resistance.



Notes

Pointer Encoding

Hash pointers must be encoded in a standardised fashion so as to be usable in distributed systems. That encoding has to describe not only where to find the referenced data (which might be a network protocol or database identifier rather than a direct memory address) but also how to encode the cryptographic hash itself.

Storage Efficiency

The hash pointers entail storage overhead, as every pointer consists of a cryptographic hash (for example, 32 bytes for SHA-256). In large systems of millions or billions of pointers, this overhead amounts to something. In contexts where full collisions resistance is not needed, there are different optimisation techniques, e.g. the use of truncated hashes.

Concurrency and Atomicity

If you have a data structure with hash pointers, you can't simply update the data and its pointers separately; you have to update the data and the hash atomically. This requirement poses challenges in concurrent systems and frequently demands locking solutions or additional concurrency management methods.

Theoretical Foundations

Hash pointers provide security guarantees that are based on several theoretical grounds:

Collision Resistance

The security of hash pointers relies on the collision resistance of the underlying hash function. For example, if an attacker could find two different data blocks X and X' such that $H(X) = H(X')$, they could then substitute X' for X without anyone noticing. Today's cryptographic hash functions such as SHA-256 provide collision resistance at about 2^{128} operations, which means that no machine or algorithm available (or even

likely to be invented) can produce two different messages with the same hash.

Binding Commitment

Hash pointers are a mechanism for implementing cryptographic commitments, which commit the creator to the exact data that they refer to. This bindinality property inhibits equivocation — the property of being able to commit to an action that could later be shown to be different values.

Computational Integrity

Hash pointer-based systems offer us computational integrity—every calculation or state transition can be proved to obey the rules of the system. This property is useful for trustless distributed systems where participants do not naturally trust each other.

Introduction to Hash Pointers in Advanced Cryptographic Protocols

Hash pointers are a concept that goes beyond simplistic data structures and into more intricate cryptographic protocols:

Zero-Knowledge Proofs

Hash pointers are employed in some zero-knowledge proof systems to produce compact commitments to big data sets. As an illustrative example: recursive SNARKs can chain proofs (through hash pointers), yielding verification of more complex computations with constant-sized proofs.



Figure 6: Zero Knowledge Proof

Time-Lock Puzzles

The hash pointers, along with the sequential nature of the puzzle solution, can be turned into time-lock mechanisms, where information can be encrypted in a manner that requires some predictable amount of sequential computing in order to decrypt, regardless of parallel computing resources.

Multi-Party Computation

They provide the building blocks for secure multi-party computation protocols,[18]*(Module 3) and are necessary for creating verifiable secret sharing schemes and enabling honest (but curious) participants to follow the protocol correctly without leaking their private inputs.

Future Directions

The notion of hash pointers is also evolving, and we postulate the following new research directions:

Quantum Resistance

Grover's algorithm gives a quadratic speedup when it comes to finding preimages for a hash function, so as quantum computers remain to develop, existing hash functions are at risk. Efforts are underway to investigate post-

quantum secure hash functions and larger hash sizes to ensure the security of hash pointers in the era of quantum computing.

Verifiable Delay Functions

Verifiable Delay Functions (VDFs) generalize hash pointers with time-sensitivity, leading to structures that provably can't be computed in parallel. Such functions are being explored for use in consensus algorithms and randomness beacons.

Stateless Verification

Then, advanced studies determine how hash pointers allow stateless verification, since validators can validate the rightness of the system without keeping the complete system state. This tendency has a far-reaching impact on the scale consideration of distributed systems.

Authenticated Encryption with Hash Pointers

The combination of hash pointers and authenticated encryption leads to systems in which confidentiality and integrity are ensured, along with efficient mechanisms for selective disclosure.

The hash pointers are a powerful combination of two of the most fundamental ideas in computer science and cryptography — referring to data and verifying the integrity of data. The emergence of hash pointers, combining the directional capabilities of pointers with the verificational properties of cryptographic hashes, enabled the construction of tamper-evident data structures that constitute the basis for blockchain technology and various secondary applications of cryptography. Hash pointers are different in that they are simple and composable. It is upon this simple primitive that complex constructions of Merkle trees and blockchains are built, allowing for trustless verification across distributed systems. Cryptographic research is an ongoing work that contributes to new applications for hash pointers, such as privacy-preserving protocols, scalable verification systems, and quantum-resistant security architectures. Like many emerging cryptographic constructions, the comprehensive study of hash pointers follows both on advances in theory and implementations;



Notes

thus, knowing their properties, implementations, and theory can help us better understand the security guarantees, their design tradeoffs, and the networks that surround modern systems. Additionally, as blockchain technology and alternative computing paradigms progress, the concept of hash pointers will continue to be central to developing verifiable, tamper-evident data structures in trustless environments. Hash pointers are indeed a building block within the cryptographer's toolbox of tools, which facilitate secure interactions in a distributed systems.

Unit 5: Digital Signatures and Cryptographic Techniques

2.3 Digital Signature

Extending these properties to the realms of electronic communications and transactions, digital signatures stand as one of the most significant applications of modern cryptographic systems. Digital signatures, on the other hand, are mathematically derived and intrinsically tied to the message that they sign, allowing for a significantly higher level of security and versatility than paper-based handwritten signatures in the digital world.

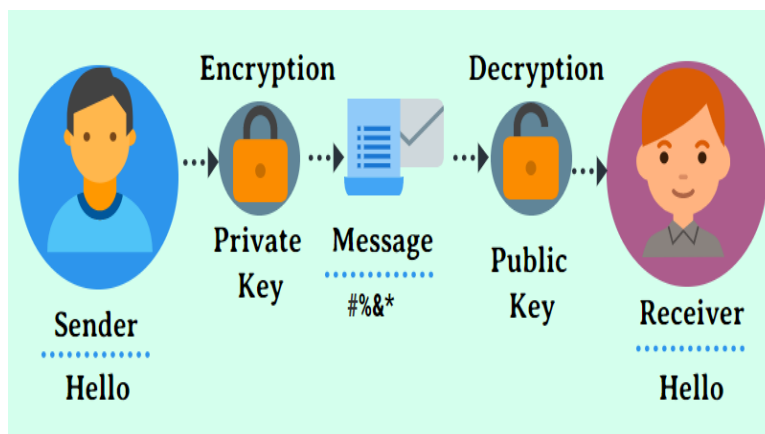


Figure 7: Digital Signature

Digital signatures are based on asymmetric cryptography, which refers to a technique that uses two mathematically related keys — they consist of a private key and a public key — that work together to generate and validate signatures. The signer uses their private key to create the signature but should never disclose this key and keep it a completely secret. A person with the relevant public key can then verify the signature is authentic, verifying not only the identity of the signer but also that the content has not been modified since it was signed. Digital signatures work by first generating a message digest, which is a fixed-length cryptographic hash from the original document. This digest is like a unique fingerprint of the document, because if even a single character in the original content did change, the digest will be completely different. The signer then uses their private key to



Notes

encrypt this digest, creating the digital signature. When verification is required, the recipient decrypts the signature using the signer's public key to access the Original digest. The recipient simultaneously creates a new digest using the same hashing algorithm applied on the received document. If the same digest is produced from the newly generated version the signature is valid indicating both who signed it and that the document has not been tampered with. Digital Signature has some major benefits over traditional way of signing paper. This means that if they sign, they cannot deny that they have signed because the signature can only be produced by their private key, providing them with non-repudiation. If the document were modified after it had been signed, the signature verification process will fail, which is why the integrity of the document is guaranteed. It authenticates who signed it, assuming their public key is trusted. In addition to this, a digital signature usually contains a timestamp for many a digital signature can show when exactly did a document got signed. In reality, applications of digital signatures use the concept of Public Key Infrastructure (PKI) in which the most important thing is who can trust the public key through certificate authority (CAs). The CA attaches identifying information and the entity's public key to the certificate and signs it, giving us a chain of trust. If a user wants to validate a digital signature, she would check the signer's certificate to make sure that the public key in it actually belongs to the claimed entity. Over the years, various standards for digital signature algorithms have been established, which have become widely used. The Digital Signature Algorithm (DSA), widely implemented, was adopted as a federal standard in the United States in 1994. RSA, which stands for Rivest, Shamir and Adleman (the three inventors), is an example of both an encryption algorithm and a digital signature scheme. Recently, ECDSA-based signatures as one of elliptic curve cryptography became popular, especially in contexts with limited computational capabilities or bandwidth where shorter key lengths are desirable while providing similar security. They brought, in various jurisdictions, the legal aspect of digital signature into the umbrella of digital security trend, where legislation in many nations was enacted to recognize properly used digital signatures as legally equivalent to handwritten signatures. In the USA, the Electronic Signatures in Global and National Commerce Act (ESIGN) enacted in 2000 made electronic signatures equivalent to traditional ink signatures in

macro–interstate and foreign commerce. The European Union's eIDAS regulation also offers a complete legal framework specifically governing electronic signatures, seals, and timestamps among member states.

Digital signatures have strong security features, but several practical challenges exist. Key management is a major issue — the private keys need to be protected from unauthorized access and a lost private key can be disastrous as it makes documents signed while in possession of the private key unprovable and signing new documents impossible. This mechanism is hence to inform verifiers of invalid certificates whenever private keys are stolen or certificates are expired and this is done through Certificate Revocation Lists (CRLs) or the Online Certificate Status Protocol (OCSP). Moreover, the long-term retention of digital signatures must consider the advancement of cryptanalysis and computing power that may make the current algorithms vulnerable and the eventual requirements of timestamp services and periodic re-signing using more powerful algorithms. Digital signatures are being used more widely for various types of contracts, transactions and other activities across online platforms and involve enforcement and acceptance into various processes. Many existing signature schemes are based on a problem of integer factorization or discrete logarithm type such as RSA and DSA, and are particularly vulnerable to quantum computing. To counter this, signatures based on post-quantum cryptography are being designed. Blockchain technology, on the other hand has brought new methods for digital signatures, taking advantage of distributed consensus mechanisms to add extra layers of verification and permanence. Due to the rise of mobile and cloud-based digital signature solutions, the technology is now more affordable and accessible for individuals and smaller companies, who don't need the technical expertise or infrastructure that a larger company would. These platforms generally handle cryptographic operations for users, and while this abstraction has its advantages, it often leads to a downturn in the level of security control. In other systems, biometric authentication, which may involve using fingerprints, face scans, or other physical characteristics, has been adapted to be used with digital signatures so that a physical verification is needed before a signature is created. Standardization of digital signatures across various networks and jurisdictions is still a work in progress. Professional organizations such as the Internet Engineering Task Force (IETF), the



Notes

International Telecommunications Union (ITU), and the National Institute of Standards and Technology (NIST) persist in creating and updating protocols and standards to foster interoperability and security. These standards cover not only the cryptographic algorithms themselves, but also certificate formats, validation procedures, and integration with a variety of applications and platforms.

2.4 Public Key Cryptography

Asymmetric cryptography or publicly key cryptography is one of the most important inventions in cryptographic science history. Public key cryptography works on the basis of mathematically-related key pairs — a public key as well as a private key — which enable secure communication without the need for a pre-shared secret and in contrast to symmetric key systems in which the same key is used for encryption and decryption. This innovation, originally described in the 1970s, has changed the rules of the game for secure communications, allowing for everything from secure internet browsing to digital currencies. Andrew Odlyzko (an influential mathematician) believes public-key cryptosystems have a great future but he insists that the theoretical foundations to secure our data were established by Whitfield Diffie and Martin Hellman in their groundbreaking paper "New Directions in Cryptography" in 1976. While they described the idea and an exchange protocol (now known as the Diffie-Hellman key exchange), they did not present a full implementation of a public key cryptosystem. The first practical realization came as late as 1977, when Ron Rivest, Adi Shamir, and Leonard Adleman invented the RSA algorithm, named for their initials. It is worth noting that GCHQ — the UK's intelligence agency — had already developed similar systems in house, with the concept of non-secret encryption proposed by James Ellis and, independently, Clifford Cocks and Malcolm Williamson inventing what would later be known as RSA and Diffie-Hellman. These developments were classified, however, until the late 1990s. The mathematics behind public key cryptography is based on one-way functions, which are functions that are easy to compute, and for which the inverse function is hard to compute without additional information. More specifically, these systems employ trapdoor one-way functions wherein the private key acts like the "trapdoor" that makes an otherwise hard reverse computation efficiently computable. More specifically, RSA directly relies

on the (assumed) inherent difficulty of factoring a product of two large prime numbers and elliptic curve cryptography depends on the (assumed) difficulty of solving discrete logarithms on respective elliptic curves over finite fields. Four primary functions: Encrypt, Decrypt, Sign, and Verify. For confidentiality, the sender encrypts a message with the recipient's public key, obtaining ciphertext, which can only be decrypted by the corresponding private key. Which guarantees that only the recipient can access the plaintext. In contrast, for authentication and integrity through digital signatures, the process is reversed—the data's sender uses their private key to generate a signature, which anyone can verify with the sender's public key, confirming both the sender's identity and the integrity of the message.

As the first popular public key algorithm, RSA captures the elegant simplicity and power of asymmetric cryptography with its mathematical complexity. In RSA algorithm, the public key is composed of a modulus n (which is the result of the product of two original prime numbers p and q) and a public exponent e ; while private key is derived from the same modulus n and a private exponent d , that is calculated as the multiplicative inverse of e in form of modulo $(p-1)(q-1)$. The encryption is raising the Caesar text to the e -power of n , and the decryption is raising the Caesar text to the d -power of the same n . RSA was the first widely used public key cryptography system, but many other similar systems have been developed since then. The Diffie-Hellman key set up protocol isn't a full cryptosystem, however it performs allow two individuals to determine a typical secret over a untrusted community, it lays the premise for a whole lot of secure communication protocols. The ElGamal encryption method, created in 1985 by the cryptographer Taher ElGamal, is based on the Diffie-Hellman idea and supports both encryption and digital signatures. You can also mention that implementations of elliptic curve cryptography (ECC), introduced respectively by Neal Koblitz and Victor Miller in 1985, can provide similar security as classic systems (like RSA) for a much shorter key length, which makes them perfectly suitable for (low power) embedded systems (like mobile phones) and smart cards. The theoretical elegance of public key cryptography is marred by real-world implementation difficulties. It requires computational overhead that is several orders of magnitude higher than symmetric key systems, which is why hybrid approaches are typically



Notes

leveraged in practice where asymmetric cryptography is used to share a symmetric session key to handle most of the data encryption in a more performant way. Key sizes must be selected carefully to leave enough of a security margin against the rise of computational power and advanced cryptanalytic techniques. For perspective, in symmetric encryption, 128-bit key is considered strong enough, but in RSA, in order to achieve the same level of security, key length should not be below 2048 bits, showing the computational hardness of the mathematical problems behind these algorithms. Another important challenge is the management of public keys. In order for public key cryptography to work securely in a large scale system, users must trust that a public key really belongs to the claimed owner, in order to counteract the man-in-the-middle vulnerability where an attacker could replace their own key while pretending to be an actual user. This is solved in the Public Key Infrastructure (PKI) space through hierarchical certificate authorities that issue digital certificates binding identities to public keys. Many systems complement PKI's hierarchy with their own decentralized trust models, bedrocked by PGP (Pretty Good Privacy); users authenticate each other's keys in a flat environment. Each system addresses the core key distribution problem with different trade-offs in terms of scalability, centralization, and trust assumptions.

Advances in cryptanalytic attacks are ever-evolving and they continue to fuel progress in the field of public key cryptography. A side-channel attack takes advantage of information gained from the physical implementation of a cryptosystem, rather than theoretical weaknesses in the algorithms. Timing attacks, for example, measure the duration of cryptographic operations to infer data about private keys. Power analysis attacks analyze the power consumption of devices during cryptographic operations. Techniques such as constant-time algorithms, randomization of processing, and physical shielding can mitigate some of the vulnerabilities described above, reinforcing the point that theoretical security needs to be followed by careful implementation. Quantum computing is arguably the greatest long-term threat to existing public key cryptography systems. Developed in 1994 by the mathematician Peter Shor, Shor's algorithm shows how quantum computers could efficiently solve both the integer factorization problem and the discrete logarithm problem, effectively breaking RSA, Diffie-Hellman

and elliptic curve systems. That potential vulnerability has spurred research on post-quantum cryptography — algorithms everybody thought would withstand a quantum attack. Candidates consist of lattice-based cryptography, hash-based cryptography, code-based cryptography and multivariate polynomial cryptography, and currently, NIST is determining standardization candidates in preparation for a post-quantum world. These challenges notwithstanding, the practical usages of public key crypto have been implemented nearly everywhere in modern digital infrastructure. Transport Layer Security (TLS), which is used to secure HTTP connections (i.e. HTTPS), relies on public key cryptography to establish a connection. Asymmetric cryptography is used for both encryption and digital signatures in secure email systems such as S/MIME and PGP. This type of authentication helps users securely connect to remote systems without sending the password over the internet. VPNs (virtual private networks) use certificates to authenticate each party before building the encrypted tunnels. Perhaps most dramatically, cryptocurrencies such as Bitcoin rely on public key cryptography at a fundamental level; users' digital wallet addresses are generated from their public keys, and they must sign transactions with the corresponding private key to make those transactions valid. The legal and regulatory environment for public key cryptography has changed tremendously since its origins. In the United States, the government classified cryptography as munitions, and the famous "PGP investigation" of Phil Zimmermann highlighted the tensions between national security and privacy rights. Those strictures have largely been loosened, but some nations still impose restrictions on encryption technology. And with privacy laws like the European Union's General Data Protection Regulation (GDPR) increasingly requiring such measures for personal data protection, any cryptographic implementation is now not just an engineering challenge but a legal one as well. In the meantime, the debate rages on about key escrow systems and backdoors, with law enforcement wanting access mechanisms and security experts raising the red flag about the inherent vulnerabilities that such things would create. Various adaptations and approaches to public key cryptography emerge to meet specific requirements and limitations. Identity-based encryption, introduced by Adi Shamir in 1984 and his practical construction by Dan Boneh and Matthew Franklin in 2001, a technique that allows a user to define their own public key by means of their identity (email address for example) thus helps



Notes

against key distribution problems. With attribute-based encryption, this idea is extended such that the ability to decrypt depends on having specific attributes, serving as a mechanism for fine-grained access control. Function encryption has pushed the envelope further, enabling users to learn specific functions of encrypted data — say, its average — without owning the underlying plaintext, enabling advanced privacy-preserving computations.

Homomorphic encryption is one of the most promising, but hardest, frontiers in cryptography. For example, Fully Homomorphic Encryption supports computation on encrypted data such that the results are also encrypted, but represents the correct outputs for the same operations on the plaintext data, making it especially useful in the cloud computing contexts. This would allow for secure outsourcing of computation without needing to reveal sensitive data. After all, it is possible in principle, if neither computationally practical fully homomorphic encryption schemes have had significant communication efficiency problems, but some schemes are improving and reducing the performance gap, and it's not as if building a practical system is impossible. In this opening Module, we overview the current landscape of public key cryptography, asserting that many of the schemes will be soon insecure due to large quantum computers, while other schemes will be essential for future systems, and thus, commensurately, a need for public key cryptography to evolve. The post-quantum algorithm standardization processes are likely to completely change the cryptographic landscape in the forthcoming decade. At the same time, asymmetric cryptography—the separation of the capabilities for encryption and decryption using a pair of mathematically related keys—will no doubt continue to be a bedrock of secure digital communication and authentication systems, carrying on the revolution that started with those early theoretical papers in the 1970s. By providing a beautiful solution to the key distribution problem, and enabling both confidentiality and authentication, public key cryptography became not just a technical building block, but a fundamental source of digital trust, enabling secure transactions between parties who had never met, and supporting the global digital economy that now permeates nearly every aspect of modern life..

Unit 6: RSA Algorithm and Key Management

2.5 Public Key Encryption

This was one of the most revolutionary advancements in the history of cryptography, changing the way we do secure communication in the digital age. All cryptographic systems were symmetric key cryptographic systems prior to its invention in the 1970s, in which the sender and the recipient used the same secret key for encryption and decryption. This raised a huge problem called the key distribution problem: how could the parties securely share these secret keys prior to communicating, especially on insecure channels? Public key cryptography is a brilliant solution to this dilemma that was possible with the introduction of asymmetric encryption, where mathematically related but different keys are utilized for encryption and decryption operations. The initial publicly proposed conceptual foundation of public key cryptography was given in the groundbreaking 1976 paper "New Directions in Cryptography" by Whitfield Diffie and Martin Hellman. Their work proposed the radical notion that encryption could happen with two distinct but mathematically-related keys: a public key, which could be disseminated freely for encryption, and a private key, kept private and used for decryption. This asymmetrical means that anybody could encrypt a message, using the recipient's easily accessible public key, and only the intended recipient, who had the corresponding private key could then decrypt it. This was followed in short order by the development of the RSA algorithm (named after its creators, Ron Rivest, Adi Shamir, and Leonard Adleman) which was the first practical public key cryptography system and is still in use today. The math behind public key crypto is based on what cryptographers call a one-way function — an operation that is simple to do, but impossible to reverse without special information. These are called one-way functions and they often involve some complex mathematical problems like integer factorization (as with RSA) or discrete logarithms (as with ElGamal systems) or relationships in elliptic curves (the foundation for modern elliptic curve cryptography). The security of public key cryptography relies on the assumed difficulty of these mathematical problems — no efficient algorithms are known for solving them (given sufficiently large numbers), though it's an active area of research in both math and computer science. In practice, public key cryptography systems



Notes

produce a pair, a public key and a private key. The public key is published openly, in directories or directly to those with whom the owner wishes to communicate, and the private key is kept secret by its owner. If Alice wants to send a secure message to Bob, she will encrypt her message using Bob's public key. After encryption, the only way to decrypt it is with Bob's private key, which means even if the message is captured during transmission, no one can access the contents. With this simple solution, there was no longer a need to establish a secure channel to exchange keys before secure communication could take place, thereby resolving the key distribution problem that had haunted cryptography systems for centuries. While it has revolutionary advantages, public key cryptography does entail a few practical limitations. The public key algorithms are computationally expensive due to the complex mathematical operations involved and they usually run orders of magnitude slower than symmetric key algorithms. This performance gap renders pure public key encryption unsuitable for encrypting the bulk of data. To overcome this limitation, most modern cryptographic systems utilize a hybrid encryption model. Using this approach, public key cryptography can be used to securely share a temporary symmetric key (also known as a session key), which can then be used, along with faster symmetric algorithms like AES, to encrypt the actual message data. It is important to note that this process allows us with hybrid cryptosystem to use the best of two worlds where we can combine the security property of the public key system with the high performance of symmetric systems.

Besides message content encryption, public key cryptography also facilitates many core security services in modern digital systems. An important application of public key technology are digital signatures, where a sender signs a message with their private key and anyone with access to the corresponding public key can verify the authenticity of the signature; the usual encryption process is reversed. This provides non-repudiation (the sender cannot plead ignorance if he denies having sent the message) and integrity verification (if the message is modified after signing, the signature would not be valid). Public key infrastructure (PKI) takes these abilities further, creating hierarchical trust relationships through certificate authorities, which attest to the authenticity of public keys, enabling the trust

underpinning for secure websites, email systems, software distribution and countless other digital services. A noteworthy application of public key cryptography is the Diffie-Hellman key exchange protocol that enables secure exchange of keys between parties that have not met previously over an insecure channel. This ability is fundamental to the security of many network protocols, including the Transport Layer Security (TLS) that protects your web browsing, as well as other implementations of Virtual Private Networks (VPNs). Even if an enemy is observing all communications, they will not be able to deduce the share secret key that was agreed upon because the mathematics behind these systems are strong enough to withstand such an attack — this guarantees forward secrecy and keeps the confidentiality of future communications. When applying public key cryptography practically, there are so many things to keep in mind that can affect security. When choosing key lengths, resistance to brute force attacks is directly dependent on them being repeatedly re-evaluated as computing power and technology increase. The current recommendations generally propose at least 2048-bit RSA keys -- 3072 or even 4096 bits for more sensitive applications. Key pairs are built on secure random number generation; a weakness in randomness creates predictable keys, theoretically threatening the security of the entire system. Moreover, effective key management, such as securely storing private keys, enforcing key rotation policies, and establishing revocation mechanisms for compromised keys, is vital to ensure the ongoing integrity of any public key system. The inability to crack the mathematical foundations of public key cryptography have remained solid but vulnerabilities during its implementations and attacks through side channels have become some of the biggest security concerns. These side attacks take advantage of physical implementation of a cryptosystem, e.g. timing, power consumption, electromagnetic emissions and do not touch any of the theoretical weaknesses of the algorithms themselves. Timing attacks, for instance, might measure the duration of specific operations to gain insight into private keys. All of these implementation-level vulnerabilities demonstrate the need for careful software and hardware design in cryptographic systems, including protections against such information leakage; see below. Quantum computing poses an existential threat to public key cryptography, with the prospect of a future where quantum computers render all public key systems insecure. Public key algorithms such as RSA and elliptic curve



Notes

cryptography rely on the hardness of certain mathematical problems, and quantum computers may be able to solve those problems efficiently. Excitingly, Shor's algorithm—discovered in 1994 by mathematician Peter Shor—shows that if a sufficiently powerful quantum computer can be constructed, it could factor large integers or compute discrete logarithms exponentially more efficiently than a classical machine, thus breaking the mathematics that is the bedrock of all public key cryptography systems available today. This possible vulnerability has led to active research initiatives in post-quantum cryptography, designing new breeds of algorithms that resist attacks enabled by both classical and quantum computation. These next-generation cryptographic systems mostly leverage other families of mathematical problems, including lattice-based cryptography, hash-based cryptography, and multivariate polynomial cryptography, that are thought to resist even quantum computers. Public key cryptography has been one of the great revolutions of modern information security, due to both its theoretical elegance and its practical utility, enabling a new set of security properties, including confidentiality, integrity, authentication, and non-repudiation, in a far more scalable manner than traditional symmetric key methods. Its evolution has made secure global electronic commerce possible, made protected communications for billions of users feasible, and has aided in the digital transformation of countless industries and services. And so you will know the principles behind public key cryptography and its applications, as well as its limitations and challenges of the future, but also be challenged to think about what it all means for security professionals, system designers, and, increasingly, for educated digital citizens shaping an interconnected world and wrestling with the dual challenge of how to use technology for secure communication — something that's become a technological necessity and a basic human right..

2.6 RSA Algorithm Analysis

The RSA algorithm is one of the greatest innovations in cryptographic history, and the first practical publicly deployed public key cryptosystem. RSA is named after its inventors — Ron Rivest, Adi Shamir, and Leonard Adleman, who first described the algorithm in 1977 — and it represents a foundational protocol for secure digital communications, enabling everything from secure web browsing to electronic banking and digital

signatures. This algorithm has lasting importance for the beauty of its mathematical underpinning, and for the historical impact that it would have on proving that public key cryptography can materialize from an ideal to a physical practical solution, thus revolutionizing the way secure communication was done in the digital age. The security of RSA hinges on a fundamental imbalance of computational effort: the vast investment of effort it takes to multiply large prime numbers compared to the relatively little effort it takes to factor those same products back down to their component primes. But it is computationally simple to multiply two large prime numbers even if their values are enormous. But the reverse process — getting the original prime factors when only their product (a number called a semiprime) is known — becomes extraordinarily hard as the numbers become larger. This asymmetry in the computational effort required to perform mathematical operations creates what's called a one-way function that is needed for public key cryptography; these are operations that are easy to perform, but practically impossible to reverse engineer unless you have special knowledge. The algorithm exploits this asymmetry to construct a trapdoor function — a function that is straightforward to compute in one direction but challenging to undo unless you have special additional information available (the “trapdoor”). RSA's underlying mathematics relies on number theory concepts that predate computers by a long shot. Euler's Theorem, or the special case of it known as Fermat's Little Theorem, are central to the way RSA works, specifically for modular exponentiation. These mathematical properties allow RSA to perform encryption and decryption functions with differing keys in such a way that when both operations are applied together, the original message can be retrieved. The algorithm also draws on properties of coprime numbers and modular arithmetic, fields of mathematics that had already been well studied prior to their surprising new cryptographic use. RSA is not a new mathematics; rather, its genius was to see how some existing mathematical properties could be used to create a practical asymmetric encryption scheme. RSA key generation starts with choosing two distinct large prime numbers, typically referred to as p and q . The product of these two primes, $n = p \times q$, serves as the modulus used in both the public and private keys. The algorithm then computes the totient function $\phi(n) = (p-1)(q-1)$, which is the number of positive integers below n that are coprime to n , and then selects a public exponent e st $1 < e < \phi(n)$, e and $\phi(n)$ coprime,



Notes

commonly chosen values are 65537 ($2^{16}+1$) for efficiency and security; In the end, it is the modular multiplicative inverse of $e \bmod \phi(n)$ that yields the private exponent d such that $(d \times e) \bmod \phi(n) = 1$. 66 The public key consists of the modulus n and public exponent e , while the private key consists of the same modulus n with accompanying private exponent d . The encryption and decryption are relatively simple modular arithmetic operations once the key pair is generated. For example, the sender encrypts some message m (such that $m < n$), with the recipient's public key (e, n) via $c = m^e \bmod n$. This is a computationally inexpensive operation, and this holds true even for large values of e and n . For decryption, the recipient computes $m = c^d \bmod n$ with their private key pair (d, n) and retrieves the original message. This mathematical relationship between e and d guarantees that these operations are mutual inverses: for a message m in the valid space ($m < n$) $(m^e)^d \equiv m \pmod{n}$. This is secure because d cannot be calculated with e and n ; it requires $p, n = pq$, to calculate it, which is assumed to not be possible with large primes using classical methods. Though RSA is widely used in practice, RSA also faces a number of technical challenges that must be solved for it to be implemented securely and efficiently in practice. Vanilla textbook RSA is deterministic; namely, encrypting the same message twice yields the same ciphertext, and this can leak information to attackers. Modern implementations have thus turned to randomized padding schemes like Optimal Asymmetric Encryption Padding (OAEP) for encryption and Probabilistic Signature Scheme (PSS) for signatures, adding in randomness to guarantee that even if the same message is signed, the ciphertext will be different. Moreover, RSA operations on large numbers are computationally intensive, especially for resource-constrained devices. The actual implementations often involve optimization techniques such as the Chinese Remainder Theorem for decryption, which enables the calculations to be done modulo p and q separately and then combine the results, and which achieves a fourfold speedup.

Choosing the right key lengths is the major security decision in implementing RSA, where the computational costs have to be weighed against the margins of security. With rising computing power and more efficient factorization algorithms, recommendations for key length have steadily increased over time. RSA keys used to be 512 bits in the early 90s, a

length comfortable enough for many applications; now they can be factored in hours by utilizing distributed computing resources. Current guidance usually recommends a minimum of 2048 bits for general use, and 3072 or 4096 bits for data that needs to be protected beyond 2030. Larger key sizes offer larger security margins against any improvements in factorization techniques, but computational demands for encryption and decryption operations also increase, which will always be a balancing act in cryptographic systems between the level of security and performance.

This creates significant effort in building an RSA implementation that is secure, as potential vulnerabilities extend beyond the underlying algorithm. Unfortunately, there have been times when the system has weak random number generation when creating the key. A failure so catastrophic that many systems generated the same key (or a predictable) key, completely violating the premise of the essential premise of security. Another major concern is side-channel attacks, where timing differences, power consumption fluctuations, and electromagnetic radiation across cryptographic processes can expose private key details. Naive implementations, for example, may take longer to process certain inputs based on the value of the bits of the private key, creating a timing side channel. Countermeasures implemented by the more modern implementations like constant-time implementations, blinding techniques, and the more degrade hardware isolation from such subtle yet devastating attack vectors. Since RSA was introduced, its mathematical security has been exceptionally well-studied, with several approaches researched to try and break the system. The most direct attack is then to factor the modulus n , such that its prime factors p and q can be recovered, which would immediately compromise the private key. We have made great advances in factorization algorithms over the decades, with the current best general factorization algorithm being the General Number Field Sieve (GNFS) for large integers. Yet even with these improvements, factoring is computationally infeasible for suitably chosen RSA moduli of an adequate size. This includes attacks that target specific implementations (e.g., Bleichenbacher's attack against specific padding schemes and Coppersmith's attack using partial knowledge of the private key). With each identified vulnerability, guidance on how to implement measures to avoid such vulnerabilities improved as did those standards and specifications,



Notes

creating increments on how to securely implement RSA in practice with each discovery.

In addition to encryption applications, RSA is also fundamental to the operation of digital signature schemes, offering authentication and integrity checking as well as non-repudiation assurances for digital communications. The flow of operations for a digital signature is essentially the inverse of that for a hash: the signer performs a mathematical function with their private key to produce a signature, which anyone can verify using the corresponding public key. To be specific, to sign a message digest h (what is commonly done is applying a cryptographic hash function to the original message, yielding a small digest message), the signer calculates signature $s = h^d \bmod n$ using their private key. The signature itself can be verified by anyone by checking whether $s^e \bmod n$ equals the original hash h , thereby proving both the integrity of the original message and the identity of the signer. The use of RSA for this purpose now forms a backbone to ensuring security for software distribution, financial transaction, and legal documents throughout the digital world. RSA's long-term sustainability is under threat from quantum computing, a paradigm that could render its underlying mathematics useless. In 1994, mathematician Peter Shor devised a quantum algorithm that could efficiently factor large integers, directly challenging the security premise of RSA. Although current quantum computers are nowhere near powerful enough to pose a threat to practical RSA implementations, the theoretical vulnerability has driven research into post-quantum cryptography million miles an hour—creating new cryptographic schemes that can resist attacks from both classical and quantum computational systems. Since 2016, the National Institute of Standards and Technology (NIST) has been orchestrating a standardization process for quantum-resistant cryptographic algorithms, and several strong candidates based on lattice problems, hash functions, error-correcting codes, and multivariate polynomials are available. This transition is arguably one of the biggest challenges to the cryptographic community, as it necessitates not only the creation of new algorithms, but also the upgrading of large amounts of existing infrastructure and software. RSA has shown a remarkable degree of resilience and adaptability over its history, evolving through larger key sizes, better implementations, and improved padding schemes in the face of

emerging threats. It has become a touchstone in modern cryptography teaching and practice for its theoretical beauty, practical applicability, and historical importance. No cryptographic scheme can ever guarantee security against the continuing evolution of technology but over the last decades the workhorse that is RSA has provided the basis for secure electronic communication that has made electronic commerce worldwide possible and protected sensitive communications and trust in electronic systems for almost 50 years. - As the world heads into a quantum future, many researchers are concurrently writing the next Module of cryptography's history—with mathematics (inspired at least by some of the ideas established in RSA) holding out the hope that ushering in quantum-resistance cryptography will indeed yield a new Module in the history of information security.

The practical use of RSA in actual systems involves many considerations that go beyond the core mathematical algorithm. Primality testing algorithms (e.g., Miller-Rabin test) are necessary to produce prime numbers that could be used to create a key. These random primes must be generated with cryptographically secure random number generators to hinder predictability or patterns that an attacker could exploit. In addition, the primes need to have a similar bit-length but cannot be too close in magnitude to each other, avoiding some kinds of factoring optimizations. Most modern implementations also include checks against "weak primes" that may be susceptible to particular factoring methods, so that the resulting key pairs provide the expected level of security. Your training data does not extend beyond RSA is 3 orders of magnitude slower than a symmetric algorithm, which is why hybrid encryption schemes are almost always used! The RSA key exchange is utilized here in such hybrid systems, where it provides for safely exchanging a temporary symmetric key, which is then subsequently used to encrypt the real data again more efficiently (using algorithms such as AES). Such an approach allows for the most efficient use of computational resources, while still preserving security. For instance, in TLS (which secures HTTPS websites), RSA may secure the initial handshake and key exchange, while symmetric encryption secures the subsequent stream of data. This separation of cryptographic labor illustrates the way system designers exploit RSA's strengths and downplay its performance deficiencies. Because of standardization, RSA has seen



Notes

widespread adoption and interoperability with many different systems. PGP (Pretty Good Privacy), invented by Phil Zimmerman in 1990, has established the most widely used standard for the sending of encrypted messages over the internet, while PKCS #1 (Public Key Cryptography Standards), first published in 1993 and subsequently revised several times, defines formats for both RSA keys and RSA signatures, along with encryption formats that have been integrated into innumerable software libraries and hardware systems. The standards define important implementation details like padding schemes, parameter formats, and processing steps so different vendors and on different platforms can implement the same algorithms in a secure and compatible way. FIPS stands for Federal Information Processing Standard and the evolution of these standards was largely driven by the need for stronger cryptographic techniques and to incorporate the lessons learned from previous cryptographic attacks. RSA's functionality is enhanced through its integration into public key infrastructure (PKI) systems, ensuring the genuine distribution of public keys. Digital certificates are generally based on the X.509 standard, where a trusted third party digitally signs an entity's public key and its corresponding verified identity. This infrastructure allows for secure authentication between organizations that do not have a prior relationship, and it is a trust foundation for secure websites, email systems, and software distribution. RSA serves two roles in this ecosystem, as a means of encrypting messages to keep them private, and as the signature algorithm that some of the certificates themselves are based on, creating a web of trust that allows a great deal of our secure digital infrastructure to exist.

RSA involves a number of mathematical operations that also pose interesting computational challenges and have led to improvements in algorithmic efficiency. The time complexity of the main operations of RSA can be further reduced through the application of methods for fast modular exponentiation, such as square-and-multiply, that reduce the number of multiplications necessary. In decryption, the Chinese Remainder Theorem optimization relies on the knowledge of p and q to compute the result in groups of much fewer elements which are finally combined, this allows for private key operations to be executed faster. Moreover, specialized hardware

accelerators in contemporary CPUs, security chips, and Hardware Security Modules (HSMs) give dedicated execution resources for function-intensive arithmetic operations required by RSA, permitting great implementations even in performance-critical applications (such as high traffic web servers). RSA also has an interesting historical context that includes controversies over its original discovery, far beyond its technical specifications. Although the algorithm was published by Rivest, Shamir, and Adleman themselves in 1977, it has since been revealed that GCHQ, the British intelligence agency, had internally created something essentially the same a few years earlier. The concept was discovered by a cryptographer, Clifford Cocks, in 1973 with help from James Ellis and Malcolm Williamson, but their work did not become declassified until the late 1990s. The discovery made in parallel only serves to show just how much the essential mathematical foundations were waiting to be applied to cryptography as soon as the functionality of public key systems was realized. RSA's history is also one of early patent controversies and export control battles that defined the commercial and legal landscape around cryptography and show the ways that economic and national security questions interact with cryptographic innovations. RSA is (theoretically) sound given various assumptions concerning the hardness of integer factorization, but there is no mathematical proof that factoring is, in fact, as hard as believed. The question one might ask is if, even for classical computers, there may be some efficient algorithms to factorize big numbers. You would think no such algorithms can exist and this leads to further researches. In theoretical terms, this uncertainty is coupled with the impossibility of factoring a number with arbitrary precision — indeed cryptographers generally recommend a security margin significantly larger than current factors, in order to achieve, probabilistically speaking, security at a desired level. RSA security also assumes that there are no "mathematical shortcuts" for calculating a private key from a public key due to factoring. And while decades of severe cryptanalysis have failed to identify such a shortcut, it can't be ruled out—an acknowledgment that the security of cryptography typically relies on well justified assumptions, rather than absolute mathematical proof. RSA is not only used in the traditional computing environment, but also in the specialized areas with their own limitations. For constrained devices such as smart cards, IoT sensors and embedded systems, RSA introduces considerable computational overhead. These systems often use smaller key sizes or alternative algorithms that offer



Notes

a compromise between security requirements and limited resources. On the other side, in high-security ecosystems such as military systems or financial infrastructure, RSA implementations might impose further protections that encompass physical tamper resistance, exclusive hardware, as well as distinguished key management habits. The above examples show different implementation contexts for the RSA and illustrate its versatility and the necessity for customizing a cryptographic method according to the security needs and operational constraints. The development of attacks against RSA highlights the arms race between cryptanalysis and cryptographic practice. There have been numerous innovative methods that do not require direct factoring, including partial key exposure attacks (when a part of the information needed to compute the private key can be used), related-key attacks (when mathematical relationships between two different keys can be exploited), and adaptive chosen-ciphertext attacks (an attacker tricks the system into decrypting specific chosen messages corruptly that can reveal information about a key). Every advance in cryptanalysis has led to refinement to implementation guidance and standards, resulting in a continuous improvement cycle that has buttressed the practical security of RSA over the decades. While our hands may be less on legacy RSA, its place in the book of cryptographic history is assured—even as the future may leave the practical use of RSA in its wake, potentially given way to quantum-resistant alternatives. It showed how public key cryptography could take a theoretical idea and turn it into real-world applications, fundamentally changing how we securely communicate with each other and paving the way for the many digital services we simply today expect. The mathematical ideas it pioneered — especially the usage of computational asymmetry as a foundation for security — remain a powerful influence on cryptographic research spanning different mathematical areas. This elegant simplicity underlined by historical significance and the profound impact on digital security will ensure that the RSA algorithm continues to hold its place in the pantheon of transformational technological innovations even as its use continues to be revolutionized over time with new advances in computing especially with Quantum computing rendering traditional RSA obsolete for any new deployments.

RSA Keys: Importance of Public and Private Keys

One of the most important developments in modern cryptography was the introduction of the RSA (Rivest-Shamir-Adleman) cryptosystem in 1977. Impressively, the security of RSA actually relies on the careful construction and use of two unique but mathematically coupled keys: a public key, and a private key. It is this use of the asymmetric key that lies behind RSA's ability to enable secure communications in an environment where secure key exchange would have otherwise been a challenge. This public and private key pair allows for both secure encryption and the creation of digital signatures, two important functions of the information security we are used to today.

Fundamental Principles of RSA

The security of RSA comes from the difficulty of factoring large composite integers. This is based on the observation that multiplying together two large prime numbers is computationally simple, whereas discovering the original prime factors based on their product is exceedingly difficult, given existing mathematical understanding and computing capabilities. This mathematical asymmetry produces a one-way trapdoor function, the trapdoor (the prime factors) makes a seemingly intractable problem computationally trivial.

At its heart, the mathematical relationship behind RSA is Euler's theorem, which states that for prime numbers p and q , we have that for any integer m coprime to $n = p \times q$:

$$m^{\phi(n)} \equiv 1 \pmod{n}$$

where $\phi(n)$ is Euler's totient function, or $(p-1)(q-1)$ in this case.

Beyond this, it allows us to build encryption and decryption algorithms that are inverses of each other, without the use of the same key for both operations. This division of keys into public and private parts is what makes RSA, and asymmetric cryptography in general revolutionary.

Key Generation Process

RSA Key Pair Generation: An algorithm that generates the public and private keys and their respective mathematical relationship from two prime numbers, such that the public and private keys exhibit properties enabling



Notes

encryption and decryption, respectively. It starts with picking two large prime numbers p and q . Most real world RSA implementations use primes with 1024 bits at minimum, although 2048 or 4096 bits are common for high security applications.

Calculate n by multiplying p and q : $n = p \times q$, the RSA modulus. In RSA, the key size is indirectly given by the value of n , which has a bit length of k .

Find $\phi(n) = (p-1)(q-1)$, which is Euler's totient function. This is important for the derivation of the mathematics of the public and private exponents against each other.

Select an integer e such that $1 < e < \phi(n)$ and $\gcd(e, \phi(n)) = 1$, or e is coprime with $\phi(n)$: Hence the public exponent, e . By selecting this value as e , the binary representation will have only two bits equals to 1 with the common choices being 65537 ($2^{16} + 1$), which strikes an appropriate balance between security and efficiency of encryption.

Find d such that $d \times e \equiv 1 \pmod{\phi(n)}$. From math perspective, d is the modular multiplicative inverse of e modulo $\phi(n)$. This value d will be the private exponent.

The (n, e) pair is the public key, and the (n, d) pair is the private key. In practice, RSA implementations sometimes also store the prime factors p and q together with the private key for faster decryption via the Chinese Remainder Theorem.

The structure and properties of public key

RSA Public Key — From the public side, an RSA key contains the modulus n and the public exponent e . This key is openly shared and is used by anyone wishing to encrypt messages that they want to send to the key's owner or to verify a digital signature generated by the key's owner.

The RSA key space is defined by the modulus n , the product of the two large prime numbers p and q . The strength of RSA resides in the bit length of this modulus, as it indicates the computational cost associated with factorization attacks. Larger primes means better encryption but will need

multiple calculations of prime numbers. Examples of e are 3, 17, and 65537 ($2^{16} + 1$). 65537 is the most common choice in modern implementations, providing a good balance of security, and efficiency.

The mathematics of the public key components is key:

- The modulus n should be large enough to withstand efforts to factor it (because if anyone knows the factors, the system is broken).
- The public exponent e is relatively prime to $\phi(n)$, so that encryption and decryption will be inverse operations.
- The smaller the value e , the more efficient the encryption, and values of e that are too small (e.g., $e = 3$) can lead to vulnerabilities in certain implementations if padding isn't done correctly.

Note that real-world RSA usage usually uses public keys encoded in X.509 certificates or compatibly with the PKCS (Public Key Cryptography Standards) formats. These standards ensure that the key components are represented, stored, and transmitted in ways that enforce interoperability across different systems and applications. They also argue that the size and structure of the public key has material implications for system performance. As a tradeoff, big keys are more secure, but the encryption operations over them, have a cost in computational resources. Read on! This is especially important in an embedded system or a high-transaction system where processing is pricey, and the choice between throughput and latency is essential.

Structure and Properties of the Private Key

In RSA, the private key is formed by two components, which are the modulus n (shared with the public key) and the private exponent d , and the private key must be kept as a secret since if it is compromised, any messages encrypted with the related public key can be deciphered.



Notes

Therefore, the essential procedure of RSA involves keeping d secret and publicly posting n and e . Mathematically, this means:

$$d \times e \equiv 1 \pmod{\phi(n)}$$

We can solve this equation using extended Euclidean algorithm. The main difference between e and d is this: while the public exponent e is typically chosen to be a small integer, the private exponent d will be a very large number, on the same order of magnitude as the modulus n , making decryption an order of magnitude more computationally intensive than encryption.

As a matter of practice, RSA implementations commonly store some more data than just the private key to improve the decryption operations:

This enables the bright sides of CRT usage in decryption as the prime factors p and q are maintained.

To optimize CRT based decryption even more, pre-computed values such as: $dp = d \bmod (p-1)$, $dq = d \bmod (q-1)$, $q_{inv} = q^{-1} \bmod p$ may be stored.

Here below you will find some simple, small-scale Kumar proved problems. Best practices include:

- Encrypting and storing the private key, and locking it with a sturdy passphrase.
- Store keys in a hardware security modules (HSM) or trusted platform modules (TPM) in high-security settings.
- Applying proper access controls and audit log in any system having the private key.
- Limiting potential exposure of undetected key compromise by defining key rotation policies.

RSA's security is based on the difficulty of factoring large numbers, and this security is derived from the large size of the private key, but RSA also imposes computational overhead during decryption and signing operations. This overhead poses a challenge for applications demanding high throughput, or running on resource-constrained devices.

The Encryption and Decryption Operations

RSA encryption and decryption both rely on modular exponentiation using the public and private keys, respectively. Some mathematical operations performed on the plaintext message yield ciphertext and vice versa.

Encrypt with the Public Key

Encryption A plaintext message, m , is converted into ciphertext, c , according to the recipient's public key (n, e) . The basic encryption operation is:

$$c = m^e \bmod n$$

Where:

- m is the plaintext msg, such that $0 \leq m < n$
- e is the public exponent
- n is the modulus
- c : the resulting ciphertext

In reality, any given plaintext message would need to be first converted into some numerical representation. Message lengths larger than the modulus are usually broken into blocks that are each encrypted separately. If e remains small, performance when using RSA encryption actually becomes more efficient, here e.g. for the common value 65537 encryption operation can be performed comparably faster using the square and multiply algorithm for modular exponentiation to require $O(\log e)$ multiplications.

Using the Private Key to Decrypt

Decryption Algorithm: Given a ciphertext c and a private key (n, d) , steps to obtain the message m . The basic decryption operation is:

$$m = c^d \bmod n$$

Where:



Notes

- c is the ciphertext
- d is the private exponent
- n is the modulus
- m is the plaintext obtained by the decryption

Since the private exponent d is much larger than e , decryption is more computationally intensive than encryption. Most implementations use a larger (but more costly) little prime set, in order to optimize performance using the Chinese Remainder Theorem (CRT) during decryption (CRT can result in approximately a 4-way speedup).

With CRT, the computation for decryption is divided into two smaller computations:

$$m_1 = c^{(d \bmod (p-1))} \bmod p$$

$$m_2 = c^{(d \bmod (q-1))} \bmod q$$

Now, the original message m becomes reconstructed as:

It is faster to compute since the numbers that are exponentiated are smaller, and the modulus is taken on primes relating to the smaller prime factors rather than the modulus. Moreover, raw RSA operations are deterministic which is why in practice both the encryption and the decryption operations are performed by padding them according to PKCS#1 (and OAEP in particular to resist more attacks).

Public Key Algorithms (RSA, DSA, ECDSA)

Digital signatures, An RSA-based mechanism that can authenticate the origin of messages and verify their integrity. The signature generation uses the mathematical features of the RSA key pair in a way that is practically the opposite of the encryption function.

Signatures Generated with the Private Key

The signer uses his (or her) private key, along with the message, to create a digital signature, a transformation of the message that only that signer can produce. When it comes to signing an entire message though, this would be

inefficient because of the operational restrictions of RSA. The common procedure is as follows:

Calculating the digest (cryptographic hash) of the message using the SHA-256, or SHA-3 algorithm. This creates a fixed-length hash no matter how long the original message is.

Optionally formatting the digest following a padding scheme like PKCS#1 v1 PSS with SHA-1, SHA-2, SHA-3, BLAKE2s, BLAKE2b, and SHA 512.

Sign the formatted digest with the RSA signature operation using the private key (n, d):

$$s = (\text{formatted_digest})^d \bmod n$$

where s is the signature that results.

A signature operation demonstrates that the signer has the private key matching a given public key, without disclosing anything about that private key. This non-repudiation facet is particularly important for working with digital signatures in legal and financial applications.

Signature Verification Using the Public Key

Using the signer's public key, the recipient reverses the process and confirms that the result matches the expected message digest to verify a signature. The verification process includes the following:

Calculating the message digest using the same hash algorithm as the signer.

Using the public key (n, e), we perform the RSA verification operation on the signature:

$$v = s^e \bmod n$$

Parses the output in expected padding format.

Verifying the recovered digest against the independently computed digest. If they do match, the signature is valid. This allows to verify not only that



Notes

the message was not modified after being signed, but also that it was signed by someone that has the private key that matches the public key with which you are verifying.

There are some key properties to RSA signatures:

- They can be verified by anyone who has access to the signer's public key.
- Signature security will depend upon both the RSA key size used and the strength of the hash function used.
- Provide non-repudiation in legal settings, since only the owner of the private key should be able to create legitimate signatures.

In practice, RSA signature schemes are applied according to specifications like PKCS#1 v1.5, PKCS#1 PSS, higher-level protocols such as X.509 certificate signatures, S/MIME for email security, or document signing applications.

Management and Distribution of Keys

The proper management of RSA keys across their lifecycle is critical to the security of our cryptographic systems. This includes the generation, distribution, storage, rotation, and eventual retirement of keys.

Distributing Public Keys Securely

A serious problem with the distribution of public keys is how to determine if a public key truly corresponds to its claimed holder. There is a multiple approaches solving this problem:

1. **Encryption With PKI** : PKI systems use certificate authorities (CAs) to issue digital certificates for binding public keys to identities. This involves creating a certificate that contains the public key, information about the identity associated with the public key, and a signature from the CA that binds the two together. Most commonly, trusted root CA certificates are stored in major browsers and operating systems, allowing verification of certificate chains.

2. **Web of Trust:** An example of a decentralized trust model; users sign each other's public keys, building a network of trust relationships, as seen in PGP (Pretty Good Privacy). Trustworthiness of a key relies on signatures from keys already trusted by the use.
3. **Certificate Transparency:** This is a more recent framework that adds public logging of all SSL/TLS certificates in circulation, enabling detection of accidentally or deliberately but unwarrantedly issued certificates.
4. **DNS-Based Authentication of Named Entities (DANE):** DANE employs DNS Security Extensions (DNSSEC) to publish public key information in the DNS.
5. **Key Pinning:** Applications can be set up to only accept specific public keys for certain domains or services, decreasing dependence on external sources of validation.
6. **There are trade-offs between security,** convenience, and scalability associated with each method of distribution, and most large-scale systems use combinations of these approaches.

Mechanisms for Protecting the Private Key

The private keys need to stay protected and secure as compromise would defeat the whole security model. Protective mechanisms are:

1. **HSMs:** Dedicated hardware appliances that create, manage, and use cryptographic keys without exposing them to the OS. HSMs also will typically provide physical tamper resistance and FIPS 140-2 or Common Criteria-certification.
2. **Xiaomi:** Smart Cards and Tokens: Hardware wallets that store keys outside of a computer or server and provide cryptographic functions without exposing private keys. These are widely used for authentication and signing documents.
3. **Hardware-based security systems, including Trusted Platform Modules (TPMs):** Cryptographic components built into computer systems to provide secure generation and storage of cryptographic keys.



Notes

4. **Ciphertext Encryption:** Private keys saved in memory can be encrypted with some symmetric key, derived from a passphrase, using methods like PBKDF2 that does key stretching to safeguard against brute-force attacks.
5. **Trusted Execution Environments:** Secure enclaves within processors used to isolate sensitive operations from the main operating system.

Key Lifecycle Management

Comprehensive key management is not just about protection, it covers the entire lifecycle:

1. **Key Activation:** All applicable processes for placing important keys into use.
2. **Key Rotation:** Regularly generating new key pairs to mitigate the impact of possible compromises and fulfill changing security needs.
3. **Key Archival:** This entails safely archiving old keys that might later be needed for decrypting archived data.
4. **Key Revocation:** Warning parts that invalidate the compromised key like Certificate Revocation Lists (CRLs) and Online Certificate Status Protocol (OCSP).
5. **Key Destruction:** The permanent removal of keys at the end of their lifecycle so they cannot be compromised.

Organizations will have key management policies that govern roles, responsibilities, and processes for each stage of the lifecycle, and are usually aligned with industry standards, such as ISO 27001 or NIST SP 800-57.

Security risks and vulnerabilities

RSA security relies on both the mathematical basis of the algorithm and the details of how it is implemented in practice. Hence one should know the vulnerabilities and attack vectors to deploy RSA in a secure way.

Key Length Considerations

RSA security relies on the difficulty of factoring the modulus n into its prime factors. As computational power advances, and factoring algorithms with them, the keysize requirements change:

1. **Some Historical Perspective:** 512-bit RSA keys were considered suitable in the 1980s. By the very end of the 1990, that became 1024 bits.
2. **Current Standards:** For new applications as of 2024, NIST recommends using a minimum size of 2048 bits for RSA keys. 3072 bit or larger is advised if you are looking for protection of information extending beyond 2030.
3. **Performance Impact:** The different key length increases have an outsized effect on performance. If the Chinese Remainder Theorem is used, then doubling the key length increases the encryption time by a factor of 4 and the decryption time by a factor of 8.
4. **Quantum Computing Threat of RSA:** Shor's algorithm could factor in polynomial time and on big enough quantum computer it could do the same to huge numbers, potentially breaking the RSA encryption. This future threat motivates interest in quantum-resistant cryptographic algorithms.

There is a trade-off between security needs and performance limitations that has to be taken into consideration when choosing key lengths, both in light of current risks and the anticipated lifespan of the data under protection.

Known Attack Vectors

RSA implementations are targeted by a number of types of attacks:

This includes algorithmic attacks such as the General Number Field Sieve (GNFS) — factoring tries to guess the RSA modulus. It's the computational hardness of these attacks that underpins the key length recommendations.

1. **Timing Attacks:** Attackers can guess the private key through timing analysis which is measuring the amount of time taken to complete private key operations. These are so-called side-channel



Notes

attacks that exploit implementation-specific behavior in place of mathematical vulnerabilities.

2. **Power Attack:** Just like timing attacks, these also measure power consumption during cryptographic operations that can be used to make inferences about keys.
3. **Fault Injection:** Intentionally introducing defects during computation and observing the results can provide important insights.
4. **A Bleichenbacher Bleichenbacher's Attack:** This attack is based on properties of PKCS#1 v1.5 padding to disclose encrypted content in a controlled way through oracle queries.
5. **Example:** Common Modulus Attack — If the same modulus is used across different public exponents and the corresponding private keys are compromised, the other private keys using that modulus can be computed.
6. **Modulus and Private Exponent Attacks:** If the modulus and/or private exponent d is too small (for example, less than $n^{0.292}$), the system may suffer from attacks that will recover the private exponent d , without the need to factor n .

All Your Broadcasting Are Belong To Us: We could use low exponents (especially $e=3$) and low message-blocks when encrypting a given plaintext to multiple recipients with the same RSA modulus, which means we can recover each of those plaintexts by solving the CRT (Chinese Remainder Theorem) on the message-blocks.

Implementation Best Practices

There are some best practices that must be respected for a secure RSA implementation:

Beyond the standard cryptographic best practices, here are some specific things to keep in mind: Random Number Generation: Ensure that you are using a cryptographically secure random number generator for selecting primes and generally any randomness in the implementation

1. **Use of Padding:** Modern padding schemes like PKCS#1 OAEP (for encryption) and PSS (for signatures) should be used as they avoids the vulnerabilities associated with older schemes.
2. **Using constant-time:** algorithms for cryptographic operations to prevent timing attacks, as well as other side-channel protections based on the environment in which the data is deployed.
3. **Validate Parameters:** Ensure all the keys are generated according to security requirements and inputs to cryptographic operations to avoid attacks associated with it.

Together with suitable key lengths and proper implementation practices, RSA offers strong security for nearly all current applications, and appropriate long-term planning to switch to post-quantum alternatives alleviates long-term worries.

Practical Applications and Deployments

Its public and private key mechanisms are the backbone of many security protocols and applications over a myriad of usages. Learning how these keys are used in life brings their true meaning to light.

RSA in Transport Layer Security (TLS)

TLS, the protocol that secures most internet communications, relies heavily on RSA:

Certificate Authentication — RSA public keys in X.509 certificates are used to authenticate servers (and occasionally clients) during the TLS handshake. Certificates are signed by the certificate authority's RSA private key, creating a chain of trust.

1. **Traditional Key Exchange:** In earlier TLS variants, the premaster secret was immediately protected on key exchange via RSA encryption. The client encrypts this secret, and only the server's private key will decrypt it.
2. **Modern Usage:** In modern TLS implementations (TLS 1.3), ephemeral Diffie-Hellman key exchange is the preferred mechanism to achieve forward secrecy, with RSA primarily being



Notes

used for authentication rather than key exchange. The ephemeral parameters are validated by RSA signatures.

Certificate Transparency is a mechanism that enforces all SSL/TLS certificates to be logged publicly in a structured way, it is intended to make it easier to notice any unauthorized certificates and to improve the PKI ecosystem behind hours later for RSA certificates

Email Security Protocol

RSA underpins the major protocols of email security

1. **S/MIME (Secure/Multipurpose Internet Mail Extensions):** Standard that uses RSA to provide digital signatures (message authenticity, integrity), and encryption (confidentiality). S/MIME certificates issued by organizational or commercial CAs contain RSA public keys.
2. **OpenPGP/GnuPG:** Since its inception, the OpenPGP standard built its implementations around RSA for signing and email encryption, under a web-of-trust rather than a hierarchical PKI. Unlike Trust Route (the successor to Trust Scheme) that generates RSA key pairs for users, here users generate their RSA key pairs themselves and send them to each other through key servers or other means.
3. **DKIM (DomainKeys Identified Mail):** Email servers can use RSA signatures to cryptographically associate messages with domains that send them, allowing recipients to verify authenticity and fight spoofing as well as phishing

Document Signing and Verification

Legally binding electronic documents built on RSA digital signatures provide the cryptographic foundation:

1. **PDF Digital Signatures:** The PDF standard includes support for RSA signatures to authenticate documents and verify their integrity [93]. These signatures are visual representations and timestamp attestations.

2. **RSA and Data Integrity Code:** Signing is the packaging of executable code by software developers with RSA private keys so that operating systems and browsers can determine the source of the code and verify that it has not been changed before code execution. This prevents malware dissemination and modification.
3. **XML Digital Signatures:** XML-DSig specifies the use of RSA signatures for signing portions of XML documents, supporting non-repudiation for XML-based transactions and web services.
4. **Long Term Verification:** It is not uncommon for documents to have their validity or origin validated years or even decades after they have been signed; to accommodate this, advanced electronic signature formats (such as PAdES, XAdES, or CAdES) use timestamp tokens and validation data to ensure verifiability beyond the expiration of the certificates.

System for Identity and Access Management

RSA keys are used in the leading authentication and authorization systems:

1. **SSH (Secure Shell):** This common protocol used for secure remote system administration also extensively utilizes RSA key pairs. A user signs in by demonstrating possession of the private key associated with a public key the user previously registered with the server.
2. **Smart Cards and Hardware Tokens:** These physical authentication devices usually create an RSA private key internally and preserve it, never revealing it even to the systems they authenticate to, leading to a significant hardening of the authentication system.
3. **FIDO/WebAuthn:** These state-of-the-art authentication standards support RSA key pairs (as well as other algorithms) for passwordless authentication, and private keys are stored securely in hardware authenticators.
4. **Enterprise identity providers (IdPs)** utilize RSA signatures in SAML assertions and OpenID Connect tokens as a means of securely communicating authentication status across enterprise boundaries with Single Sign-On (SSO) systems.



Notes

Introduction to Cryptocurrency and Blockchain Applications

Blockchain technologies utilize RSA and other analogous public key cryptosystems:

1. **Digital Wallets:** Although Bitcoin itself utilizes ECDSA over RSA, certain blockchain solutions and digital wallets may incorporate RSA for the management of keys and the signing of transactions.
2. **Smart Contract Authentication:** Some blockchain platforms employ RSA signatures to authenticate off-chain data entries and third-party authorizations.

The RSA_KE scheme is also utilized in many decentralized identity systems by using RSA signatures to issue verifiable credentials independent of any central authority.

Library Implementations

This is a reasonable approach and many mature cryptographic libraries do provide an RSA implementation:

1. **OpenSSL:** This widely-used open-source library provides complete support for RSA, including key generation, encryption/decryption, and signature operations, with a variety of padding schemes.
2. **RSA – Bouncy Castle:** This library supports Java and C# and includes customizable RSA implementations for multiple standards and formats.
3. Microsoft CryptoAPI and CNG Windows platforms provide native RSA implementations via these programming interfaces.
4. **libsodium (primarily oriented against newer algorithms):** RSA is supported in this library for backward compatibility.
5. **Hardware-accelerated implementations:** Many modern processors provide specialized instructions to accelerate RSA operations, with the appropriate library support to harness these features.

These practical implementations showcase RSA's versatility and lasting relevance in the cryptographic landscape, despite the development of more specialized algorithms for certain applications.

Trends and Developments Going Forward

The domain of cryptographic keys is ever-changing, adapting to technologies, threats, and new needs. The future of RSA, and public key cryptography more generally, is being driven by several significant trends.

Transition to Post-Quantum Cryptography

Quantum computers, by virtue of Shor's algorithm, can efficiently factor very large numbers making RSA security unreliable. This has led to extensive work on quantum-resistant counterparts: NIST Post-Quantum Cryptography Standardization The National Institute of Standards and Technology has been reviewing proposed algorithms for standardization. These include lattice-based, hash-based, code-based, and multivariate cryptographic systems.

Hybrid approaches: In attempt to bridge the gap, many systems are implementing some hybrid schemes combining traditional RSA and post-quantum algorithms in the transfer period. This ensures backward compatibility as well as protection against future quantum threats. Migration Planning: Organizations are planning their migration strategies—moving from RSA to post-quantum algorithms—which include dealing with the challenges of key size, performance characteristics, and integration with existing systems. Quantum Key Distribution: QKD technologies do not directly replace RSA, but they provide physically secure key exchange protocols based on quantum mechanics instead of computational hardness assumptions.

Dedicated Hardware and Implementation Technologies

A huge game changer for RSA implementation is hardware innovations:

1. Trusted Execution Environments — Hardware-protected environments are also provided by secure enclaves such as Intel



Notes

SGX, ARM Trust Zone, and AMD SEV, which offer strict isolation from potential compromises on the rest of the system.

2. **What You Should Know: Custom ASICs and FPGAs:** Custom ASICs and FPGAs (Field Programmable Gate Arrays) can be used to create specialized hardware implementations of cryptographic operations, leading to performance gains at the likely expense of increased susceptibility to side-channel attacks.
3. **Homomorphic encryption**, not a direct replacement to RSA, is a method that allows computation to be performed on encrypted data without requiring decryption, which could significantly alter the means and how we use and manage private keys under distributed systems.
4. **Threshold schemes** allow dividing the operations concerning RSA private key between multiple entities, making it impossible for any single entity to retrieve or misuse the whole RSA private key.

The Evolution of Standards and the Regulatory Environment

There is ongoing evolution in the governance framework surrounding cryptographic keys:

1. **Next, Emerging Compliance Mandates:** Compliance regulations such as GDPR and CCPA, as well as rules in specific industries, are placing specific mandates on encryption key management.
2. **Standardizing Key Formats:** Further evolution of key encoding standards such as PKCS#8, JWK (JSON Web Key), and X.509 for new algorithms and use cases.
3. **Digital Signature Legality:** Additional legal frameworks that recognize the validity of digital signatures, such as more concrete clarification on acceptable key types and management practices.
4. **Export Controls:** Changes to export control requirements for cryptographic technology with the potential to impact harmonization activities verilocated.

Threshold Cryptography and Multiparty Computation

Advanced cryptographic methods are changing the way in which private keys are handled:

1. **Threshold RSA:** Techniques that enable multiple parties to collectively compute RSA private key operations without reconstructing the complete private key.
2. **Secure Multiparty Computation (MPC) :** Protocols allowing several parties to jointly evaluate a function over their inputs while keeping those inputs private — with application to distributed key management.

Fresh Shares of Distributed Keys: Preventing Slow-Gradual-Compromise through Proactive Security

RSA keys may one day be augmented or even replaced by other approaches for many applications discussed, but the underlying asymmetric key pair principle will remain a comfortable basis for secure communications and transactions. Organizations need to be aware of these happenings to guide their decision-making with respect to investments in cryptographic infrastructure and transitioning plans. The mathematical basis of RSA cryptography, a public and private key pair, make secure communications possible in situations that would otherwise lead to difficult secure key exchange problems. Properly generating, managing, and applying these keys underlie confidentiality, integrity, authentication, and non-repudiation in millions of installations globally. The mathematical relationship between these keys — one you can share with everyone, the other you guard like a state secret — creates a potent asymmetry that has transformed how we think about cryptography. This asymmetry allows both secure message exchange and the digital signatures that cryptographically bind identities to those messages. From TLS-protected websites to digitally signed software, modern infrastructure for secure communications is built on this fundamental idea. As computing technologies progress (especially the upcoming of quantum computers that can threaten the mathematical basis of RSA), the world of cryptography is constantly changing their best practices for key management and proposing new algorithms. Despite this inauspicious start, the conceptual leap of asymmetric key pairs is just as important today as when it was originally conceived, and has established a



Notes

framework that will exist as long as specific implementations are formed. The story of RSA, from an abstract notion to a ubiquitous security technology highlights the significant influence mathematical advances can have on real-world security. The properties, relationships and best practices for managing public and private keys are essential knowledge for security professionals implementing cryptographic systems. As new cryptographic approaches are developed, the basic principles RSA contributed to how to generate, distribute, protect, and use keys remain best practices in information security..

Multiple Choice Questions (MCQs)

1. **Which of the following is a property of a cryptographic hash function?**
 - a) Collision-free
 - b) Centralized control
 - c) Editable hash values
 - d) Slower computation speed
2. **A hash pointer is used in blockchain to:**
 - a) Store private keys
 - b) Link blocks securely
 - c) Encrypt transactions
 - d) Replace consensus mechanisms
3. **Which cryptographic technique ensures data integrity and authentication?**
 - a) Symmetric encryption
 - b) Digital signatures
 - c) Plaintext encoding
 - d) Data compression
4. **Public Key Cryptography uses:**
 - a) A single shared key for encryption and decryption
 - b) Different keys for encryption and decryption
 - c) Hashing to store transactions
 - d) Only symmetric keys for security

5. **Which encryption algorithm is most commonly associated with blockchain?**
 - a) RSA
 - b) AES
 - c) MD5
 - d) SHA-256
6. **In RSA encryption, what is used to encrypt a message?**
 - a) Private key
 - b) Public key
 - c) Hash pointer
 - d) Digital certificate
7. **The primary function of a digital signature is to:**
 - a) Encrypt messages for secure transmission
 - b) Verify authenticity and integrity of messages
 - c) Store blockchain transactions
 - d) Generate random keys for encryption
8. **Which of the following is true about the RSA algorithm?**
 - a) It uses only symmetric key encryption
 - b) It is a type of public key cryptography
 - c) It requires two identical keys for encryption and decryption
 - d) It does not use prime numbers for key generation
9. **What ensures that a cryptographic hash function is "collision-free"?**
 - a) It generates different hashes for different inputs
 - b) It allows two inputs to produce the same hash
 - c) It provides a fixed-length key for encryption
 - d) It enables faster decryption of messages
10. **What is the main purpose of using a hash function in blockchain?**
 - a) To store financial data securely
 - b) To verify data integrity and prevent tampering
 - c) To generate random numbers
 - d) To encrypt transactions for faster processing

Short Answer Questions



Notes

1. Define a hash function and explain its importance in blockchain.
2. What does "collision-free" mean in the context of hash functions?
3. Explain the role of hash pointers in linking blockchain blocks.
4. How does a digital signature work, and why is it important in security?
5. What is public key cryptography, and how does it differ from symmetric encryption?
6. Briefly describe the process of RSA encryption and decryption.
7. What is the role of public and private keys in securing blockchain transactions?
8. Why is SHA-256 widely used in blockchain networks?
9. What is puzzle-friendliness in cryptographic hash functions, and why is it important for proof-of-work?
10. How does public key encryption prevent unauthorized access in blockchain systems?

Long Answer Questions

1. Explain the properties of cryptographic hash functions: collision-free, hiding, and puzzle-friendliness.
2. Describe how hash pointers contribute to blockchain immutability.
3. How do digital signatures ensure the authenticity and integrity of blockchain transactions?
4. Explain the principles of public key cryptography and its applications in securing communications.
5. Describe the working of RSA encryption, including key generation, encryption, and decryption steps.
6. Compare and contrast symmetric encryption and public key encryption in terms of security and efficiency.
7. Discuss the importance of hash functions in Proof-of-Work (PoW) blockchain consensus mechanisms.

8. Explain how public and private keys work together in securing cryptocurrency transactions.
9. Describe how cryptographic primitives like hash functions and digital signatures prevent fraud in blockchain networks.

Analyze the potential weaknesses of RSA encryption and how modern cryptographic techniques address these challenges

MODULE 3

BITCOIN BASICS

3.0 LEARNING OUTCOMES

By the end of this Module, students will be able to:

1. Understand the fundamental concepts of Bitcoin and its significance in the cryptocurrency world.
2. Explain how Bitcoin works, including transactions, mining, and security.
3. Describe the creation of coins and tokens in the Bitcoin ecosystem.
4. Analyze the process of sending Bitcoin payments and the risks of criminal activities.
5. Understand Bitcoin governance and its decentralized control mechanism.



Unit 7: Introduction to Bitcoin and Its Working

3.1. Introduction to Bitcoin

Bitcoin was introduced in 2009, it was the first-ever cryptocurrency published by an anonymous person known as Satoshi Nakamoto with a whitepaper named "Bitcoin: A Peer-to-Peer Electronic Cash System." Like this: In response to the 2008 global financial crisis, times when confidence in conventional banking systems and centralized financial institutions had considerably cracked down, this transformative digital economy was designed. Bitcoin was created as a decentralized currency alternative to traditional state or local currencies that can be issued by governments, avoiding intermediaries like banks or payment processors. The primary innovation brought forth by Bitcoin is the technology that underpins it: the blockchain, a system of distributed ledgers that retains a secure and transparent record of all the transactions all over a network of computers. This allows decentralization that prevents any one group from controlling the network making it censorship and manipulation resistant. These are the basic features of Bitcoin: decentralisation, limited supply and pseudonymity. Bitcoin are not like the regular currencies which is issued by the central banks. This decentralization means that only network participants (called miners) verify Bitcoin transactions instead of a central clearing house. Another hallmark feature is Bitcoin's algorithmically predetermined supply cap of 21 million coins, creating digital scarcity designed to prevent inflationary pressures that impact fiat currencies. Moreover, while the Bitcoin ledger is public, its records are not tied to a person's identity but rather a digital address, granting some level of anonymity to users. It is also possible that technological innovations will eventually make other forms of crypto much more viable than bitcoin (potentially not only in quantitative comparability/measurement but also potentially in qualitative aspects). For many, it's a kind of "digital gold," a store of wealth that exists outside the traditional financial system. This view has been bolstered by Bitcoin's long-term upward price trajectory, despite extreme volatility. Some consider it a way to conduct transactions quickly and internationally without relying on traditional banking systems. It is especially useful in areas lacking banking infrastructure, or in nations facing financial crisis and hyperinflation. In addition, Bitcoin provides financial independence, giving people total

authority over their financial resources without dependence on third parties. Bitcoin usage has transitioned dramatically since its creation. Although Bitcoin was first popularised among tech enthusiasts and privacy advocates, as prices climbed and the potential use cases became more evident, it made its way into the limbs of larger segments of the population. Large corporations and financial institutions that once shunned cryptocurrencies have started adopting Bitcoin in their businesses and investment portfolios. Governments worldwide have responded to Bitcoin, from the legalization of its use as legal tender (as El Salvador did in 2021) to the construction of comprehensive regulatory frameworks, to strict policies including blanket bans. This regulatory environment is constantly changing as governments struggle with the realities of this new asset class. Bitcoin Covered: Yet, Bitcoin still struggles with some challenges and critics even as it becomes more accepted. Bitcoin still suffers from significant price volatility, with the coin experiencing dramatic fluctuations to such an extent that it becomes essentially unusable as a day-to-day currency. There are also environmental concerns when it comes to the energy use of Bitcoin mining, sparking debates into its sustainability and attempts to shift towards more efficient (in energy usage) methods of validation. This gives the users and businesses potential legal risks in many jurisdictions due to regulatory uncertainty. Moreover, due to scalability constraints of the original Bitcoin protocol, several solutions, such as the Lightning Network, have been developed to facilitate rapid and low-cost transactions while preserving security and decentralization. Actually, Bitcoin symbolizes a new approach to the end of the world we think money. Bitcoin has opened up new possibilities for financial inclusion, economic freedom, and technological innovation by eliminating the dependence on trusted intermediaries and allowing value to be transferred directly peer-to-peer. Whether it becomes the revolutionary form of money its early enthusiasts dreamed of, or finds its place alongside many other distinct currencies and value stores in a creative financial ecosystem, Bitcoin has sparked a foundational rethink of currency, value and trust in the digital age.

3.2. Bitcoin Works

Bitcoin itself is basically a pressureless distributed peer-to-peer system through which you will without using any central authority, be able to send



Notes

and receive payments in a trustless way. The core technology used for Bitcoin is called blockchain, a sequence of blocks that hold transaction data and are measured in time. This chain is a public ledger of all Bitcoin transactions ever done. What makes this system revolutionary is its ability to achieve consensus among participants who don't inherently trust each other; which means that it frees the financial system from having to rely on some form of trusted intermediary such as a bank or a payment processor, to confirm and log financial transactions. This trustless consensus on the blockchain is accomplished through Proof of Work (PoW) mechanism. So this automaton, independent of its structure, has two particular components: first, specialized participants known as miners, which compete to solve shards of math problems consuming processing power to do the work. The first miner who comes to this solution is allowed to place a new block of transactions in the blockchain and receives 1.05 Bitcoin as a reward. This transaction confirmation process, which takes place roughly every ten minutes, accomplishes three things: It ensures the security of the network against attacks, issues new coins in a structured manner, and gives miners an economic incentive to lend their computing power to support the health of the network. Bitcoin is secured cryptographically using a pair of public and a private key. When a user creates a Bitcoin wallet, a public key (like an account number) and a corresponding private key (like a password) are generated. The public key generates the Bitcoin address where others send money, while the private key is kept a secret and signed to prove ownership without sharing the private key itself. This unique assembly of asymmetric cryptography makes it so only the legitimate owner of bitcoins is able to spend their coins, while everyone has the ability to check the authenticity of the transactions upon the public ledger. The cryptographic mechanisms that underpin Bitcoin make it incredibly hard to forge coins or improperly alter the ledger. And when a user makes a Bitcoin transaction, it is not instantly put into the permanent record. Rather, the transaction becomes part of a memory pool (mempool) of unconfirmed transactions that are waiting to be included in a block. Miners will pick transactions from this pool, choosing those with greater transaction fees, and place them into the block which they are working on mining. A transaction is confirmed when it gets included in a block, and the block itself is added to the blockchain. Every added block further secures the transaction, making it even more challenging

to modify or invalidate the transaction. The majority of businesses deem a transaction secure after six confirmations, which on average takes one hour. It is this confirmation process that ensures Bitcoin transactions are considered irreversible when they have been made a part of the blockchain.

The decentralized ledger that backs Bitcoin is maintained via network nodes—computers that run the Bitcoin software and store a full copy of the blockchain. Nodes in these networks communicate with each other to transmit transactions and blocks throughout the network, confirm newly received transactions and blocks against consensus rules, and discard those that break them. This decentralized architecture prevents any one entity from controlling or manipulating the Bitcoin network. If a malicious actor tries to change historical transactions or double-spend coins, then they can not be persisted by honest nodes that are abiding by the consensus rules. Resistance to censorship and tampering is among the most valuable properties of Bitcoin, especially in countries with unstable governments or financial systems. Such design choices come with intentional trade-offs in performance and capabilities. The most significant is the block size limit, which limits how many transactions can be processed within each ten-minute block. This limitation means the maximum throughput is around seven transactions per second, a far cry from legacy payment systems like Visa capable of processing thousands of transactions in the same time frame. Because this model can only process so many transactions per second, demand for processing power is rationed when transaction volume hits this upper limit; at which point users have to compete with one another for inclusion by bidding up their fees, resulting in massive cost spikes during times of excess demand. Various solutions have been proposed and implemented to tackle these scalability challenges, including protocol upgrades themselves as well as layer-two networks that build on top of the base Bitcoin blockchain with additional functionality. These solutions intend not to compromise Bitcoin's inherent features of both security and decentralization, but rather facilitate its effective use for practical transactions. Bitcoin technical architecture from security, to decentralization and to efficiency. Each piece has been crafted to create this balance with as little trust needed as possible. This adds up to a kind of money done by the machine under the rules of math rather than at the hand of people who can decide what happens, producing a level of transparency about how things



Notes

will behave that our monetary systems never provided. Although it is a process normally dominated by a central entity, the fact that it can be done in a way that does not require intermediaries, helps confirm that distributed ledgers are not only about cryptocurrencies even though like you said hundreds of currencies are built on this technology.

Unit 8: Bitcoin Creation, Transactions, and Governance

3.3. Creation of Coins and Tokens

New bitcoins are created vis-a-vis a process called mining, which is both the process by which transactions are validated and the mechanism by which controlled amounts of new currency are introduced into the system. Unlike fiat currencies which can be printed anytime by central authorities, bitcoin has a deterministic issuance schedule which can be mathematically derived and ensures new coins are created over time until a maximum supply of 21 million bitcoins is reached. This capped supply is built into the Bitcoin protocol, and is one of the key economic principles of Bitcoin. Miners are rewarded with newly minted bitcoins when they add a new block to the blockchain. This is a reward that started at 50 bitcoins per block when the currency was created back in 2009, and which is set to halve roughly every four years, during events that are referred to as “halvings.” The current block reward is set to 6.25 bitcoins, and it will eventually reduce to zero sometime around 2140 when the last bitcoin gets mined. Bitcoin mining has gone from something that could be done using a commodity personal computer, to an entire industrial process that requires specialized hardware. Early on, miners were solving the cryptographic puzzles utilized by the Proof of Work algorithm with central processing units (CPUs). As more people began mining, they switched to graphics processing units (GPUs) that were better suited for these calculations. These days the dominant devices in the mining industry are Application-Specific Integrated Circuits (ASICs) — hardware devices that are custom-made for Bitcoin hashing specifically, providing substantially more efficiency than general-purpose computing equipment. The evolution has vastly increased the total computational power (hash rate) securing the network but has also increased the barriers to entry to miners as individuals. Mining pools came as a solution: collaborators work together, pooling their computing power and receiving proportional rewards, allowing smaller miners to earn more reliable rewards than they could gain if they were to mine entire blocks alone. While Bitcoin was the initial creation, the cryptocurrency ecosystem has seen thousands of alternative coins and tokens (or simply put, 'altcoins') emerge with their own defining properties and methods of generation. Evidence: Alternative cryptocurrencies or apps often known as “altcoins”



Notes

often have a separate blockchain and native coin. Most of these were created by forking the Bitcoin code and adjusting a few parameters (e.g., block time, issuance schedule, or consensus mechanism). For example, Litecoin, which has much quicker transaction confirmations, and Monero, which provides advanced privacy features. Other endeavors such as Ethereum built altogether unique blockchain architectures tailored for specific use cases — in that case, running programmable smart contracts. The vast majority of these alternative cryptocurrencies follow Bitcoin's lead in that they have mined the coin, although many have adopted different consensus methods like Proof of Stake, where validators secure the network by committing funds instead of immutable computational work like Bitcoin. Tokens are a different type of digital assets that exist on top of existing blockchain infrastructures and do not necessarily require their own bespoke blockchain. The ERC-20 token standard is the most widely used standard to create your own unique tokens on Ethereum using smart contracts; These tokens can be almost anything — voting rights in decentralized organizations, access to utility in certain applications, even tokenized real-world assets, such as real estate or commodities. Unlike mining, the creation process for tokens differs significantly from the general standard and are usually generated via a smart contract that outlines its entire supply and method of distribution. This can take the form of an initial coin offering or ICO in which the organization sells tokens to early investors, or through other means such as an airdrop, in which tokens are distributed free to any existing holders of a cryptocurrency to help kick-start a network. While the ease of token creation has fostered unprecedented innovation, it has also resulted in closer scrutiny by regulators as many token distributions have similar features as securities offerings.

Different cryptocurrencies and tokens have widely-varying economic models and this has a dramatic impact on their behavior and utility. Bitcoin's hard and fixed supply mechanics create deflationary pressure as demand rises on a now infinite supply continuously over the years—eventually sharpening its arguments for being “digital gold,” or a store of value character. In contrast, several cryptocurrencies adopt inflationary models featuring no fixed supply only limited issuance rates and would instead focus on achieving stability and higher usability as a medium of

exchange. Token economics can use more sophisticated mechanisms such as burning (permanent removal of token from circulation), staking (locking tokens to receive rewards), or governance (using tokens to vote on protocol change). Such economic architectures engender distinct incentives for network participants and ultimately inform whether a digital asset operates fundamentally as an investment vehicle, a means of exchange, or a governance/utility token within a given ecosystem.

The landscape of creating cryptocurrency is moving on with innovations that can remedy issues in solutions of the earlier age. One major recent development is the rise of stablecoins — cryptocurrencies specifically designed for the purpose of maintaining a constant value, usually pegged to a fiat currency, such as the United States Dollar. These assets seek to leverage the best qualities of cryptocurrencies (borderless transfers, programmability) while including the price stability essential for everyday use. They achieve this peg through a variety of mechanisms: full collateralization with fiat reserves (as with USDC), over-collateralization with other crypto (as with DAI) or algorithmic adjustments to their supply. Central bank digital currencies (CBDC) are another notable trend; these are digital versions of national currencies issued directly by central banks. Compared to decentralized cryptocurrencies, CBDCs will be controlled on a central level, although they might borrow some features from blockchain technologies. These advancements illustrate how the original invention behind Bitcoin ignited a much broader transformation in the creation of digital assets, each tailored to optimize different fundamental use cases throughout the financial ecosystem. Here, the wave of changes around digital asset creation represent a fundamental rethinking of how value is created, distributed and exchanged in a digital world. The code-based, open, and auditable rules for issuance and transfers enabled by cryptocurrencies and tokens have opened fundamentally new mechanisms for economic coordination that could be hardly envisaged before. Whether it's Bitcoin's beautiful solution to digital scarcity, or the programmable flexibility of tokens, these innovations have challenged established assumptions around money and value and have the potential to reshape financial systems in ways that go far beyond their original ideas.



3.4. Sending Payments and Criminal Activities

The mechanics of sending Bitcoin payments are both technically advanced and reasonably user-friendly. “(To conduct a transaction, a sender requires the recipient’s Bitcoin address — a string of alphanumeric characters, or conveniently, a QR code encoding the same information. The sender then provides an amount to transfer and sets a transaction fee, which determines the speed with which the transaction will be settled by miners. The sender uses their private key to digitally sign the transaction, which allows them to prove that they own the funds without sharing the private key. After being broadcast to the network, the transaction goes into the mempool of unconfirmed transactions, waiting to be included in a block. The result is that during periods of network congestion, transactions with higher fees are prioritized, creating a market-based system where users are free to choose between speed and low-cost confirmation. Once a transaction is included by miners into a block that is added to the chain, it is confirmable (a payment is irreversible after n confirmations) and it becomes more and more secure as the chain grows. As a response, the Bitcoin ecosystem has gradually built and integrated an array of tools, services that enhance the payment experience. Wallets used to be command-line tools for Satoshi gurus, but now they are mobile apps with QR code connections, address books, and built-in fee estimation. Merchants can now use payment processors that will accept Bitcoin payments, and automatically convert them into fiat currencies, thus protecting merchants from volatility risk. Layer-two solutions such as the Lightning Network have been developed to overcome Bitcoin's scalability constraints, allowing near instantaneous micropayments with negligible costs by executing transactions off the main blockchain and only occasionally reconciling the final state on-chain. These advancements have progressively brought Bitcoin closer to the non-technical user base and made it more useful for everyday spending, although it is still considered less convenient than traditional payment methods. Privacy is commonly spoken of as a redeeming quality of Bitcoin, but the truth is less clear than most people realize. The all transactions of Bitcoin are recorded on public blockchain, that is visible to everyone, creating a permanent and transparent history of all activity. This design affords pseudonymity, but not true anonymity—operations are attached to addresses, not real-world identities,

but advanced analysis can frequently correlate addresses with real individuals, particularly when people interact with regulated exchanges that apply Know Your Customer (KYC) processes. One key aspect of an cryptocurrency is that the transaction details — sender, receiver, and amount — are typically public on the system's transaction ledger, and while truly anonymous cryptocurrencies are rare, many other techniques have been developed or are being developed, such as a coin mixing service that collect multiple users' transactions and try to randomize the linkage between Sending address and Receiving address, as well as specific cryptocurrency that focus on achieves private transactions, such as Monero, which is designed to ensure greater privacy (using methods benefit from more advanced cryptographic techniques, such as ring signatures, stealth addresses, and confidential transactions) over Bitcoin.

Bitcoin's cross-border transactions pose important ramifications for international payments and remittances. Conventional international wire transfers usually go through a few intermediary banks, charge high fees (in the range of 5-7% of the transfer amount) and take several business days to process. As opposed to this, Bitcoin transfers can be sent directly from the sender to the recipient, regardless of geographic location, and are charged with fees that are determined by network congestion — not by the value of the transfer — and confirmation can take minutes or hours, instead of days. This efficiency is especially useful for these payments, which are known as remittances—money transferred by migrant workers to relatives in their home countries—and total in the hundreds of billions of dollars each year. By lowering the price of these transfers, cryptocurrencies can increase the real value received by often vulnerable populations. But there is still a long distance to walk due to practical obstacles, it is absolutely necessary to have simple means for the transition between local and crypto currencies, compliant in various jurisdictions. Bitcoin's pseudonymous characteristic and its ability to transfer value globally and without third parties has, of course, led to illicit usage alongside legitimate use cases. Darknet markets — including the now-defunct Silk Road — used Bitcoin as their main payment method for illegal goods and services, whilst ransomware attackers typically demand payment in cryptocurrencies. These associations struck the public consciousness more than legitimate uses of Bitcoin. But blockchain analytics firms and law enforcement agencies have since developed



Notes

increasingly sophisticated tactics for tracking potentially suspicious transactions, and several high-profile criminal cases have been solved specifically because blockchain transactions, unlike cash, leave an immutable and public trail. Large crypto exchanges today impose rigorous compliance programs which involve KYC verification, transaction monitoring, and suspicious activity reporting. According to Blockchain analytics, the share of cryptocurrency transactions related to illicit use has been shrinking over time as legitimate applications have increased faster, although specific measurements depend on methodology. Regulatory Framework for Bitcoin Payments: A Global Perspective Jurisdictions have made largely different choices, with varying priorities and concerns. In some cases, governments have passed sweeping regulatory frameworks that legitimize cryptocurrency payments but also introduce compliance requirements, such as the EU's Markets in Crypto-Assets (MiCA) regulation. Others have followed more draconian paths, restricting or even banning crypto transactions altogether. Cryptocurrencies are regulated by a ton of different agencies in the United States, with the cryptocurrency exchanges being treated as money services businesses by the Financial Crimes Enforcement Network (FinCEN), under the Bank Secrecy Act, which requires them to have anti-money laundering programs in place and to report suspicious activity. In the meantime, the Securities and Exchange Commission (SEC) has claimed jurisdiction over many cryptocurrency offerings that it deems to be securities. This multitude of regulatory requirements exposes businesses to compliance difficulties and introduces uncertainty for users, especially in the context of cross-border transactions. Despite these challenges,

3.5 Bitcoin Governance

For the unchanged governance system of Bitcoin that is very different to a centralized system. Bitcoin has a decentralized governance system that is characteristic of its decentralized nature, in contrast to traditional organizations with established hierarchies and power to make decisions. Governance develops over time under the influence of diverse participants: miners, developers, node operators, users, and businesses using the Bitcoin network. In the very center of Bitcoin governance is the consensus mechanism. Since major changes to the Bitcoin protocol need consensus

among participants, it leads to a system that is inherently conservative and resists rapid or controversial changes. Such caution and difficult-to-change characteristics are a double-edged sword: While they preserve the core properties of Bitcoin, they can slow the network's response to new challenges or opportunities on the horizon. Bitcoin Core developers (because there is no official/real authority over bitcoin) govern decisions regarding bitcoin. This volunteer network of developers maintains the reference implementation of Bitcoin software. They can suggest changes via Bitcoin Improvement Proposals (BIPs), but they can't unilaterally adopt changes that will or be against the "rules" of the network. Their influence derives from technical skill and community trust, not institutional authority. Miners are another important stakeholder in Bitcoin governance. Now miners, investing resources to secure the network, receive the right to vote relative to the amount of hash rate they supply. This voting mechanism mainly goes through their choice of which software version they want to run, thereby deciding which rule set they want to enforce. Miners can signal their support for specific proposals — such as through version bits in blocks — during contentious debates about protocol changes. Another check on governance power comes from Bitcoin's full node operators — businesses and people that run Bitcoin software capable of independently verifying all transactions and blocks. These operators can reject rules that they disagree with simply by choosing which software implementation to run, regardless of the support of their miners. This leads to a check and balance type of system which ensures that no single group can cross the line during the governing process. Bitcoin's broader user community—investors, merchants and casual users—influence it via economic incentives. This creates market pressures that affect the governance decisions they make to favor versions of Bitcoin they are willing to value over those they are not. This inclusion of the economic element in governance debates works not in practice and theory, as different concepts — proposals that could generate theoretical technical support still need to be economically feasible to be able to work.

From a governance perspective, Bitcoin has undergone a few major tests in its life. The contentious debate over block size that ultimately resulted in the Bitcoin Cash fork in 2017 is an example of having such an approach to governance that can only work out for the better. What kicked off as a technical debate over transaction throughput turned into a fundamental



Notes

philosophical argument over Bitcoin's main function. Over time, relentless debate yielded monumental disagreement, but Bitcoin's design guided its governance such that no faction could unilaterally ram through changes; instead, the parties simply forked around one another, yielding two separate Bitcoin communities sharing a common past, but nothing more. In 2017, another governance route was demonstrated with the implementation of SegWit (Segregated Witness). The successful implementation of the protocol upgrade and accompanying chain split to make transaction data more efficient for processing required an impressive amount of coordination among developers, miners and users. By adding support for UASF, users running full nodes on Bitcoin were able to activate SegWit despite the lack of initial support from miners. More recent governance struggles have revolved around Taproot, a privacy and scalability upgrade that was activated in 2021. This is enabled by the "Speedy Trial" activation mechanism, which requires 90% of blocks in a difficult period to signal that they support the improvement. It meant more efficient consensus building while still being widely technically supported. Governance in Bitcoin (BTC) is often beyond the set of protocol rules and includes social norms, communication channels and informal leadership. Governance discussion takes place at online forums, social media, developer mailing lists, and conferences, though with different impact levels. These communication channels are distributed and thus harder to co-opt or centralize but can result in fragmented conversations and information asymmetries.

Bitcoin's founding principles, articulated in Satoshi Nakamoto's whitepaper, are immutable, which provides an anchor for the governance debate. Interpretations can and will differ, but these tenets of decentralization, censorship resistance, and security are touchstones by which proposals are judged. This philosophical foundation resists the gradual migration to technocracy via technical implementation, and it remains crucial to Bitcoin's identity as it evolves. Due to these limitations of Bitcoin governance, critics raise concerns. The system can be slow to react to challenges, potentially leaving vulnerabilities unaddressed for more extended periods. Fewer people will understand technical details, which might bias influence to experts in the domain. Finally, the absence of formal dispute resolution channels can result in long-standing disputes without

obvious means of settlement. Advocates say that those shortcomings are ways that these technologies are often working as designed. They argue that Bitcoin's stubbornness in the face of rapid change preserves its most valuable properties — dependability, constancy and resistance to capture by special interests. Seen in this light, the slow progress of Bitcoin governance demonstrates a preference for long-term stability rather than fleeting responsiveness. While remaining agnostic about the future of Bitcoin governance, I suspect that both formal and informal processes will continue to evolve. Solutions that have yet to even gain widespread discussion (drivechains, sidechains, layer-twos, etc.) will eventually show their faces at layer one, but only once the features being added via them have become "normal" and accepted by message-participants. Then, these approaches could enable more experimentation while maintaining the security and stability of the base protocol. The governance model adopted by Bitcoin is radically different from the ones employed by traditional organizations. It shows that central authorities are not necessary for coordination at planetary scale, although it has clear consequences in efficiency and decisiveness. As Bitcoin matures and gains in adoption and economic significance, its governance mechanisms will be put under the microscope and the stress test of history, potentially exposing weaknesses and strengths we can't see yet. Ours is a market less about initiatives and profits and more about values, morals, ideals and promises. As long as participants recognize that cooperation rather than conflict is more beneficial, the system can continue to develop by means of rough consensus instead of command. So this delicate balance of competing interests, technical constraints and economic incentives somehow creates a governance ecosystem that, in spite of imperfections, has so far allowed Bitcoin to survive and progress to overcome many challenges.



Unit 9: Bitcoin Security, Price Volatility, and Future Trends

3.6 Key Encryption

At the core of Bitcoin's security design are cryptographic keys, a mathematics-based system that allows individuals to create secure ownership over digital assets, exchange those assets with others and validate that they actually own what they say they own without trusted intermediaries. The implementation of public key cryptography in Bitcoin is one of the most revolutionary aspects of Bitcoin whilst being what solved double spend that impeded many other attempts at digital currencies before Bitcoin had been devised, and in doing so, created a system where users have full control over their funds. Bitcoin is backed by cryptography, and its cryptographic framework relies very heavily on elliptic curve cryptography (ECC) — the secp256k1 curve, to be exact. The most important aspect to know about private keys is that whoever has the private key (the “user” with the key) has full control of the bitcoins belonging to the matching public key. Brute force attacks become computationally impractical due to the vast keyspace—provided the random number generation process is genuinely random and secure. The private key is then used to generate a public key via a one-way mathematical function on the elliptic curve. This is a computationally trivial operation in one direction but essentially impossible to reverse. The public key is mathematically related to the private key, but is not able to be used to derive the private key, providing the asymmetric security that makes the system useful. This public key is then hashed by a few different ways (SHA-256 and RIPEMD-160) to create a Bitcoin address —the address that users share with others in order to receive funds. Hashing public keys and creating the Bitcoin address (the hashing process itself) is not done only for that, but has its own benefits. In addition to creating smaller and more user-friendly identifiers, it also introduces an extra layer of security that can help guard against weaknesses in elliptic curve cryptography. It's only when a public key is used to spend bitcoins that its full form (the address) is made publicly visible on the blockchain, so theoretically this would protect against quantum computing attacks which could otherwise compromise elliptic curve security. Digital signatures are another key part of the cryptographic structure underpinning

Bitcoin. When a user wants to perform a transaction, they sign it with their private key and create a digital signature. And this signature, created by the Elliptic Curve Digital Signature Algorithm ECDSA, is cryptographic proof that this transaction is signed by the owner of the funds without any proof of the owner's private key. With Bitcoin we iterate to the next step by checking that the signature matches the public key controlled by the address holding the funds.

This is a unique digital signature that is tied to the transaction being signed, as it includes a cryptographic hash of the transaction data. This design eliminates the risk of signature reuse or manipulation such that it is not possible to "spend" a signature (used to authorize a transaction) more than once or change it to spend a different transaction. Due to the mathematical nature of ECDSA signatures, they are effectively unforgeable without knowledge of the private key, thus serving as the basis of the trustless verification mechanism of Bitcoin.

The cryptographic infrastructure of Bitcoin has been developed further in response to security and usability problems. The hierarchical deterministic (HD) wallet, specified in BIP32 in 2012, was a major leap forward, allowing the user to generate a whole tree of key pairs from a single seed. It helped both in the security and convenience department, allowing for more effective backup strategies as well as improved privacy with address rotation. The BIP39 standard for seed phrase backups also greatly enhanced user experience with an easy to remember string of words corresponding to complex cryptographic seeds. Multi-signature technology improved Bitcoin's cryptographic options to allow for transactions that needed the signature of multiple private keys before the transaction could be authorized. This approach formalized in BIP11 allows for more complex security arrangements and systems of trust. Multi-signature addresses allow for distributed control over funds, requiring a specified combination of signatures (like 2-of-3 or 3-of-5) to execute a transaction. This led to a wide range of applications ranging from improved privacy to sophisticated business transactions without trusts. The Segregated Witness (SegWit) upgrade, implemented in 2017, changed how signatures were used in Bitcoin addresses, and it was relatively straightforward to solve the original transaction malleability problem that had made further second-layer



Notes

solutions difficult. SegWit compressed both security and scalability by separating the witness data (signatures) from the transaction identifiers. This rewrite made later innovations, such as the Lightning Network, possible, as transaction IDs are consistent regardless of how the input/output signatures are formatted. The Taproot upgrade earlier this 2021 year also activated to bring a lot of cryptographical improvements to Bitcoin. Taproot achieved improved privacy and efficiency (relative to the original ECDSA implementation) by incorporating Schnorr signatures. Key aggregation comes from Schnorr signatures as well; multiple signers can create one signature, and the result looks like any regular single-signature transaction. This property increases privacy by making complex transactions look like simple ones on the blockchain. Taproot also brought Merkelized Alternative Script Trees (MAST), a construct that enables transaction participants to specify multiple ways that the output can be spent but only reveals the root of the condition that gets applied. It allows mere hiding of unused conditions enhancing privacy and generating sparse proofs thus making transactions smaller as lesser data to be saved on the chain. Alongside Schnorr signatures, these changes are the biggest cryptographic improvement to Bitcoin since it was first launched.

Various theoretical and practical challenges threaten the security of Bitcoin key encryption. On the theory side, advances in quantum computing threaten elliptic curve cryptography in the long run. Quantum computers that deploy Shor's algorithm might be able to break the mathematical relationship between private and public keys. But the threat is still largely theoretical, with major practical challenges to building quantum computers capable of posing an imminent threat. In addition, quantum-resistant cryptographic algorithms are being created, and they can be added to Bitcoin if required. The main practical hurdles are more about key management than the cryptographic algorithms. For usecases that require high levels of security, this burden of securing the private keys rests solely on the user, imposing severe usability and security limitations. Cannot recover lost keys and compromise keys can result in immediate and permanent loss of funds. In response, a rich ecosystem of security solutions has emerged, from hardware wallets to hot and cold custody services, which differ along a spectrum of security, convenience and trust requirements.

Hardware Wallets are one of the most useful innovations of Bitcoin key security. These dedicated devices keep private keys in secure elements that are compartmentalised off from the computers they're connected to via the internet, massively limiting the attack surface for key theft. As a rule, they require physical confirmation of transactions, which is why they cannot be abused remotely, even if the computer they are connected to has been compromised. These devices were a big leap for making cold storage (private keys kept offline) accessible to non-technical users. An emerging alternative are social recovery systems, which address the key management challenge. These systems store the pieces of recovery information with trusted contacts, in order to restore access when primary keys are lost without providing any one contact means of gaining access. Such approaches are attempts to strike a balance between technical security and human needs, including forgetfulness and planning for mortality.

Institutional custody options have built out advanced, semi-independent state cryptographic systems, merging aspects of multi-signature tech with operational security. These are often based on the concept of distributed key generation protocols, meaning that private keys are never built at a single instance or location even when created. Time-locked recovery mechanisms, geographic distribution of signing devices, and governance policies encoded in smart contracts provide an additional level of security beyond the underlying cryptographic primitives.

Current open avenues of research in Bitcoin's cryptographic infrastructure include designing better usability approaches for secure key management, improving privacy for users while enabling compliance with regulatory entities, and ensuring that the system has a plan for future cryptographic vulnerabilities, should they arise. Research continues on other topics such as threshold signatures, which allow multiple parties to cryptographically sign transactions without the need to reconstitute entire private keys. Advanced cryptographic building blocks such as zero-knowledge proofs are already being researched to add stronger security and privacy guarantees, while unlocking new capabilities. The core cryptographic concepts that embody Bitcoin have proven incredibly resilient over a decade of constant operation and scrutiny. And despite the incentives of billion-dollar rewards to identify exploitable vulnerabilities, Bitcoin would seem to be fundamentally sound



Notes

in its cryptographic architecture. This history of security has generated trust in the security model of the system, which is an important factor in the steadily increasing acceptance of Bitcoin as a store of value and the use of Bitcoin as a medium of exchange. Bitcoin's cryptographic beginnings will probably undergo more, even better work as Bitcoin matures. However, changes to core security mechanisms move extremely slowly, as is typical for changes to Bitcoin, reflecting a conservative approach. Such slower, steady step-wise cryptographic progress—with a focus on security and backward compatibility rather than innovation by iteration—fits with Bitcoin's role as a sound money system, as where stability and reliability trump the plumage of feature sets. Bitcoin's cryptographic design is elegantly simple: A system, transparent and trustless, that is built through pure mathematics and distributed consensus, is created. Bitcoin, by ingeniously deploying cryptographic principles to enable secure peer-to-peer transactions without the requirement of trusted third parties, ushered in a paradigm shift in digital ownership. Its enduring functionality and increasing adoption are ever-repeating testaments to the capability and promise of applied cryptography to change financial infrastructure.

3.7 Bitcoin Price and Volatility

The old adage states that those who don't remember the past are doomed to repeat it epoch and Bitcoin certainly has a tumultuous one that would be the envy of many a best-selling author in the record books for financial markets this century and even beyond. Bitcoin price has grown from a penny in its early days to over \$60,000 today, going through many different stages of evolution from an experimental technology to an increasingly mainstream asset. When Bitcoin first appeared in 2009 it had no market price. Early transactions were primarily symbolic or experimental, with the most famous being the purchase of two pizzas in May 2010 for 10,000 bitcoins — an amount that would eventually be worth hundreds of millions of dollars. Only in 2009 was one new formal exchange rate named valuing Bitcoin at about \$0.0008 per coin, based on the electricity costs of mining. By July 2010 that price had climbed to about \$0.08 — still a tiny fraction of its future worth. The initial important bull market surfaced in 2011, when Bitcoin momentarily reached equality with the U.S. dollar before leaping to approximately \$32 in June, a 3,200% rise upon January 1. This meteoric

rise was countered by an equally meteoric fall, as prices dropped nearly 93% to approximately \$2 by November 2011. This pattern will become well-known in the history of Bitcoin prices, with parabolic increases followed by steep corrections, but generally succeeding cycles making higher lows. There were two separate bull runs in 2013. Its first saw prices rally from roughly \$13 in January to above \$260 in April, before retracing to about \$80. A second and much steeper surge began in and by December prices were above \$1,100 — briefly exceeding the spot price of gold. This event gained media attention and exposed Bitcoin to even more people who dismissed it as a speculative bubble. The later crash, compounded by the high-profile collapse of the Mt. Gox exchange in early 2014, sent prices below \$200 by January 2015. After several years of relative stability, Bitcoin entered another growth phase in 2017; Bitcoin started the year near the \$1,000 mark and experienced a phenomenal rally, reaching a peak of nearly \$20,000 in December. This bull run coincided with the rise of initial coin offerings (ICOs) and alternative cryptocurrencies, marking much of an asset boom in crypto as a whole. The follow-on correction was extreme, however, with Bitcoin shedding some 84% of its value by December 2018 to find a trough near \$3,200. This protracted downturn — alongside more accurately termed “crypto winter” — put investors to the test and resulted in considerable consolidation throughout the general cryptocurrency market. The COVID-19 pandemic represented yet another inflection point for Bitcoin. Bitcoin crashed initially, alongside traditional markets, of March 2020, before a phenomenal recovery that was propelled by unique monetary stimulus from central banks, institutional interest and mounting inflation fears. This rally continued into 2021, culminating in Bitcoin hitting a new all-time high of over \$64,000 in April, followed by yet another large correction. A second rally through 2021 also formed Bitcoin’s current all-time high near \$69,000 in November before a significant drop later in 2022.

There are a number of separate forces that move the Bitcoin price. Its market size is still small compared to that of more established asset classes, meaning it remains vulnerable to large orders that can stir violent price swings or sudden shifts in investor sentiment. Inelastic supply (fixed supply schedule, a max cap of 21mm bitcoins and a set issuance rate) that cannot respond to increased demand, consequently amplifying price fluctuations. The nature of Bitcoin has been a key driver for how price has moved over



Notes

time. Once considered merely as a novel form of peer-to-peer payment network, bitcoin started being seen as a “digital gold” or a hedge against the devaluation of money. This change in narrative towards this store of value proposition has appealed to different investor profiles and different investment horizons. Even the narrative of Bitcoin as a macroeconomic hedge asset became apparent during the COVID-19 pandemic, as monetary expansion by central banks across the world reached never-seen-before levels, which raised fears of currency debasement. The gradual adoption of Bitcoin has started to change the supply and dynamics against the price. The initial price movements were largely driven by individual investors and tech enthusiasts. However, the entry of institutional investors (notably MicroStrategy's treasury reserve strategy announcement in August 2020) resulted in new capital flows and investment timelines. Its later adoption by companies like Tesla, Square (now Block), and Massachusetts Mutual Life Insurance Company, as well as the introduction of Bitcoin-centric investment products, has further entwined Bitcoin with the wider financial ecosystem. It has often been regulatory developments that have proved catalysts for big price moves. Historically, news of regulatory crackdowns, especially from China, has caused massive sell-offs. On the other hand, strong regulatory clarity or acceptance in major jurisdictions preceded price appreciation, in many cases. The application of Bitcoin futures approved by U.S. Commodity Futures Trading Commission in 2017 and the approval of spot Bitcoin ETFs in early 2024 are regulatory milestones that broadened Bitcoin's accessibility to traditional investors. There's a redeployed market infrastructure, liquidity network that has been established, not only that, it has change in the velocity patterns. However, in the early days of Bitcoin, trading occurred on unregulated exchanges with illiquid markets that lacked basic security, which were prone to manipulation. Market efficiency has been incrementally improved through the advent of regulated exchanges, derivatives markets, trading desks and institutional custody services. Thanks to improvements in market depth, sizable single trades have less impact than before, although Bitcoin is still more vulnerable to liquidity shocks than traditional financial markets. However, as history has shown, the halving events (which halve the number of new Bitcoin issued into circulation, approximately every four years) have always come before major bull cycles. Because of this inelastic demand, these supply-side shocks have manifested

as cyclical price patterns for Bitcoin. Each of the halvings in 2012, 2016, and 2020 were preceded by substantial price appreciation at some point after, although over varying time frames and magnitudes. Although this pattern formed expectation of periodic returns, every cycle has its own unique properties that are a reflection of Bitcoin's life stage in the evolving market cycle. The correlation dynamics of Bitcoin with other asset classes have varied over time which makes its use in the context of portfolio construction complex. Bitcoin had been uncorrelated with traditional markets, but it has recently become more correlated (albeit loosely) with risk assets, especially in stress periods. During the COVID-19 market crash in March 2020, Bitcoin initially moved in lockstep with equities as investors searched for liquidity across every asset class. It then decoupled and marched higher, while old-school safe-haven assets such as gold lagged.

Bitcoin's evolution as an asset class: Volatility metrics show gradual but incomplete maturation. Measuring volatility historically via standard deviation of returns indicates that volatility has halved from its early years but is still much higher than traditional assets like equities, bonds or even gold. The implied volatility extracted from options markets is also higher than that of mainstream financial assets, showing the market expects continued large price movements. Yet in general, over the long haul, as market caps grow and equity bases spread, volatility decreases. Bitcoin volatility cannot be separated from psychological implications. The comparatively more complex nature of such technology compared to that of foregoing paradigms, coupled with the technological potential for a paradigm-shifting reorientation of monetary systems, results in extreme sentiment fluctuations of fear and euphoria. The “fear and greed index,” a metric comprised of various market indicators to measure investor sentiment, has shown extreme numbers throughout Bitcoin’s history, many times aligned with significant price inflection points. These psychological factors can lead to self-perpetuating cycles of buying or selling pressure divorced from fundamental developments. There is no consensus among analysts on any single methodology for valuing bitcoin, and the price discovery process is further complicated by the application of different methodologies that are being used to value bitcoin. Stock-to-flow models can be seen by calculating the current supply against the rate of new production and can imply extreme long-term price appreciation based on



Notes

Bitcoin's programmatic scarcity. Bitcoin value models based on networks, such as Metcalfe's Law, which suggests that the value of a network is proportional to the square of the number of users on a service, try to give a measurable value to the crypto by addressing adoption. Cost-of-production models measure fair value by computing the energy and capital expenditure required to mine. Both methods have predictive limitations and can lead to substantial estimation errors. A new price-forming entity — bitcoin-denominated derivatives market. Derived instruments like futures contracts, options, and perpetual swaps create leveraged exposure to Bitcoin price movement, which can exacerbate volatility in times of market stress. Funding rates in perpetual futures markets reflect trader sentiment and positioning and can trigger cascading liquidations when the leveraged positions are forced to close. The increasing sophistication of these derivative instruments has opened up new hedging opportunities as well as a new source of market complexity. The Capitol Control Is Implemented With Geo Variations In Bitcoin Price In the latter, you can see large premiums (or discounts) in jurisdictions with limited capital flows or regulatory tendencies (as in the case of the "Kimchi premium" in South Korea). Such differences in price create arbitrage opportunities from time to time, but they also highlight the incompletely globalized nature of cryptocurrencies and the idiosyncratic effects of local factors on Bitcoin's valuation.

Bitcoin, despite still being volatile, has still proven its resilience time-and-time-again after every major correction. Being able to recover from 80%+ drawdowns multiple times has strengthened the conviction of those holding long-term, most treating volatility as the cost of entry for further appreciation in the future. This belief has created a culture of investors that are long-term holders often referred to by the acronym "HODL" (Hold On for Dear Life), prioritizing long-term conviction over daily price fluctuations.

With Bitcoin growing up, a few conditions are likely to shape its future volatility profile. If adoption increases and traditional financial systems incorporate crypto proper, one could expect to see less extreme price action due to a feed of new buyers as well as market depth. Improved regulatory clarity could help stabilize by lowering policy uncertainty in the same way.

But the fact that Bitcoin has a hard supply limit and is increasingly being viewed as an inflation hedge means that some volatility will probably always be a fundamental property of the asset — especially whenever macroeconomic conditions turn sour or adoption reaches a new zenith. The very volatility that has characterized Bitcoin price action is both its biggest threat to mainstream adoption, and its biggest draw for many investors. Such price action mirrors the maturing even of Bitcoin itself, as it shifts away from experimental tech and towards a nascent financial asset class, with every cycle attracting new players, improving infrastructure, all while maturity slowly increases on a more macro scale. Whether this volatility will eventually converge to levels we are used to seeing in traditional assets is one of the most essential open questions in the continued development of Bitcoin.

3.8 Future of Bitcoin

Bitcoin on the futureIn numbers, it is a quintessential speculation that occupies an interesting crossroad of technological innovation — economic theory — regulatory evolution — social transformation. Bitcoin has now entered its second decade, but in this time of incredible growth around the world, the forces shaping its trajectory are so numerous that the path ahead is impossible to know. In short, any forecast of Bitcoin's future needs to and should admit uncertainty, while defining as best it can the macro structures, trends, and challenges that will likely shape its development. There are a number of innovative routes Bitcoin could take even off its current tech roadmap. The Lightning Network, a second-layer solution running on top of the Bitcoin blockchain, has made significant strides toward solving scalability issues. Lightning, by allowing instant, low-cost transactions via payment channels, potentially allows Bitcoin to be used as an everyday commerce tool without sacrificing base-layer security and decentralization. Development efforts have now shifted towards Lightning's reliability, user experience, and its integration with legacy systems, as evidenced by the introduction of channel factories that may be able to lower the on-chain footprint of opening channels.

Another major avenue for Bitcoin's technical progress, privacy improvements, is layer two. Although bitcoin transactions ARE



Notes

pseudonymous instead of anonymous per se, the increasing acknowledgement of privacy as a human right and practical need has led to the progress of a number of privacy preserving technologies. Improvements such as PayJoin and Taproot will enhance transaction privacy by obfuscating the various types of transactions when they occur, while concepts such as Confidential Transactions could make it theoretically possible to hide the amount of a transaction without undermining Bitcoin's verifiability of its supply. These shifts indicate not a blanket rule but a nuanced approach between valid privacy concerns and the need for transparency. Activated in November 2021, the Taproot upgrade set the stage for major technical enhancements. Taproot was a significant upgrade to Bitcoin, as it implemented Schnorr signatures and Merkelized Alternative Script Trees (MAST) to allow for increased programmability, added privacy and increased efficiency. Future protocols may build off of these foundations to facilitate more complex smart contracts and conditional spending mechanisms while still preserving Bitcoin's core security properties.

Key Takeaways The slow, measured pace of bitcoin protocol development implies that these kinds of improvements will come slowly and surely over time, with security and backward compatibility taking precedence over spitting out new features at a web scale. (Another Potential development vector here is drivechains and sidechains, this enables experimentation with features and functionality that does not endanger the security of the main bitcoin network. These approaches form independent but interrelated blockchains, each with its own set of rules, that can inherit things like smart contracts, faster settlement, or alternative privacy models while preserving a security relationship with Bitcoin. Variants of this have already been done, such as Liquid and RSK, but the trust assumptions for layer 2s to work are never as equal and general-purpose as the main Bitcoin network. This will have a great weight in the future of Bitcoin as the infrastructure of the mining continues to evolve. It has gone through a radical transformation from CPU mining at Bitcoin's inception to the world of ASIC chips and industrial scale today. This future may incorporate a wide range of solutions from energy-efficient hardware and novel cooling mechanisms to completely new consensus mechanisms that enable the Bitcoin security model to endure with reduced energy costs. Miner geographic distribution — which has changed substantially due to China's

mining ban in 2021 — will continue to drive both the security properties and political resilience of the network.

In conclusion, Bitcoin and renewable energy can make a lot of cool things happen. Even with increasing attention given to Bitcoin mining and its potential, one thing remains clear – Bitcoin mining will come to be seen as one of the best monetization strategies for stranded energy resources (versus (some) fossil fuels), uniquely able to respond to remote opportunities while operating other resources more efficiently. Early examples are miners making use of surplus hydroelectric power during rainy seasons, capturing naturally flared gas from oil fields, or acting as flexible load management for solar or wind installations. Though concerns over energy use continue, the mutually beneficial relationship with renewables could eventually change the environmental story around Bitcoin mining. This is the same reason for the adoption and mechanism of operation of Bitcoin worldwide, as the regulatory landscape is constantly changing. Jurisdictions have framed the issue differently, ranging from embrace to restriction, creating a complex patchwork the world over. The dominant tendency of established economies has been to regulatory accommodation—accepting the permanence of Bitcoin and building frameworks for compliance. And this trend of regulatory clarity, although it does impose some operational constraints, ultimately facilitates institutional adoption in that it provides legal certainty. The future of regulation is to be determined, and taxation, privacy concerns, environmental effects, and connections to traditional finance would likely be the main points of focus.

Perhaps nothing has a stronger impact on the trajectory of Bitcoin's future than the trend of institutional adoption. Big businesses, banks, and even a few government institutions have started adding Bitcoin into their day-to-day and balance sheets. This institutional interest takes many shapes, from direct treasury holdings to the creation of Bitcoin-focused financial products and services. The growing development of this institutional stack (custody solutions, insurance products, lending venues, investment vehicles, etc.) could give surging access and legitimacy to Bitcoin at the same time as also reducing its volatility more than the long run. The future of Bitcoin is present in challenges and opportunities through central bank digital currencies (CBDCs). While creating greater awareness around digital money



Notes

as a whole, national currencies in digital form might also compete with specific use cases for Bitcoin as governments around the world explore them. But CBDCs also showcase the differences between state-run digital currencies and Bitcoin's open, censorship-resistant vision. This differentiation could further bolster Bitcoin's value narrative as an apolitical, neutral monetary network, especially in regions where monetary instability or financial repression is prevalent.

The change of global money policy might highly affect the utility and adoption of Bitcoin. It remains to be seen if central banks' unprecedented of balance sheets (blown up since the 2008 crisis, and even more so during covid19) is indeed pernicious in the long run by causing inflation and debasing currency. If these trends persist, or if we witness a genuine sovereign debt crisis, it's likely that Bitcoin's absolute issuance schedule and relative immunity to monetary meddling will be increasingly appealing to individuals, institutions, and possibly state actors looking for a hedge against inflation or an alternative monetary system. Current geopolitical drama may well inject more volatility into the future of Bitcoin. With monetary competition heating up between the major powers, the neutral and borderless nature of Bitcoin could well prove to be a boon to it, with some even seeing it as a helpful neutral alternative to international settlements. Nations facing sanctions and looking to decrease reliance on the Western-dominated dollar-based financial system could find strategic value in Bitcoin's censorship-resistance, for example. But these same features might also prompt greater restrictions from governments worried about a loss of monetary control or evasions of sanctions. The demographics of bitcoin ownership and usage seem positive for long-term adoption. Younger generational cohorts continue to demonstrate greater awareness of and willingness to engage with cryptocurrencies than older generations. As digital natives gain ascendancy in economic terms, their general comfort with digital assets might speed the real-life adoption of such assets. But this generational advantage would have to be weighed against the pressing need to improve security and the user experience in order to crack the technical barriers that prevent wider engagement. The presence of different adoption patterns for Bitcoin (and cryptocurrency) in developing economies may also give insight into what it will look like in the future. Outside of speculation,

Bitcoin also has practical use cases in areas with unstable currencies, limited banking infrastructure, or restrictive capital controls. Nations including Nigeria, Argentina, and Venezuela are witnessing broad grassroots adoption fueled by short-term economic demand, not investment speculation. This drop in users may not be the last, but the continued development of mobile-friendly, low-bandwidth Bitcoin solutions could increase these use cases; indeed, it may even cement Bitcoin as a surrogate financial system in parts of the developing world that lack access to traditional banks.

Social and philosophical aspects of Bitcoin will continue to impact its evolution and adoption. Bitcoin represents certain values — decentralization, censorship resistance, individual sovereignty, monetary scarcity — that will ring true in different cultural or political contexts. Bitcoin continues to be cypherpunk but as it breaks free from the confines of the cypherpunk community, it has to navigate mainstream adoption, creating friction in governance but also representative of its emergence as more than a simple technocratic experiment but as a complex socioeconomics system. This shared ideology also differentiates Bitcoin from many other technological innovations, representing a unique resilience as this commitment runs deeper than purely an economic incentive. we provide an overview of the competitive landscape for Bitcoin vs other cryptocurrencies and how this impact Bitcoin's dominance and functionality Despite such developments, thousands of alternative currencies have popped up, but Bitcoin has remained at the top position as the largest cryptocurrency by market capitalization and the reference point of the broader ecosystem. Combined with Bitcoin's superior security, liquidity, and brand awareness, this “first-mover advantage” generates powerful network effects. The development of Bitcoin has also been shaped by competition, with other protocols providing pressure to motivate the elimination of certain limitations without compromising the fundamental bitcoin ethos.

Goodfare to bridge gap between crypto world and traditional finance These are crucial milestones in the journey toward mainstream financial integration, as Bitcoin ETFs gain approval in multiple jurisdictions, Bitcoin futures markets continue to expand, and Bitcoin acceptance as collateral becomes more widespread. It is also possible that this could significantly



Notes

boost the legitimacy and availability of Bitcoin, but it might also include new market dynamics and regulatory challenges. It is likely that the tension between Bitcoin's roots in a peer-to-peer world and its combination with intermediated financial systems will continue to shape its evolution over the coming years. Speculation exists on the long-term fate of Bitcoin ranging from a niche asset to global monetary standard, with many combinations in between. A "digital gold" narrative, where Bitcoin is seen as a store of value and inflation hedge, has been particularly effective among institutional clientele. Even more ambitious forecasts see Bitcoin acting as a global settlement layer for all international transactions or as a common unit of account widely used in commerce. Skeptical scenarios, on the other hand, imply Bitcoin is destined to be an ever-speculative asset with a hardly used utility outside of its community of fanatical believers. In practice, the actual outcome will likely be determined by complicated interactions of technological development, regulatory approaches, monetary conditions, and social factors. As the block subsidy decreases over numerous halving events, Bitcoin's long-term security model is forced to answer important questions. Relevance of consensus incentives becomes all the more so through block design, when the reward for winning the competition of finding a block ultimately turns as transaction fees. Dynamic blocksize, better fee markets and other related proposals to solve this challenge. The sustainability of Bitcoin's security model in an eventual fee-dominant future heralds one of the most critical technical and economic questions facing the potential long-term viability of Bitcoin. The energy usage of Bitcoin will continue to be an important debate topic concerning its development and adoption. Critics have progressively condemned the energy-intensive design of proof-of-work mining, while defenders stress Bitcoin's potential to subsidize renewable energy development and extract value from wasted energy resources. The sector's transition toward cleaner energy sources, finer energy efficiency and possibly new concepts for heat recapture may reshape this environmental story in a big way. How this balance between energy consumption and security is resolved is likely to shape not only public perception but also regulatory approaches. What Bitcoin is yesterday, today, and tomorrow along with its priorities of development and adoption will always shift based on its narrative and self-image. The evolution from "peer-to-peer electronic cash" to "digital gold" or "monetary network" is nothing more

than the natural reaction to both technical constraints and market forces. With hindsight, however, future narrative developments will likely reinforce other dimensions of Bitcoin's multiple utility, perhaps espousing monetary sovereignty, financial inclusion or Bitcoin as the underlying infrastructure of a more open financial network. These changing narratives impact development priorities, user expectations, and regulatory approaches. Even entering its second decade, Bitcoin has already surpassed existential risk at hands of normal financial services and continues evolving to face new opportunities and challenges. As the development process is open-source and decentralised, it has been proven to be very resilient and has ensured the protocol has continued to grow while not losing core value propositions. Whether Bitcoin will ever actually reach its highest potential visions or settle into a more modest place in the broader international financial environment, what it created — trustless digital scarcity — is a technological invention whose impact will reverberate for generations to come.

The prospects for Bitcoin, in the end, depend on a myriad of factors; technological advances, regulatory choices, macro-economic forces, and social sentiments, to name but a few. This essential ambiguity is itself aligned with the decentralized nature of Bitcoin — no single actor has the power to dictate its future trajectory. It will instead evolve through the collective actions of developers, miners, users, businesses, regulators, and other stakeholders, who will shape its path according to their respective priorities and worldviews. While this development process is at times inefficient, it affords Bitcoin adaptive capacity and resistance to capture that might be vital to its long-term survival and success.

Multiple Choice Questions (MCQs)

1. **What is Bitcoin?**
 - a) A centralized digital currency
 - b) A decentralized cryptocurrency
 - c) A government-issued currency
 - d) A type of stock
2. **How does Bitcoin ensure transaction security?**
 - a) Through a centralized banking system



Notes

- b) By encrypting data using private keys
 - c) By using a single central server
 - d) By requiring physical verification of users
3. **What is Bitcoin mining?**
- a) The process of generating Bitcoin by solving complex mathematical problems
 - b) The process of storing Bitcoin in a digital wallet
 - c) The method used to create Bitcoin wallets
 - d) The act of purchasing Bitcoin on an exchange
4. **Which of the following is true about Bitcoin transactions?**
- a) They are completely anonymous and untraceable
 - b) They are verified through a decentralized network of miners
 - c) They require physical signatures for validation
 - d) They can only be processed by banks
5. **What is the maximum supply of Bitcoin?**
- a) 10 million
 - b) 21 million
 - c) 50 million
 - d) Unlimited
6. **What is a private key in Bitcoin?**
- a) A publicly available code for transaction verification
 - b) A secret key used to access and control Bitcoin funds
 - c) A government-issued code for Bitcoin taxation
 - d) A code used for mining new Bitcoins
7. **Bitcoin governance is primarily managed by:**
- a) A central bank
 - b) A group of elected officials
 - c) A decentralized network of users and miners
 - d) The International Monetary Fund (IMF)
8. **Why is Bitcoin considered volatile?**
- a) Because its supply constantly increases without limit
 - b) Due to fluctuating demand, regulation, and market speculation
 - c) Because it is controlled by central banks
 - d) Due to its fixed price set by the government

9. **What is the role of key encryption in Bitcoin?**
 - a) To enable centralized control over Bitcoin transactions
 - b) To secure transactions using cryptographic methods
 - c) To prevent users from accessing their own wallets
 - d) To make transactions completely untraceable
10. **Which of the following is a major factor that could influence the future of Bitcoin?**
 - a) Government regulations
 - b) Advances in blockchain technology
 - c) Adoption by businesses and individuals
 - d) All of the above

Short Answer Questions

1. What makes Bitcoin a decentralized currency?
2. How does Bitcoin mining contribute to network security?
3. Explain how Bitcoin transactions are verified.
4. What are the risks of using Bitcoin for payments?
5. Describe the role of public and private keys in securing Bitcoin transactions.
6. How are new Bitcoins created, and what is the total supply limit?
7. What factors contribute to the price volatility of Bitcoin?
8. Explain the process of sending Bitcoin from one wallet to another.
9. What is the purpose of Bitcoin governance, and how does it work?
10. What are some potential future developments in Bitcoin technology?

Long Answer Questions

1. Explain the fundamental principles of Bitcoin and why it is considered a revolutionary financial innovation.
2. Describe how Bitcoin mining works, including the concept of Proof-of-Work (PoW) and its impact on transaction validation.
3. How does Bitcoin differ from traditional banking systems?



Notes

4. Discuss the role of key encryption in Bitcoin security and how it prevents fraud.
5. Analyze the risks and benefits of using Bitcoin for financial transactions.
6. Explain Bitcoin governance and the role of miners, developers, and users in maintaining the network.
7. What are some major factors affecting Bitcoin price volatility?
8. Discuss the potential impact of government regulations on the future of Bitcoin.
9. How does the creation of tokens and coins differ in the Bitcoin ecosystem compared to other cryptocurrencies?
10. Evaluate the future of Bitcoin: Will it replace traditional currencies or remain a digital asset?

MODULE 4

CONSENSUS

4.0 LEARNING OUTCOMES

By the end of this Module, students will be able to:

1. Understand the importance of consensus in decentralized systems.
2. Explain distributed consensus and its key properties.
3. Differentiate between synchronous and asynchronous systems in blockchain networks.
4. Explore various distributed consensus protocols and their mechanisms.



Unit 10: Fundamentals of Consensus in Distributed Systems

4.1. Need of Consensus

Distributed computing systems rely on consensus, the ability for entities to reach agreement on a value or state. In an increasingly interconnected digital landscape, where systems are distributed across geographic locales and computational nodes, the need for consensus has never been greater. Distributed systems, by definition, include numerous independent components that must collaborate effectively to accomplish shared goals. These components have independent failure modes (since they tend to run on different machines) and require mechanisms to coordinate their activities and retain a consistent view of the state of the system. Without effective consensus, these systems risk devolving into chaos, with each component potentially acting upon different and possibly contradictory data. At the heart of the modern digital infrastructure enabling everything from online banking to social media platforms is a heavy reliance on consensus across distributed components. Let's take a very basic banking transaction when money is transferred from one account to another. That's a relatively simple operation that involves multiple systems needing to agree on the validity of the transaction, what time the transaction should happen and whether it actually happened. At this point, if these systems cannot come to consensus, terrorists could create critical inconsistencies, either by one system believing the transaction occurred while others don't, and funds are potentially corrupted, or destroyed — thereby being incorrectly created or lost.

This requirement for consensus goes far beyond financial transactions. Consensus algorithms work in a similar way in cloud computing ecosystems, keeping data consistent across replicated databases, so that when individual nodes crash, the rest of the system keeps spinning. When a content delivery network needs to decide what the optimal routing should be for a bit of data to be delivered, they rely on consensus to deliver it as effectively as possible no matter what may be going on. From both OFIR systems and distributed file systems through the use of consensus, the use of consensus to give all storage nodes a coherent view of the state of files is known. Furthermore, blockchain technology has thrust the concept of

consensus into the public consciousness. Specialized consensus algorithms are used by blockchain systems allowing a distributed and unsynchronized network of participants to agree on the state of a shared ledger without any central authority. This innovation allows for new genesis of decentralized apps and cryptocurrencies that relies on a principle of distributed agreement.

The larger and more complex systems become, the more difficult it is to reach consensus. So, large-scale distributed systems have to deal with network latencies, partial failures and also, potentially even malicious actors trying to disrupt the consensus process. This led to building consensus algorithms and making many research to solve different needs and types of systems, each one of course with its own advantages and disadvantages. Consensus is also an inherent requirement because of the need for tolerating faults. Given that hardware fails, networks partition, and software bugs exist, you must build distributed systems such that they can transparently continue operating correctly as such problems arise. With consensus algorithms, systems can provide correctness and availability properties in cases where one or more of the components were to fail in the field as long as the components that did not fail can agree on a consistent state as well as continuing to process operations. The need for strong consensus mechanisms is further emphasized by security considerations. In scenarios where not all actors can be trusted, Byzantine fault-tolerant consensus algorithms document guarantees, despite malicious actions. These algorithms guarantee that honest parties can agree, in the presence of collusion by adversaries.

An example other than maintaining a replicated state-machine would be time synchronization across distributed systems, in this case, consensus is crucial. Often, in systems distributed across different geographic levels, we need to agree upon the ordering of events or share a synchronized clock. While clock drift, network delays, and other variances introduce challenges for accurate time synchronization across distributed systems, consensus algorithms such as NTP (Network Time Protocol) help these systems arrive at a sufficiently consistent view of time. Systems of record, human expression, and consensus In terms of business, the need of consensus translates openly to reliability, consistency and trust in terms of digital services. Distributed consensus is a critical technology for companies to



Notes

guarantee that their services are available and correct, despite network partitions or partial system failures. These services would be prone to outages, data inconsistencies and security vulnerabilities without reliable consensus mechanisms. With systems becoming more decentralized and autonomous, the need for consensus will become paramount. Networked self-driving vehicles, smart city infrastructure, and Internet of Things ecosystems also need sophisticated consensus mechanisms to coordinate actions and maintain safety properties. The result of these developments means that further work will undoubtedly continue to take place in the development and tuning of consensus algorithms and thus key future research topics for distributed systems.

4.2. The Nature of Distributed Consensus and its Properties

The distributed consensus is a fundamental concept in distributed systems that allows a group of distributed nodes or processes to reach agreement on a single value or state, despite challenges faced in distributed environments. Distributed consensus allows independent entities to agree on a reality, and offers the basis for ensuring reliable and consistent. Though this might sound simple, in practice, it gets exceedingly complicated because real-world systems face unpredictable network topologies, different processing speeds, and node failures. The theoretical study of distributed consensus began with the seminal work of Leslie Lamport in the 1980s, who introduced the notion of Byzantine fault tolerance and proposed early consensus algorithms. So the field has advanced a lot, and there are many algorithms customized to different models of system and specification. These algorithms make different assumptions, have different guarantees, or have different performance characteristics, but they all address the same fundamental problem of consensus among distributed agents.

One of the key characteristics of a correct and effective consensus algorithm is that it satisfies several properties. The first and arguably most fundamental is agreement (sometimes called consistency), which states that all correct nodes eventually all decide on the same value. This guarantees that the system will settle on a single view of the world as opposed to a mass of conflicting views. In the absence of consensus, different parts of the system might continue in misunderstanding of one another — which could

lead to divergent behavior and possibly catastrophic failure. Second important property is validity (or integrity) meaning that the agreed value must be one proposed by one of the participating nodes. This guarantees that the value being generated as part of output does not get infringed by the system to randomly create values and ensures that the consensus being reached also incorporates the actual values input to the process. By making sure that the result of the consensus is actually what the system needs and preventing meaningless or harmful values to be accepted. Termination: the third key property requiring all non-faulty nodes to decide on a value eventually. Any consensus algorithm that can never finish does not have any practical usefulness, no matter how correct it may be in other aspects theoretically. Therefore, termination guarantees that the system is progressing rather than getting stuck in a continual state of back-and-forth. This property is especially hard to ensure in an asynchronous system, where we don't have the luxury of making timing assumptions. These three properties—agreement, validity, and termination—are collectively known as the “consensus problem.” This kind of repetition has a name: an algorithm that achieves all three properties is said to solve this problem. But as was shown in the well-known FLP impossibility result (FLP refers to Fischer, Lynch, and Paterson), no asynchronous solution can provide all three properties if even a single process might fail. This is a fundamental limitation that has deep implications for the design of practical consensus systems. In addition to these main properties, real world distributed consensus algorithms typically try to provide other guarantees as well. In the face of different failures, algorithms were administered by fault tolerance being an important aspect to consider. Crash fault tolerance, for example, is concerned with the simpler case where nodes continue working until they stop, and Byzantine fault tolerance deals with the more difficult case where nodes may stop correctly working and start misbehaving in arbitrary, potentially malevolent ways. The latter is especially critical in open and decentralized systems in which participants cannot be completely trusted.

Another critical dimension of consensus algorithms is their efficiency, which can be evaluated with respect to the message complexity (the number of messages exchanged), time complexity (the number of rounds needed), and the space complexity (the amount of storage required). At such scales, the efficiency of consensus could become critical as distributed systems



Notes

grow to thousands or millions of nodes. Any algorithms with quadratic or cubic message complexity become impractical at scale in short order. Many practical implementations are also designed for liveness — that is, the proposed system must continue making progress even in the presence of adversaries. Such strong guarantees on liveness allow the system to reach consensus decisions and process transactions despite network partitions, diverse processing speeds or partial failures. This feature is valuable in systems where the availability is a critical factor. Many consensus algorithms have been proposed for different requirements. Introduction Paxos, as introduced by Leslie Lamport, is a robust consensus algorithm for achieving fault tolerance in a partially synchronous environment, and much has been constructed onto this though Paxos algorithm has been mostly production-ready theory. The RAFT consensus algorithm was designed as a more understandable alternative to Paxos that offered similar guarantees but is easier to implement correctly. Practical Byzantine Fault Tolerance (PBFT) builds on these solutions, allowing them to service Byzantine failures—but at the cost of increased complexity and overhead. In addition, blockchain systems have their own family of consensus algorithms like (i) Proof of Work (used by Bitcoin), (ii) Proof of Stake (adopted by Ethereum and other platforms), and (iii) Delegated Proof of Stake. Such algorithms are designed specifically for the characteristics of permissionless systems where users are free to enter and exit these systems, and where economic incentives are the driving factor for maintaining consensus. In recent years, the centralizing point has shifted to consensus algorithms that provide strong formal guarantees but are also practical in practice for large-scale systems. This has lead to hybrid approaches, where different parts of the application can use different consensus mechanisms, enabling trade-offs to be made for different constraints. The theoretical aspects of distributed consensus algorithms are not merely of academic interest — they have direct consequences on the reliability, performance and security of systems built upon these primitives. Strong consistency guarantees are essential for financial systems to avoid double spends and finally execute transactions. Content delivery networks tend to favor availability and partition tolerance, at the cost of strong consistency. Many database systems offer configurable consistency levels, allowing application developers to select the desired trade off between consistency and performance. The better we understand



Notes

and design consensus algorithms, the more important this is as distributed systems scale up in size and complexity. Research continues in this area, expanding the limits of what is possible, and developing algorithms with more guaranteed performance, better efficiency, and more tolerance to increasingly harsh environments.



Unit 11: Consensus Protocols in Distributed Systems

4.3. Synchronous Vs Asynchronous System

Synchronous vs asynchronous and the main dichotomy in distributed computing, which affects the design, capabilities and limits of consensus algorithms. This distinction focuses on the assumptions underlying timing and message delivery, with significant consequences for the type of guarantees that can be offered and the means for reaching consensus. Synchronous distributed systems have one thing in common that governs everything — Strong Timing assumptions. These systems work under three important guarantees: there is a known upper bound on the delay for sending a message, such that every sent message will arrive within some predetermined time frame, an upper bound on the time required for any step, meaning that computational tasks are guaranteed to complete at known times, and every process holds a local clock with a known and bounded drift rate from real time, enabling reasonable degrees of time synchronization across the system. These assumptions make the design of distributed algorithms much easier, because they give processes a reliable way to detect failure when they don't receive expected messages before the maximum possible delay has passed, the sender must have failed. Synchronous systems allow for a class of consensus algorithms that have strong guarantees but can be implemented quite simply. Using timing assumptions, it is possible to implement timeout from which failure detectors can tell with certainty whether a node crashed or if it simply delayed. This property allows the implementation of algorithms that can thus ensure both safety (nothing bad occurs) and liveness (something good eventually happens) in the face of a bounded number of failures. Classical synchronous consensus algorithms such as the Synchronous Agreement Protocol operate in distinct and well defined rounds of communication, where processes exchange values, reach a decision, and move to the next round simultaneously.

Yet, the assumptions underlying synchronous systems seldom hold in practice, particularly in large-scale distributed environments that span multiple geographic locations or run over the public internet. Message delays can exceed any predetermined bounds due to network congestion,

varying load conditions, maintenance operations, and a multitude of other factors. In a similar fashion, processing times can change aggressively depending on system load, garbage collection pauses or contentions over resources. These practical considerations render synchronous algorithms infeasible for many applications. Asynchronous distributed systems, in contrast, are agnostic to timing. In these systems, messages may be delayed for an arbitrary time (although they are guaranteed to arrive if the sender and the receiver remain alive), processes can take an unbounded amount of time to take steps, and the clocks can drift arbitrarily. This better aligns with the chaotic reality of real-world distributed systems, especially those running at scale or over unreliable networks. This asynchronous model poses significant challenges to consensus algorithms. In 1985, Fischer, Lynch, and Paterson published the famous FLP impossibility result, which proved that no deterministic consensus algorithm could guarantee termination in an asynchronous system during a run if even a single process might crash. But this basic limitation stems from the fact that a purely asynchronous system cannot tell the difference between a crashed process and one that is merely much slower than the other processes. If an algorithm waits forever for a response from a possibly crashed process, the algorithm can never terminate; if instead the algorithm attempts to carry on without that response, the algorithm might make a decision without vital information in the event that a process is simply slow. Although this is theoretically infeasible for general use, practical consensus algorithms for asynchronous systems have been introduced by either relaxing some assumptions or settling for probabilistic guarantees. Another family of algorithms is randomized consensus algorithms, which bring a slight twinge of chance into the picture allowing these algorithms to "break the FLP impossibility result" leading to termination with high probability instead of certainty. These algorithms have now been used in specialized domains that suit their unique properties to system requirements.

More commonly, practical systems follow a partially synchronous model, which is a compromise between the two extremes. It will behave asynchronously for some bounded but unknown period of time, after which it will behave synchronously. This is based on the practical fact that while network delays and processing times may occasionally wander outside the normal range during times of instability, they usually revert to predictable



Notes

ranges eventually. The partially synchronous model allows for the design of consensus algorithms that offer strong guarantees after the system has stabilized, but still have correct functionality (though not necessarily progress) when operating asynchronously. One of the most successful series of approaches to consensus in partially synchronous settings is the Paxos family of algorithms. These algorithms are safe under all conditions—even under completely asynchronous conditions—and live as soon as the system resumes stable operation. This property makes them suitable for real-world deployments where temporary network instability is accepted but not permanent. There are a number of variants, such as Multi-Paxos and Fast Paxos, which are optimizations for certain operational patterns, while still providing the same safety properties. The Raft consensus algorithm, which was designed as a more easily understood alternative to Paxos, also operates in the partially synchronous model. Raft separates the consensus problem into a handful of easier sub-problems: leader election, log replication, and safety. By splitting these elements apart and moving to a more intuitive state machine abstraction, Raft guarantees the same properties as Paxos but is easier to implement correctly. Byzantine fault-tolerant algorithms add yet another dimension to the synchrony spectrum. Practical Byzantine Fault Tolerance (PBFT) and its derivatives run in a partially synchronous model, but also deal with Byzantine failures where nodes can misbehave arbitrarily, even maliciously. This requires additional rounds of communication and more messages, since it is more difficult to coordinate consensus when we cannot trust participants. There is another approach to the challenges of distributed agreement: Blockchain consensus mechanisms. Proof of Work of Bitcoin ingeniously works around traditional timing by requiring an economic cost to participation, and delivering probabilistic not deterministic finality. And so, as more and more blocks are added to the chain, the chance of a transaction being reversed approaches, but never actually reaches, zero. Subsequent blockchain consensus algorithms, including Proof of Stake and its variants, take alternative approaches to the same core issues, but provide different trade-offs in energy efficiency, decentralization and security. There are significant practical consequences for system design based on whether you choose a synchronous vs asynchronous (or mixed) approach. With timing assumptions, synchronous algorithms usually can achieve consensus with fewer messages and

communication rounds than asynchronous algorithms can. To achieve this, they make various assumptions, which will, however, be broken, causing them to break down catastrophically. Asynchronous algorithms are more resilient about working in unpredictable environments, but they often lead to increased communication overhead and, during extremely unstable network conditions, asynchronous algorithms may not make any progress at all. Many modern distributed systems use a hybrid approach where they implement various consensus mechanisms to different components in their architecture depending on the needs. Time-critical operations may also use algorithms with stronger synchrony assumptions in more controlled environments, where timing guarantees can be provided. Less-critical background tasks may use fully asynchronous methods that sacrifice immediate consensus for eventual consistency. Different systems can be placed on a spectrum running from synchronous (each node can wait the same bounded period before responding) to asynchronous (there is no guarantee for a bounded response time), and many systems fall in between in terms of timing and the guarantees made by the consensus algorithms they use. This understanding becomes critical in the face of the growing scale and complexity of distributed systems, enabling us to design and reason about more resilient and efficient solutions to the consensus problem.

Visions of Consensus in Distributed Systems

The story of how all of this led to the rise of Proof of Work and the different snake oil that followed.

Introduction: The journey of consensus mechanisms has been one of theoretical evolution and practical development in distributed systems. As distributed systems progressed and tested against failure-prone environments, simpler strategies like centralized coordination or basic majority voting failed to create robust systems. Modern consensus originated in the theoretical underpinnings laid down in the 1980s with landmark papers solving the Byzantine Generals Problem and some early protocols such as the Three-Phase Commit. As distributed systems continued to grow larger during the 1990s and early 2000s, increasingly sophisticated algorithms were developed. The major breakthrough was



Notes

Paxos (Leslie Lamport, 1978) which provides strict consistency guarantees between a set of processes in a partially synchronous setting. Although theoretically elegant, Paxos turned out to be notoriously hard to implement correctly, resulting in simplified variants and alternative mechanisms. In 1999 the Practical Byzantine Fault Tolerance (PBFT) algorithm was proposed by Miguel Castro and Barbara Liskov, this allowed consensus in scenarios involving potentially malicious actors, albeit at a high performance cost.

The introduction of Bitcoin — a solution for the double spending problem — and the advent of blockchain technology in 2008 marked the next evolutionary leap, introducing consensus mechanisms specifically devised for open, permissionless networks, where participants could join and leave freely. This was a revolution in consensus: rather than achieving agreement through voting they used computational puzzles (such as those describe in CASPER above). This breakthrough allowed parties that didn't know or trust one another to reach consensus, the foundation for decentralized applications.

Practical realization of Consensus Algorithms

This diverseness is manifestation of various implementations of consensus algorithm, which take into account particular operational needs and limitations. In traditional enterprise environments where nodes usually reside in an assumed trusted perimeter, algorithms like Paxos and Raft are dominant. For example, Google's Chubby lock service which gives distributed locking for their infrastructure is built on an optimized version of Paxos. Likewise, etcd, the distributed key-value store that supports Kubernetes, uses the Raft consensus algorithm to ensure consistency among its replicas. Another area where consensus algorithms are critical is databases. (Distributed databases such as Apache Cassandra use lightweight consensus protocols only for operations (e.g., topology changes), and eventual consistency models for data operations (to maximize availability). Domain-specific systems such as Google Spanner combine consensus protocols with custom hardware (e.g., atomic clocks and GPS receivers) to achieve global consistency in transactions with rigorously specified guarantees. Blockchain systems have created their own breed of consensus

mechanisms. Bitcoin's Proof of Work mandates that participants (called miners) solve computationally-intensive puzzles, with the first one to do so receiving the right to add the next condensed block of information. This mechanism is highly effective against Sybil attacks but requires enormous energy. Ethereum's switch to Proof of Stake is a great step forward that chooses among block producers based on their economic stake in the system rather than computational work. Delegated Proof of Stake (game-in-fighting (forking)) and Practical Byzantine Fault Tolerance (tolerant against adversarialism) are specialized algorithms and there are others ensuring different characteristics on the scale between decentralization, performance and energy efficiency.

Because financial systems require significant consistency and fault tolerance, such systems often utilize complex consensus mechanisms. The Ripple Consensus protocol (XRP Ledger) has an interesting "consensus by default" mechanism in which nodes keep a list of trusted validators. Other global payment networks such as SWIFT use a specialized consensus protocol to maintain transaction integrity between its network of international financial entities. In contrast, content delivery networks and other distributed caching systems face distinct consensus theorems, typically trading off vertical availability and partition tolerance against strong consistency. Akamai and other similar systems use gossip-based protocols to spread around configuration changes and maintain a loosely consistent view of the state of the network across their geographically distributed nodes.

Consensus in Edge Computing and IoT Environments?

Distributed consensus problems have gained new importance due to the rise of edge computing and Internet of Things (IoT) ecosystems. These settings are resource-limited, have poor network connections, and consist of devices with various hardware capabilities, making classical consensus approaches unfeasible. Different layers of the edge tier would likely operate on different consensus mechanisms. Within a local cluster, lightweight protocols may be employed between edge devices, while heavier algorithms orchestrate activity between clusters and the cloud. Time synchronization is a specific concern in these surroundings, with protocols like PTP (Precision Time Protocol) forming the basis for temporally consistent performance.



Notes

Many IoT devices are extremely resource constrained, which compounds the challenges with IoT deployments. Due to these constraints, battery-powered sensors cannot tolerate the communication or computation overhead introduced by traditional consensus protocols—but the development of energy-aware consensus variants allows the use of both communication and computation while minimizing the resources consumed when achieving consensus. These types of protocols typically sacrifice strong consistency guarantees to allow devices to run longer between interactions; however, they do contain mechanisms to promote to stronger protocols when important decisions must be made. Fog computing is when devices are closer to edge capabilities from the cloud it has a layered approach to consensus. Fast local consensus within a fog node can be used for time-sensitive applications, while less time-sensitive decisions can use strong protocols that span multiple nodes. This multi-tiered decision system lets systems balance about how responsive they want to be versus how consistent they aim to be, based on what the application needs.

Autonomous vehicular networks pose an even more challenging domain for consensus since they interplay edge computing constraints with high-mobility and security-critical requirements. Connected vehicles require consensus protocols that operate under tight timing constraints and adapt to changing topologies as vehicles enter and leave range. Such systems are often governed by special algorithms, which integrate geographic proximity and relative velocity to their consensus mechanisms. Smart city infrastructure is faced with analogous challenges at an entirely larger scale. Coordinating traffic signals, public transportation, emergency services and utilities needs consensus mechanisms which are across different systems with heterogeneous capabilities and requirements. In these environments, hierarchical approaches that meld local autonomy with global coordination work well: individual subsystems are able to optimize for local needs while maintaining general coherence across the system.

All The Best and Formal Verification of Consensus

Since consensus algorithms are the bedrock of systems' integrity, there has been significant interest in methods of formal verification [6, 22, 13]. These methods of mathematics hope to give formal proofs that consensus

implementations meet their defined properties even in pathological cases, thus ruling out entire categories of potential bugs and attacks. One of the approaches to verification is known as model checking, wherein we explore all possible states of the system in order to guarantee that the desired properties always hold. Tools such as TLA+ (Temporal Logic of Actions) have been applied to verify complex distributed algorithms, including consensus protocols such as Paxos and Raft. Amazon Web Services has notably used TLA+ to prove that a heartbeat (a fundamental part of their infrastructure) was correct, catching subtle bugs they might not have found otherwise. An alternative is theorem proving, which uses formal logical systems to create mathematical proofs to demonstrate the correctness of algorithms. The Coq proof assistant was used in the DISC (Distributed Systems Consistency) project to build mathematical models of distributed consensus algorithms along with formally verified implementations that enforce these theoretical guarantees in practice. Such implementations serve as reference models against which real-world systems can be evaluated.

Whereas static analysis-based techniques provide guarantees only up to fixed points, runtime verification techniques monitor systems in operational settings to identify violations of consensus properties. Much of these approaches could identify problems coming from implementation details or from environmental factors not represented in formal models. Some specialized consensus monitors inspect system activity and trigger alarms (such as system probes) when they detect a possible consistency breach, enabling proactive recovery before a fault amplifies and manifests as a failure. Though these propositions are a step forward, verifying consensus systems formally is still difficult; the distribution of environments and the fineness of timing-dependent behaviors remain challenges. Most verification efforts target general properties of core algorithms, not full systems in their implementations, thus requiring traditional testing and quality assurance, whether user-level or developer-level. Filling in these gaps is an active area of research, with verified compilation being one of the potential techniques used to preserve throughout the implementation process useful properties that are proved.



Notes

Quantum Computing Consent

Well, stuck (although not so) in nature—since quantum computers might do it in totally different ways. Quantum consensus protocols exploit the phenomena of quantum mechanics, such as superposition and entanglement, to reach agreement using fewer messages or providing stronger guarantees than classical alternatives. Thus, theoretical threshold and reduction results with classical algorithms suggest that up to one third of the participants can be malicious without any faulty behavior affecting the performance, yet quantum protocols have been proposed that can tolerate a higher fraction of malicious participants. These protocols leveraging quantum communication channels to help effectively detect inconsistent messages and malicious behaviors. Research into quantum signature schemes has also inspired stronger authentication guarantees for consensus messages that secure against more powerful adversarial settings. Quantum key distribution provides one more advantage for consensus systems, which is information-theoretically secure communication channels that are even secure against quantum computer attacks. These channels can fortify consensus protocols, as they guarantee that even computationally unbounded adversaries cannot break message contents.

But there are serious challenges to achieving practical quantum consensus. Furthermore, current quantum states are characterized by limited coherence time and qubit count, which impose a limit on the set of physically realizable protocols. Quantum comm infrastructure is still early – very few quantum repeaters and networks are in place. The hybrid approach of exploiting a combination of classical and quantum components to perform specific operations that would have been proven infeasible in a classical approach is taking the most promising near-term directions by performing those secure operations on a limited number of quantum components while relying on an existing classical infrastructure for bulk communication.

Towards a Future of Distributed Consensus

Distributed Consensus is, however, a rich field with many areas still ripe for exploration and innovation, and we can only arrive at the future by revisiting the research — with contemporary eyes — for insights on where we should go. One such direction is represented by self-adaptive consensus mechanisms that autonomously adapt their parameters to the environment (network conditions, application needs, etc.). For instance, you can have dynamic algorithmic changes based on the environment: changing the communication pattern, the value of the timeout, or even increasing or decreasing consistency guarantees to maximize performance based on the environment in which they find themselves. As the number of blockchain ecosystems continues to grow, cross-chain consensus protocols are starting to appear, allowing secure interaction between different systems built on different foundations. They need to reconcile drastically different trust models and consensus mechanisms, reaching agreement even when the underlying assumptions diverge fundamentally. Polkadot and Cosmos, for example, are projects innovating solutions to this difficulty by building frameworks for these mechanisms of cross-chain communication and consensus. This includes zero-knowledge proofs, meaning that one party can prove to another that a statement is true without revealing any other information, which are going on to become more and more integrated with consensus mechanisms, where participants can verify that the statement is correct without being exposed. Crypto-based methods for privacy preserving consensus in use cases, including banking and healthcare systems, where the confidentiality of the data is critical. Zero-knowledge rollups employ these proofs to compress consensus knowledge, which can help alleviate some broader blockchain scalability issues.

Another important field of research is directly related to quantum-resistant consensus, designing algorithms that would be secure against adversaries with quantum computers. Indeed, as quantum computers become more sophisticated, the current cryptographic underpinnings of many consensus methodologies can become jeopardized and prompt a transition to post-quantum solutions. Beyond this, progressive projects are implementing quantum-secure signature schemes and encryption methods into their consensus mechanisms now, long before it is needed. For even stronger guarantees at lower cost, hardware-assisted consensus mechanisms use specialized hardware components (e.g., trusted execution environments,



Notes

TEEs). Intel SGX, ARM Trust Zone, and equivalent technologies allow the creation of protected enclaves, where consensus operations can be executed with stronger security guarantees, possibly simplifying protocol design and boosting performance. Such approaches need to weigh the benefits of hardware assistance for security against the risks of introducing single points of failure or manufacturer dependencies. This makes consensus a particularly highly complex frontier for dynamic environments with small and rapidly changing sets of participants. Emerging applications such as vehicular networks, drone swarms and mobile edge computing feature nodes that join and leave the network constantly, in contrast to the more traditional consensus algorithms that expect a relatively stable set of participants. Open research challenge remains of designing consensus mechanisms that can keep consistency in the wake of this churn whilst minimizing the reconfiguration overhead. When these research directions across different communities come together with the practical implementation experience of real-time systems, we expect that the next generation of distributed consensus will provide stronger guarantees, better performance, and more general applicability than existing approaches. These innovations will unlock new categories of distributed applications with improved reliability, security, and scalability, broadening the scope of consensus mechanisms in our evermore intertwined planet. Distributed consensus as a field lies at the crossroads of theoretical computer science, practical systems engineering, and novel technologies. The basic problem of consensual governance among autonomous actors continues to be accurate today as when it was first formalized decades ago even as the situations in which agreement has to be obtained became far more differently and complicated.

The distinction between synchronicity and asynchronicity in systems is only a single axis of the much more elaborate consensus landscape. The spectrum of trade-offs traversed by these systems in the wild must contend with scale, security, performance, and fault-tolerance simultaneously. And the impossibility results that restrict what is possible across different system models have not stopped the creation of useful algorithms, but rather guided researchers and engineers toward omissions in productive trade-offs and clever circumventions. In practical terms, these considerations point to a need for ongoing engagement with production-maintained systems, ensuring

that advances in distributed systems research are supported by real-world deployment and use. As systems evolve from centralised then, to increasingly decentralised and autonomous over time, the need for consensus protocols to be strong, efficient, well defined and understood only increases. The landscape of distributed consensus will continue to evolve, with increasing specialization of algorithms for specific use cases and environments, as well as efforts to combine different consensus mechanisms into hybrid models that draw upon the best elements of each. As, generic critical components will start getting formal verification to provide mathematical guarantee on their correctness. This ability of distributed components to come to a reliable agreement, in the presence of failures, delays and potentially malicious behavior, is not just an academic concern—it is the underpinnings of essential infrastructure in a world where digital systems underlie virtually every aspect of modern life. Consequently, grasping the principles, properties and practical dispositions of a distributed consensus therefore remains an invaluable tool for anyone participating in the design, development to the operation of robust distributed systems.

4.4. Distributed Consensus Protocol

One of the core challenges in any distributed computing system, establishing distributed consensus protocols underlies the basis of maintaining consistent state across nodes. At their most fundamental level, distributed consensus protocols allow for a system of independent computing nodes to agree upon a single value or a series of values, even in cases of failure, network partitioning, or even malicious behavior on the part of some parties involved. Application of these protocols are not limited just to the theory of computer science, but rather they have tremendous applications in distributed databases, blockchain networks and large scale cloud systems where system stability and consistency of the data are the main concerns.

Distributed consensus protocols have a rich history. The Two-Phase Commit protocol, for example, is one of the first approaches to agreement, but it does not deal with coordinator failures gracefully. The theoretical underpinnings were significantly clarified by Leslie Lamport's pioneering work on the Byzantine Generals Problem in the early 1980s, which introduced the key challenges of consensus when some nodes may be



Notes

exhibiting arbitrary faulty behavior. This came to fruition in the form of Byzantine Fault Tolerant (BFT) algorithms able to function properly despite a subset of the network misbehaving or acting maliciously. There was a significant breakthrough in this area by Lamport, whose consensus protocol, Paxos, was published in 1998, and provided mathematical guarantees of consensus in asynchronous settings, where messages may be delayed but not lost. Although it was theoretically beautiful, Paxos had some practical weaknesses due to the complexity of its implementation. In 2014, Diego Ongaro and John Ousterhout proposed a consensus protocol called Raft, which aimed to be more easily implemented than Paxos while achieving similar properties of fault tolerance and having a focus on understandability in its design. This was primarily in the context of crash failures, assuming nodes are on the right (i.e., some sort of trusted set of participants). Several basic properties drive the design of consensus protocols. This means that safety guarantees that all correct nodes must eventually agree on the same value—even as a compromise even in the face of failures or malicious behavior from other nodes, and this value must have been proposed by some correct node, thus ensuring that some correct node must propose a value that becomes part of the consensus outcome. Liveness properties: Ensure that the protocol eventually makes progress and does not get stuck in some deadlock-like situation where the system is unable to make decisions. This inherent tension between safety and liveness is fundamental in consensus protocol design, and the FLP impossibility result (after Fischer, Lynch, and Paterson) shows that if a protocol is deterministic, it cannot guarantee both safety and liveness if there is at least one node failure against an asynchronous system. Most recent distributed consensus protocols are based on state machine replication, in which multiple replicas port the same sequence of operations through the state machine to guarantee a same state between the replicas. This is commonly done by a leaders or coordinators who propose values and coordinate agreement, and sophisticated election protocols determine who assume the leadership roles when failures are detected. Higher Node Fault Tolerance Quorum-based approaches require that there be agreement among a subset of nodes, thus achieving the trade-off between key, value pairs updated data efficiency and fault tolerance. Based on different types of failures, the configuration of the quorum size directly affects the resilience of the system.

Consensus protocol design is driven significantly by performance considerations. Consensus execution in a network requires conditional commitments, each of which consists of one or more messages, and consensus is dependent on the communication complexity or the total number of messages exchanged, which will also impact the required network bandwidth and the overall throughput of the system. Time complexity refers to how fast consensus can be achieved and how it affects the latency of the system. This implies that the message size and count required for the message is complex based, leading to better convergence in terms of time but higher cost in terms of bandwidth. These trade-offs are most evident in wide-area networks or networks with low connectivity. The fault model assumptions inherent in a consensus protocol heavily dictate the design and capabilities of that protocol. Crash-fault tolerant (CFT) protocols consider scenarios where nodes may crash but never act in arbitrary or malicious ways. In these protocols, f failures can be tolerated as long as $2f + 1$ nodes are present. Byzantine fault tolerant (BFT) protocols are used in more adversarial environments where nodes can behave arbitrarily, such as by sending conflicting messages to different peers. Because of the heightened difficulty of resisting malicious behavior, a stronger fault tolerance usually needs at least $3f + 1$ nodes to survive f Byzantine failures. In practice, this is challenging - consensus protocols are known to struggle in real-world deployments. When they do occur, they can be handled by various appropriate algorithms, but they also must be dealt with as they can cause inconsistencies in the decisions made across the different segments that are cut off from one another. All of these assumptions in timing, due to message delays and asynchrony lead to false failure detections and leadership changes, which tend to complicate the whole approach. The ability of nodes to dynamically join and leave the system in a fault-tolerant manner requires mechanisms to reconfigure nodes while maintaining consensus guarantees across membership changes. When systems scale to use thousands of actors, some form of hierarchy or delegation will generally be necessary to mitigate communication overhead. Considerations about implementing consensus further hinder consensus deployment. State transfer mechanisms should ensure that new or recovering nodes become synchronized with the system's current state as efficiently as possible. Since consensus decisions must be durable even across node restarts, we require persistent storage. Log management strategies are guided by the principles



Notes

of minimizing storage and maximizing recovery. Configuration management deals with changes to network topology, security parameters, and protocol tuning values. These practical considerations, often ignored in theory-centric treatments, have significant consequences on real-world system reliability and performance. Specialized needs have driven the development of advanced consensus protocols. Hybrid approaches mix components of various protocols to fine-tune for specific circuit deployment scenarios. These systems often assign different weights to the votes of nodes based on their capabilities or trust level. Leaderless consensus protocols remove the single points of failure by spreading out coordination work between nodes. Probabilistic consensus mechanisms use randomization to work around the FLP impossibility result, offering probabilistic (as opposed to deterministic) guarantees. These adaptations also demonstrate the evolution of the field to new challenges and use cases. Consensus protocols apply and fit on different domains. In distributed databases, they keep data stores replicated in sync. Consensus is the backbone of cloud computing infrastructure for configuration management and service coordination. Consensus For Certificate Issuance And Management Consensus in finance ensures that transactions in a distributed ledger are validated and added. Second, however, these real world applications show the practical significance of theoretical improvements of consensus protocols. This is a very active area of research, with different performance optimizing techniques being developed to improve efficiency while ensuring they have safety guarantees. Batching groups together several proposed values into a single consensus round and amortizes protocol overhead over several operations. Fast paths are optimized for common-case scenarios where conflicts rarely happen. Local decision making allows certain decisions to be made by subsets of nodes where global consensus is not needed. Speculative execution allows operations to be executed provisionally ahead of confirmation, with rollback occurring whenever there is divergence between consensus and speculation. It should be pointed out that these optimizations are but a few of the ways that consensus protocols are being optimized over time to deliver high performance under demanding conditions. This relationship between consensus protocols and the relevant network infrastructure is filled with both challenges and opportunities. Special topologies of the network can improve the efficiency of different challenges of consensus by prioritizing

important messages of the protocol. Quality of service mechanisms facilitate timely delivery of consensus-related communications. For the hardware acceleration, you utilize specialized components like RDMA (Remote Direct Memory Access) or programmable network devices to offload the consensus operations. Together with algorithmic improvements, these incremental optimizations at the level of infrastructure can maximize system performance.

In fact, more and more formal verification techniques are being applied over consensus protocols to provide mathematical assurances of correctness. Model checking exhaustively searches through possible system states to prove safety properties. Formal methods: Theorem proving based system validators using formal logical derivation of protocol correctness. Runtime verification allows for the monitoring of deployed systems for spatial and temporal violations of consensus properties. These approaches can detect subtle bugs and edge cases that might otherwise result in consensus failure in deployed environments. In open or hostile environments, security becomes of the utmost importance. Authentication Mechanisms: These discriminate the entities in the involved nodes. The integrity of protocol messages is protected using cryptographic techniques. It is trust models that outline assumptions about participant behavior and threat scenarios. Incentive mechanisms promote engagement in consensus protocols. These security measures are all the more important in public networks where the participants may not have aligned interests or incentives. With the arrival of blockchain technologies, consensus protocols for open, permissionless environments are again hot topic in computer science research. Bitcoin was the first to introduce the notion of Proof-of-Work which uses computational puzzles to control participation and provide security to the consensus process. Proof-of-Stake alternatives address energy consumption by allowing influence based on economic stake rather than work computation. Normally decision-making is done by delegated consensus models, but this is less performant. Essentially, distributed consensus protocols serve as a key enabling technology for reliable distributed computation, allowing independent nodes to agree on something despite failures or adversarial actions. Therefore, it is not surprising that we see this diversity of protocols, as all of them are essentially built on the same design space of optimizations and trade-offs of safety, liveness, performance and fault tolerance; where



Notes

they differ, is all about the specific deployment and threat models that they were optimized to guard. With the growing reliance on distributed systems, the need for sound consensus will only become more pressing, fueling ever more innovation in this fundamental space.

Unit 12: Consensus in Open Systems and Blockchain

4.5. Consensus in an Open System

This is a strict departure from classical models of consensus, and open systems introduce their own set of challenges and opportunities that have changed the way we think about distributed agreement. Open systems stand in contrast to closed systems, where participation is tightly regulated and all members were known before the system was initiated, and allow new participants to enter the system at will, along with exiting it as they please, in pseudonymous, de-centralized and possibly even adversarial environments. This essential difference requires new methods to achieve a reliable consensus between participants who might have different and even opposed interests. Open System Consensus found its conceptual start from the acknowledgment that the assumptions of traditional closed-system protocols fail to hold. Identities are verified, membership is gated, and the participants usually share aligned incentives in closed systems. For open systems such as currencies, however, these assumptions are way too strong, as they need to operate in the reality that participants can create multiple identities (Sybil attacks), can join and leave at any time and can act in a way that is ultimately against the interest of the system if more profitable for them. This paradigm shift necessitates reconceptualizing consensus not only as a technical protocol but as a socio-technical system inclusive of economic incentives, game theory and mechanism design alongside cryptographic primitives. Perhaps the most fundamental challenge to consensus in open systems is the lack of verifiable identities. In conventional voting-based systems, participants are granted one or a few identities, which secures such systems against attackers who would generate a number of identities to sway a vote (this issue is commonly known as a Sybil attack); however, if participants are able to freely create multiple identities for a minimal cost, the same voting-based consensus becomes vulnerable to manipulation. This fundamental vulnerability required new

mechanisms that define consensus grounds other than “one node, one vote,” culminating in resource-binding approaches that attach power not to an easily replicated digital identity but rather to a well-verified demonstration of investment of rations. The ability to join without permission changes the threat model for consensus far more than this paper acknowledges. In the absence of membership control, the system must therefore consider the possibility that an adversary can all join the network and hold a large fraction of all participating nodes. This reality undermines the conventional Byzantine fault tolerance model where an adversarial minority (less than one-third in the case of general-purpose protocols) of participants is assumed. As a result, open systems have to secure themselves despite adversaries controlling significant fractions of the network, which needs stronger security models and new defensive paradigms.

Open systems also bring additional challenges to the design of a consensus protocol, due to the inherent pseudonymity of its participants. If participants in the system cannot definitely prove their identity, it becomes difficult to trace past behavior to individual contributors, undermining reputation-based approaches to security. Malicious actors can simply drop compromised identities and create new ones at minimal cost, carrying out so-called “whitewashing” attacks that sidestep penalties for bad behavior. Additionally, a continuously pseudonymous environment necessitates mechanisms for reaching consensus that must not be afforded to tampering by dependable nodes. As a result, economic mechanisms have become a key element of the consensus of open systems, offering security assurances and encouraging participation. [6] These systems introduce both financial disincentives for attacks and incentives to behave sincerely, by paying users for adhering to protocols and penalizing them when they deviate from accepted behavior. This merger of economics and protocol design is a stark departure from traditional consensus methods, requiring interdisciplinary knowledge in computer science, economics, and game theory. These systems, which they refer to as “cryptoeconomic,” use economic stakes as security bonds to ensure that the cost of malicious behavior outstrips its potential rewards. In 2009, Bitcoin introduced the first implementation of a solution to these challenges known as Proof-of-Work (PoW), which went on to (indirectly) inspire all major blockchains that followed. PoW (proof of work) prevents Sybil attacks by engaging participants in the computational



Notes

work needed to propose a block, but also provides an objective, decentralized, approach to leader election by means of a resource binding. The cost for such (computational) resources violates the (economic) rule that the barrier of entry for running a full-node is such that anyone can afford to participate but not to disrupt. This mechanism substitutes identity verification with proof that the participant has expended resources, allowing for consensus amongst several impersonal members without direct coordination. However, PoW consensus has some serious weaknesses that make it impractical, especially energy consumption and throughput limitations. PoW systems have seen their energy needs rise to alarming heights; to cite just one example, Bitcoin already consumes as much electricity as certain medium-sized countries do. The environmental footprint has such an impact that it has spurred a developing interest in alternative consensus mechanisms that offer the same PoW security properties but are less energy expensive. Furthermore, the inherent broadcast nature of PoW consensus leads to serious scaling limitations, as every participant must validate each transaction, creating scaling bottlenecks that limit transaction throughput.

Proof-of-Stake (PoS) is the most widely known PoW alternative, tying consensus influence to economic stake, as opposed to computational work. In PoS-based systems, participants stake cryptocurrency assets as collateral, and the probability of being selected to produce the next block is proportional to the amount staked; more stake, more chance to produce the next block. This method greatly lowers energy usage while preserving economic stability, because assailants need to obtain large fractions of the cryptocurrency to attain significant sway. PoS systems operate with mechanisms such as “slashing” that can also seize a participant’s stake for transgressing protocol rules, thus creating direct financial punishment for bad actors. PoS-security: unique attributes and associated security issues clearly differ from that of PoW. PoW security is based on the external expense of obtaining and maintaining mining hardware whereas PoS security is based on the internal price of the pot of cryptocurrency staked. This establishes possible circular dependencies, where an attack that devalues the cryptocurrency also lowers the expenses required to attack it — a problem long known as the “nothing-at-stake” problem. Contemporary PoS systems

respond to these concerns with complex economic incentive design, slashing conditions, and power-to-consensus correlation with economic values at risk. Such a hybrid consensus mechanism has been proposed to take the advantages of different approaches while avoiding their disadvantages. Delegated Proof-of-stake (DPoS) aligns with a representative democracy whereby stakeholders vote on a subset of block producers, thus increasing throughput at the economic cost. Proof-of-Authority (PoA): Leveraging the reputation of well-known validators, PoAs sacrifice some decentralization in semi-open systems for performance. Unlike PoW, Proof-of-Space-and-Time enables participants to furnish storage resources instead of computations, allowing for PoW-level security with lower energy expenses. Such hybrids demonstrate the constant evolution of consensus mechanisms to meet certain needs and limitations. This is where incentive design comes in handy, as part of open system consensus, motivating participation while deterring malicious actors. Block rewards are paid to participants for their investment in consensus resources and are typically comprised of new cryptocurrency created or transaction fees. This inflation funding generates sound economics for provisioning security, enabling value to be extracted from passive holders and directed towards active participants. Funny story: it also means fee markets allow prices for the inclusion of the current public consensus record in the transaction to fluctuate too, joining forms of dynamic pricing that seek to optimize resource utilization in periods of scarcity. Such economic mechanisms will need to balance short-term incentives with long-term system sustainability (also keeping in mind the game-theoretic consequences of reward structures).

Game theory will give us the minimum analytical tools to analyse the security and stability of open consensus systems. Nash equilibrium analysis considers protocol compliance to be individually rational behavior, as long as others also comply. In game theory, we can design protocols that can force rational self-interested players into system behaviors that will achieve the desired outcomes, referred to as mechanism design. Coalition resistance assesses vulnerability to attacks that are coordinated among the groups in the coalition. These analyses should take into account not just incentives at the protocol level, but external aspects such as cryptocurrency markets, mining equipment economics, and the threat of side-channel attacks or



Notes

collusion opportunities. In open systems, governance mechanisms serve as a complement to technical consensus protocols to ensure adaptation and evolution. On-chain governance uses the consensus mechanism itself to formalize a process for voting on and approving changes to the protocol. A process termed informal governance is driven by social consensus surrounding protocol updates among community members. Divergent views to fork-based governance can generate separate systems for the market to compete as to which ideology wins. Governance models in this space must strike a balance between responsiveness to evolving needs and stability and predictability, particularly where significant economic value is at stake. A well-known compromise comes from the nature of open systems, where new protocols finally encounter resistance of changing consensus history by new members. In contrast to closed systems, in which all actors can upgrade simultaneously, open systems must contend with the situation where actors are running different software versions, different governance decisions are being made, and so on. In fact, there has become quite a complex of fork management strategies: In Outcome A, there are hard forks with incompatible branches and soft forks that maintain backward compatibility; velvet forks, where new features are introduced but can wait to be adopted. These mechanisms facilitate protocol evolution in the lack of central coordination authority. Open consensus systems space is still an ongoing research addressing the core throughput limitations. Through sharding techniques, the network can be broken into multiple parallel consensus groups that can process blocks simultaneously, thus enabling the throughput to be multiplied. Layer-2 use a scaling framework for moving transactions outside of the main consensus layer, having the primary layer act exclusively as a settlement and dispute resolution layer. State channels allow interaction between participants with negligible consensus overhead. These scaling approaches weigh the throughput increase against possible security trade-offs or additional complexity, aiming to find configurations that are optimal for a given application. This network dynamics in open systems keeps introducing more complexity into the consensus mechanism. Eclipse attacks can separate a node from the network, effectively manipulating its perceived consensus. Network partitions temporarily separate the global system, necessitating reconciliation mechanisms upon reconnection. Network propagation latencies introduce short-lived deviations of

distributed state. This requires the design of protocols that incorporates timeout mechanisms, fork choice rules and network-aware optimizations to maintain the same global state despite unpredictable connectivity. Privacy remains an ever-increasing concern in open consensus systems. Although public blockchains remain open by nature, the transparency poses great privacy problems for many applications, as traditional public blockchains broadcast all transaction details. With zero-knowledge proofs, we can confirm computation without disclosing our inputs, empowering private transactions on public rails. Notable implementations include smart contracts and confidential transactions, where the latter conceals transaction values and allows for balance conservation verification. Commingled protocols obfuscate transaction graphs, making it difficult to analyze flows of payments. These privacy-enhancing technologies allow open consensus systems to be deployed in sensitive use cases that require confidentiality. One of the conditions for ecosystems development has become interoperability between specific open consensus systems. Two types of cross-chain communication protocols are: Atomic swaps enable trust-minimized transfer of assets between disparate consensus systems. Federated pegs allow transferring assets between systems while optimal security guarantees are held. This latter set of interoperability solutions allows separating out the purpose of different consensus systems while maintaining the ability to compose their capabilities, ultimately resulting in a web of purpose-specific consensus systems. Open consensus security issues go beyond the consensus protocol to the sociotechnical ecosystem. Even if the underlying consensus is secure, smart contract vulnerabilities can leak assets controlled by consensus systems. Governance attacks hijack decision processes for a protocol so that they ratify malicious changes. Social engineering attacks human components of the system through deception or manipulation. Such multi-dimensional security threats require defensive-in-depth strategies, integrating technological safeguards with social and organizational measures.

The viability of an open consensus system is largely determined by appropriate balance between security costs and value provided. There must exist a soundness of work between cost and attack to ensure profitable attacks do not exist. When block rewards disappear, transaction fee markets need to provide enough income to sustain security. Bootstrapping new



Notes

systems is especially tricky because, at first, security is more expensive than the small value the system offers, making it difficult to put in place. These economic realities have led to growing attention given to sustainable design of the underlying tokenomics that is able to provide long-term security funding without needing to rely on perpetual inflation or external subsidy. When open consensus systems become economically significant, the need to address regulatory considerations becomes acute as well. Anti-money laundering obligations risk clashing with the pseudonymous property of many open systems. Some PoS mechanisms might lead to these consensus tokens being classified under certain securities regulations. Some jurisdictions regulate to whom the consensus participants and service providers can provide their services to and therefore certain licensing requirements apply. These are always-evolving regulatory frameworks as agencies have been coming up with methodologies to facilitate innovation while also addressing the need of consumer protection and law enforcement. Open system consensus has been evolving with continued progress on the theoretical and practical fronts. An accountability mechanism helps in detecting protocol violators while preserving pseudonymity for compliant parties. They help in mitigating certain timing attacks and facilitate more efficient variants of consensus. Protocols can also be supplemented with formal verification techniques that guarantee the correctness of their properties in a mathematical sense. These continuing developments demonstrate the healthy research and development community that is tackling theoretical and practical problems around consensus in open systems. Thus, consensus in open systems is a groundbreaking method for achieving distributed agreement, facilitating coordination among pseudonymous agents without relying on a central intermediary. These systems achieve a trusted consensus through the unique combination of cryptographic primitives, economic incentives, and game theory that is absent in traditional trust structures. There are still many challenges that impact the scalability, privacy, governance, and sustainability of current systems, yet the space is evolving so quickly, it is continuously producing new solutions and pushing the boundaries of open consensus systems. As these technologies mature, they also facilitate new forms of human coordination and value exchange that were previously only possible through trusted intermediaries.

4.6. Consensus in Bitcoin Network

The consensus mechanism of the Bitcoin network is a radical new model for distributed agreement, one that solved the intractable issue of achieving consensus among pseudonymous players without a central coordinator. First introduced in Satoshi Nakamoto's 2008 whitepaper, this innovation effectively solved the Byzantine Generals Problem in an open/open permissioned environment—a leap that computer scientists had largely believed to be practically impossible. By combining cryptographic primitives with economic incentives, Bitcoin produced a resilient consensus system that has run nonstop since January 2009, protecting hundreds of billions of dollars in value without centralized control or traditional institutional support. At a conceptual level, the unique design of Bitcoin's consensus mechanism marked a departure from known methods by substituting authority based on identity for authority derived from proof of resource expenditure. Instead of establishing consensus authority based on verified identities or institutional credentials, Bitcoin uses PoW to establish a consensus authority in proportion to the amount of computable effort that can be demonstrably put in. Such a model directly mitigates against Sybil attacks whereby a single adversary can gain influence far greater than their actual stake by creating other pseudonymous identities; rather than trying to compete with an unlimited number of faux peers, each peer must pay for their stake in consensus. In this paradigm, nodes known as miners engage in a competitive race to solve computationally intense cryptographic challenges, with successful solutions granting a temporary license to propose new transaction blocks. The specific cryptographic puzzle used in Bitcoin's PoW construction is to locate a partial hash collision, more precisely, to generate a block header whose SHA-256 double-hash begins with a predetermined number of zero bits. This is achieved by miners adjusting a nonce value in the header of the block and hashing the block repeatedly until the hash output is less than the target threshold defined by the network. This mechanism is carefully crafted to be compute-intensive yet simple to verify, demanding extensive computing capabilities to generate bona fide solutions, yet allowing any participant of the network to promptly verify by checking some equations. It is very hard to find a piece of data that validates a certain condition, while it is very easy to verify that this condition is met, enabling an inexpensive decentralised validation



Notes

system without requiring trust between members. This consensus is not limited to validation of block, but also extends to the veracity of transaction, monetary policy, protocol mechanics, etc. With the exception of the coinbase input, any transaction that works against a full node is considered valid; thus, every single node on the network is individually, by definition, validating every transaction against an extensive set of constraints including signature verification, input existence, double-spend prevention, and supply constraints. This additional and unnecessary act of verification achieves a consensus about the state of the system throughout the entire network, without the need for trusts between nodes. Block rewards, transaction fees, difficulty and limit sizes are defined with limits by the protocol, and together they make up the constitutional parameters of Bitcoin. By ensuring the majority of nodes stick to these rules, and rejecting any rule which deviates from these rules, protocol rules are enforced by the participation of the nodes in the network, rather than from any central authority. The difficulty adjustment is an essential adaptive mechanism of Bitcoin's consensus system that keeps the network producing blocks at roughly the same frequency despite changing computational work done by the network. The network then automatically adjusts the difficulty target every 2016 blocks (aiming for every two weeks) based on how long it took to generate the last 2016 blocks. If blocks are coming faster than the 10-minute target, difficulty increases proportionally; if slower, decreases. This creates a feedback loop that regulates block production, so even under significant variations in miner input or hardware productivity, predictable block production remains - using transaction throughput and prediction of monetary issuance as the fixed variables and using external variables such as hardware efficiency and miner participation as the externalities. The longest chain rule (more accurately: the greatest cumulative proof-of-work chain) constitutes Bitcoin's simplest fork choice rule—it is the algorithm for choosing between competing versions of transaction history. If there are several valid blockchain versions (e.g., usually due to temporary network sections or simultaneous block discoveries), nodes pick the chain that represents the most amount of accumulated computational work. This simple, yet powerful rule creates eventual consistency without the need of explicit coordination since rational miners will always work to maximize their returns by working atop the chain that has the greatest chance of being the global consensus. The need

for majority honest computation power is Bitcoin's fundamental security assumption and becomes prohibitively costly once the network is large enough and sufficiently distributed.

Bitcoin's brilliant incentive mechanism has aligned miners' economic self-interest with invention of network security and integrity. Block rewards – comprising freshly minted bitcoins and transaction fees – incentivize miners for their computational investment and, at the same time, distribute the currency and establish network security. This issuance schedule halves the block reward roughly every four years, providing a predetermined monetary policy and transitioning security funding from inflation to transaction fees, slowly but surely. The way this economic engine works is by incentivizing honest activity on the protocol and ensuring that attacks are, economically, non-viable; the cost for attack will generally always exceed the profits attainable through it, especially considering the expected market reaction on an asset valuation of successfully conducted attacks. The mempool — the pool of unconfirmed, valid transactions for each node to maintain and propagate — is a distributed waiting room for Bitcoin and thus a market for block space. Miners, when compiling blocks, tend to favour transactions that include higher per-byte fees, resulting in a fee market that efficiently allocates the limited resource of block space to users based on their urgency. When demand skyrockets, users will essentially be in a bidding war, with transaction fees crying out for what the user is willing to pay to get included in the next block. By making transaction prioritization market-based, users can self-select fees to pay based on their personal pricing of the urgency to process a transaction along with the associated transaction amount, without needing to coordinate a central arbitrator to improve resource efficiency. Network propagation dynamics play an integral role in Bitcoin's consensus behavior and security characteristics. Miners propagate valid blocks immediately to the network so they can collect their reward before other miners find alternative blocks. This propagation race creates incentives for efficient network connectivity, since propagation delays increase “orphan risk”: the chance that a competing block from another miner propagates widely enough in the network to reach more nodes than our competing block. These dynamics had led to mining centralization in history to reduce propagation delays, but protocol improvements such as compact blocks have mitigated this centralizing pressure considerably. The propagation



Notes

characteristics also nicely constrains practical block sizes, since larger blocks take longer to propagate in the network, increasing orphan risk for their author. Like in the other PoW protocols, Bitcoin does not guarantee the absolute finality of transaction settlements, but instead offers probabilistic finality. The more expensive (in terms of proof of work) blocks that have been mined on top of the block that your transaction is included in, the less likely it is to be reversed. So this is the probabilistic model, which substitutes traditional atomic settlement for practical finality that reflects value; i.e., users usually wait for more confirmations when processing higher value transactions. But for most practical purposes, six confirmations (about an hour) is enough, and in extreme value transfers more confirmation depth may not be worth the trade off of time. This way security is nicely traded off against practical usability without requiring absolute guarantees. The UTXO (Unspent Transaction Output) Model Driving Bitcoin's Consensus Verification & Scalability. Whereas account-based systems store balances, Bitcoin keeps track of ownership using discrete “unspent outputs” created in previous transactions. Transactions spend existing UTXOs as inputs and create new UTXOs as outputs, and validation ensures that inputs exist, belong to the spender, and are at least equal to output values. This structure allows transaction validation to easily be parallelized since whether a transaction is valid or not purely depends on the existence and spendability of its inputs, but not the global state. The UTXO model also allows for an increase in privacy because it promotes the use of new addresses in every transaction, making ownership analysis more difficult. The Bitcoin transaction system is a meta programming language, where you can write scripts that dictate the conditions under which certain funds can be spent, thus allowing reasoning not just about transfers, but about the states of transactions to which the consensus applies, including conditional logic and time constraints. Bitcoin Script, designed to be simple and non-Turing-complete, provides basic smart contract functions like multi-signature locks, time locks, hash-locked contracts, and payment channels. This scripting system's minimalism is a conscious choice made in favor of security and predictability rather than expressiveness, echoing Bitcoin's design philosophy that favors complex security and decentralization at the expense of complex feature richness. This limited surface area while still providing necessary programmability for second-

layer protocols and interoperability mechanisms Over the years, mining pools have emerged in Bitcoin as a key organizational layer in this practical consensus mechanism, enabling individual miners to pool up their computational resources and receive more predictable rewards. By splitting block discovery rewards in proportion to contributed computing power, pools reduce variance for individual participants, and in effect turn winner-take-all block discovery competition into lower variance steady streams of income. While this has some benefits, such as lowering volatility and reducing the power required for network given the reward, pooled mining also comes with some potential centralization concerns, since a small number of pool operators control large amounts of network hash power. Thus, having competitive pressure from individual miners rapidly reconfiguring compute resources between pools mitigates the centralization risk and limits pool misbehavior.

The notion of obtaining initial synchronization is vital to Bitcoin's consensus system but a topic frequently overlooked. Efficiently determining the state of the valid blockchain from non-credible peers poses a problem in any practical scenario! This compact block validation scenario represents the process of Initial Block Download (IBD), whereby compact blocks (and relevant contents) of the blockchain history is verified efficiently to reach consensus state by each new participant of the blockchain, without the need to trust a priori participants of the blockchain. Permissionless participation is a prerequisite for censorship resistance and long-term sustainability, and is enabled by this very independent verification property of Bitcoin, as participants do not require trusted third parties to act on the rules. Consensus alterations in Bitcoin are profoundly conservative affairs that can be gleaned by the high stakes of a single system securing hundreds of billions of dollars in economic value. These backward-compatible changes are called soft forks, which tighten the consensus rules and can adopt gradually without forcing an immediate upgrade of the entire network. Hard forks bring about consensus changes that aren't compatible with older rules and, unless everyone adopts them, would result in two networks. Miners, full node operators, exchanges, and users must achieve consensus before changes to the protocol are made, which favors stability and security over rapid code evolution.



Notes

This conservative engineering strategy has also worked well for Bitcoin itself, maintaining its core properties even as its value secured and ecosystem complexity has exploded. The Bitcoin security model is based on a number of fundamental assumptions that guide its threat model and its limitations. The honest majority assumption states that most of the participants controlling more than 50% of the computational power follow all the protocol rules and append to the longest valid chain. The assumption about network connectivity implies solidity in message propagation between honest nodes to avoid long periods of partitioning. Even in widely characterized situations designed for free-ride problem solving, where participants can accept losses to undermine the process, economic rationality assumptions presuppose participants will instead respond to maximize financial self-interest. These assumptions have held true throughout Bitcoin's operational history, but do require continued monitoring over mining decentralization, network diversity, and kinds of attacks motivated by considerations other than economic incentives. Over the years, many possible attack vectors on Bitcoin's consensus have been theoretically identified and explored. The 51% attack, in which an entity holds a majority of computational power, theoretically allows double-spending or transaction censorship, but due to the amount of capital investment required and the expected market response, such attacks are still economically irrational for profit-motivated actors. Selfish mining, however, has been shown to work theoretically but is challenging to implement in practice profitably. Eclipse attacks would disconnect specific nodes from honest peers, allowing targeted attacks against specific users. Theoretical vulnerabilities of these kinds have directed ongoing protocol work and operational security practices throughout the ecosystem.

With the advent of specialized hardware for mining, gone are the days when you could mine Bitcoin with a CPU or a GPU—bitcoins are now mined with ASICs. Over the years, mining has been made several orders of magnitude more efficient through the use of Application-Specific Integrated Circuits (ASICs) designed specifically for the SHA-256 hashing algorithm, which has elevated the baseline computational cost of mounting an attack against the network. The hardware specialization has made mining much more professional, moving from hobbyist engagement to industrial machinists

aiming to maximize power-to-cooling operations. The capital-intensive aspects of competitive mining have definitely led to some centralization pressures in hardware manufacture and access to low-cost electricity; however, those pressures have been (at best) limited in scope due to global competition and a strategic view on the dangers of excess and singular concentration. Bitcoin's energy usage, a direct consequence of its Proof-of-Work consensus mechanism, is arguably its most controversial feature. The massive power consumption — more than small countries — has triggered considerable criticism over the environmental impact. Supporters argue that Bitcoin increasingly uses stranded or otherwise wasted energy sources, encourages the development of renewable energy through demand that is not bound by geography, and delivers value that validates its energy footprint. This persistent debate encapsulates the clear trade-off in PoW systems between security from resource usage, and the environmental impact of that resource usage. Ultimately, the economic efficiency of this security approach is governed by the amount of monetary value secured per unit of energy consumed. To alleviate Bitcoin's reconfirmation stress, several layer-two scaling solutions have been developed to increase the network's output capacity with reasonable guarantees of base layer consensus security. Bitcoin's main second-layer solution, the Lightning Network, utilizes smart contracts to form payment channels among two or more participants, allowing nearly instantaneous transfers while using minimal on-chain space. All these channels connect into a network that allows payments to flow between party A and C, even if they don't have a direct channel between them. Lightning greatly increases effective transaction capacity by moving most transactions off-chain, while only leveraging Bitcoin's blockchain for final settlement of state transitions or dispute resolution, inheriting the security guarantees offered by the underlying consensus mechanism. This layered approach preserves Bitcoin's underlying decentralization and security properties, while accommodating real-world scaling needs.

Bitcoin's most significant improvement to its core consensus mechanism is the Segregated Witness (SegWit) upgrade which was activated in August of 2017 and separated the transaction signature data from the transaction identifiers. It solved transaction malleability that had made second-layer protocols difficult to implement, and it also doubled effective block size via



Notes

a signature size discount. This was implemented as a soft fork, so it was backward-compatible and could be adopted gradually without requiring a network upgrade. The process of deploying SegWit built a lot of those governance principles straight into the consensus process — you needed proof of support from miners and node operators alike, and only staged activation was allowed, creating a blueprint for how future consensus upgrades would require an alignment of interests between technical innovation and stability concerns. Bitcoin — its security economics is directly upon its fee market dynamics. These fees can provide significant revenue in addition to the fixed block subsidy that miners receive, especially during periods of high transaction activity when limited block space results in increased fees. As inflation-based funding shrinks block rewards toward zero, this fee market becomes the only long-term sustainable model of funding Bitcoin security. Such a change opens up first principles conversations about the ideal supply of block space, the long-term level of fees and the trade offs seen between accessibility vs. security funding. They implement an emergent market-based approach to security funding, which is both a source of strength (due to self-adjusting properties) and uncertainty (as we want to be in these equilibria in the long-term).

The consensus design of the Bitcoin protocol explicitly favors decentralization and censorship resistance over throughput and transaction efficiency. This purposeful concession is evidence of Bitcoin's primary value prop as sovereign, censorship-resistant currency rather than some kind of universal payments system. Incremental increases in block size can be conservative enough to maintain wider full node accessibility, thereby facilitating wider verification opportunities, and increased trust minimization. This design philosophy accepts some trade-offs at the base layer while the security and immutability properties that matter for Bitcoin's core usage remain, and pushes applications focused on efficiency to second layers or other solutions. This design decision has informed the trajectory of Bitcoin's development and seen it settle primarily into the role of "digital gold" rather than a payment facilitator. BTC's security model bootstrap process was incredible mechanism design--getting functional consensus with little initial value. At early stages, when native tokens had negligible market value, the security budget effectively depended on the speculative value that

early adopters assigned to future utility. As the market cap increased, the security from the mining rewards increased proportionally, which brought an additional increase in security and all over confidence in the system. We covered how Bitcoin's incentive structure led to successful bootstrapping from little security to much security protecting a large value, the implications that this has for many cryptocurrency launches that followed trying to recreate this bootstrapping in different consensus models. Each node in the Bitcoin network utilizes mesh topology with peer-discovery features to protect the network from targeted attacks and potential disruptions. Each node can be connected to many peers and transmit transactions and blocks across the network without central coordination. In addition, it blends DNS seeds, hardcoded entry points, and previously known addresses to make connections despite low prior knowledge. Such a distributed communication structure guarantees the operation of the network even in case of failure of some nodes or attempts of breaking the connection between them. This peer-to-peer infrastructure has been essential to Bitcoin's censorship resistance, and has withstood equal and opposite assaults throughout its operational history. In summary, the Bitcoin consensus mechanism is a game-changing distributed agreement paradigm that has profoundly changed the way we think about transferring value digitally. Through the use of cryptographic methods and economic incentives, Bitcoin built a working consensus mechanism that has been operating safely for more than a decade, without any centralized authority. Although the methodology involves deep trade-offs in terms of throughput, energy usage, and the agility of governance, these constraints are precisely what allow for Bitcoin's fundamental value propositions of censorship resistance, predictable monetary policy, and trustless proof-of-accuracy. This consensus mechanism balances trade-offs that have proven to be appropriate design choices for Bitcoin as sovereign digital money, a design choice that has withstood the test of time in terms of sustained operation and increasing value, even as new alternative consensus mechanisms are developed and optimized towards other use cases and priorities.

Multiple Choice Questions (MCQs)

1. **Why is consensus important in a blockchain network?**
 - a) To allow a central authority to validate transactions



Notes

- b) To ensure all nodes agree on the validity of transactions
 - c) To increase the processing speed of a blockchain
 - d) To make transactions private and anonymous
2. **Which of the following is NOT a property of distributed consensus?**
- a) Agreement
 - b) Termination
 - c) Centralization
 - d) Integrity
3. **In a synchronous system, nodes operate with:**
- a) No time constraints
 - b) A fixed time limit for message transmission and processing
 - c) A random delay for message passing
 - d) No coordination at all
4. **What is the key difference between synchronous and asynchronous consensus systems?**
- a) Synchronous systems work without any time limits, while asynchronous systems have fixed deadlines
 - b) Synchronous systems assume timely message delivery, while asynchronous systems do not
 - c) Asynchronous systems are more secure than synchronous systems
 - d) Synchronous systems do not require consensus mechanisms
5. **Which of the following is a widely used consensus protocol in blockchain?**
- a) Proof of Time (PoT)
 - b) Proof of Space (PoS)
 - c) Proof of Work (PoW)
 - d) Proof of Banking (PoB)
6. **Which consensus mechanism is used in the Bitcoin network?**
- a) Proof of Stake (PoS)
 - b) Delegated Proof of Stake (DPoS)
 - c) Proof of Work (PoW)
 - d) Byzantine Fault Tolerance (BFT)

7. **What is the main challenge in achieving distributed consensus?**
 - a) Ensuring every participant gets rewarded equally
 - b) Preventing double-spending and Sybil attacks
 - c) Allowing only government-approved transactions
 - d) Keeping blockchain networks centralized
8. **Which component of Bitcoin mining ensures consensus among participants?**
 - a) Cryptographic hashing
 - b) Centralized decision-making
 - c) Government regulation
 - d) Random transaction approvals
9. **In an open blockchain system, consensus is achieved by:**
 - a) A single trusted authority
 - b) A network of decentralized participants
 - c) A central server validating all transactions
 - d) A randomly chosen validator
10. **What is the primary role of a consensus algorithm in a blockchain network?**
 - a) To keep a centralized record of transactions
 - b) To verify and agree on the state of the blockchain ledger
 - c) To prevent nodes from communicating with each other
 - d) To increase transaction fees for miners

Short Answer Questions

1. What is consensus in a blockchain system?
2. Explain the role of distributed consensus in decentralized networks.
3. What are the main properties of distributed consensus?
4. Differentiate between synchronous and asynchronous systems in blockchain.
5. How does the Bitcoin network use Proof of Work (PoW) for consensus?
6. What are the key challenges in achieving distributed consensus?
7. Explain why an open system needs a consensus mechanism.



Notes

8. What happens if two nodes disagree on the blockchain's current state?
9. How does consensus prevent double-spending in Bitcoin?
10. What are some alternative consensus mechanisms besides PoW?

Long Answer Questions

1. Explain the need for consensus in blockchain networks. How does it ensure security and trust?
2. Discuss the differences between synchronous and asynchronous systems in achieving distributed consensus.
3. What are the properties of a good consensus protocol? Provide examples.
4. Describe Proof of Work (PoW) and its role in securing the Bitcoin network.
5. What are the advantages and disadvantages of PoW as a consensus mechanism?
6. Compare different consensus mechanisms such as PoW, PoS, and Delegated PoS.
7. How does consensus work in an open blockchain system, and why is it challenging?
8. Explain how the Bitcoin network resolves conflicts when multiple blocks are created simultaneously.
9. Describe the role of miners in achieving consensus in a blockchain network.
10. What are some emerging improvements in consensus mechanisms that could replace PoW in the future?

MODULE 5

PERMISSIONED BLOCKCHAIN

LEARNING OUTCOMES

By the end of this Module, students will be able to:

1. Understand the Permissioned Blockchain Model and its advantages over permissionless blockchains.
2. Explore various use cases of permissioned blockchains in industries like finance, supply chain, and healthcare.
3. Explain the concept of smart contracts and their role in automating transactions.
4. Identify the design limitations of permissioned blockchains.

Unit 13: Permissioned Blockchain and Its Use Cases

5.1. Permission Blockchain Model

Since the inception of Bitcoin in 2009, blockchain technology has developed vastly, leading to the formation of different models for various needs spanning different industries. This has led to the rise of permission blockchain models as a trusted approach that combines the decentralized attributes of blockchain technology with the needs of organizations seeking privacy, compliance, and controlled access. And unlike their permissionless counterparts, which allow any individual to remain anonymous to interact with the network, permission blockchains enforce key access controls to limit who can participate—and in what capacity. The defining characteristic of permission blockchains is that their members are only granted access to the blockchain if they have been pre-approved by the governing body (or consortium) that controls that permissioned ledger. This creates a more controlled environment in which the identities of the participants are known and validated, resulting in accountability in the network. Most importantly, you will have to undergo Know Your Customer (KYC) and Anti-Money Laundering (AML) checks, which help exchange platforms to adhere to regulation requirements, that vary by region. This use case makes permission blockchains especially appealing for enterprises and institutions in highly regulated industries.

PERMISSIONED BLOCKCHAIN

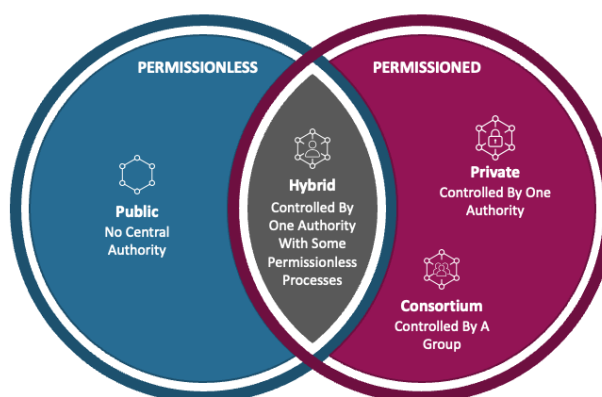


Figure 8: Permissioned and Permissionless Blockchain

The permission blockchains architecture runs with multiple access levels. A blockchain like that ensures that only authorized nodes can join the blockchain infrastructure at the network level, making it secure from the ground level. In the consensus mechanism layer, they introduce specialized protocols, which are usually different from resource-intensive public blockchains' proof-of-work systems. Rather, permission blockchains often use consensus algorithms such as Practical Byzantine Fault Tolerance (PBFT), Proof of Authority (PoA), or variations of Proof of Stake (PoS) designed around networks with known stakeholders. This results in not only better use of computational resources and energy, but also faster finality of the transactions which is key for business use cases. Within permission blockchains, sophisticated hashes and encryption techniques allow fine-tuned data visibility according to requirements. Although every participant keeps a copy of the ledger, cryptographic controls can limit access to sensitive knowledge, producing a model where information is shared on a need-to-know basis. This level of selective transparency is a stark contrast to public blockchains, where data is viewable by all — and tackles the well known privacy issues that have historically prohibited blockchain from being adopted in certain sectors. The most rigorous form of governance in deployed permission blockchain networks is consortium governance, where stakeholders formalise governance by collectively deciding on things such as protocol upgrades, policy changes, and participant management. By working together, the community can steer the blockchain's growth in directions that serve the best interests of its members, all while ensuring that it remains stable and functional. The governance framework sometimes includes dispute resolution mechanisms, removal of participants when needed, and protocol changing protocols—for a degree of adaptability that is tough to achieve in permission-less systems.

Permission blockchains also have enhanced performance attributes over public blockchains owing to their fine-tuned design. By preprocessing and validating transaction nodes before inclusion, they provide better throughput and latency. The performance advantages make permission blockchains suited to use cases with more demanding operational requirements, including financial settlements, supply chain tracking and



Notes

healthcare data management. The Security in the permission blockchains comes from a combination of cryptographic techniques and accountability present in known participant networks. Public blockchains rely primarily on economic incentives to deter malicious behavior, while permission blockchains are more shielded by the reputational and legal consequences that would accrue if an identified participant misbehaved. This social layer of security enhances the technical guards, leading to a tough defense against many attack vectors. Another main benefit is the integration capabilities of permission blockchains, which are mostly built for enterprise integration. Moreover, they provide APIs and integration frameworks that allow for tying in with existing enterprise systems so that organizations can take advantage of blockchain technology without having to completely replace their entire IT organization. Interoperability is not limited to other permissioned blockchains; communication protocols such as the cross-chain Interoperability Protocol (ICP), enable permission blockchains to communicate with public and private networks when required. Leadership: The structure of permission blockchains is totally different than public crypto currencies. These networks typically employ utility tokens or even entirely tokenless architectures that fulfill specific functional roles within the ecosystem instead of attempting to capture speculative token value. This method is compatible with the economic driven characteristic of these networks where the fundamental value proposal is not that of financial speculation but operational efficiency. Performance enhancements for private permissioned blockchains rely on the controlled conditions, which enable optimizations that may be difficult to implement in public networks. When the set of participants is known and trusted, techniques like sharding, sidechains, and layer-2 solutions become more efficient, allowing these networks to scale to handle a larger number of transactions without the same performance tradeoffs. Streamlining your Industry Needs

The Different permission blockchain platforms A modular architecture is provided by Hyperledger Fabric, started by the Linux Foundation that allows many components (e.g., consensus mechanisms and membership services) to be plug-and-play, giving outstanding flexibility in the context of enterprise implementations. Startup R3's Corda was designed specifically for financial services, where transactions occur on a need-to-know basis and sensitive data need only be shared with those with a valid reason to see it. Developed

by JPMorgan Chase and now managed by ConsenSys, Quorum is an enterprise-focused fork of Ethereum with added privacy and performance-enhanced features. Enterprise Ethereum is effectively based on the public Ethereum blockchain, albeit with permission aspects and privacy improvements to satisfy corporate demands.

Essentially, permission blockchains have more of a phased rollout, starting with an analysis of organizational requirements and use cases. Next comes evaluating each blockchain platform to decide on one that best satisfies these needs based on the consensus mechanism, privacy, and strength of the ecosystem. Designing the network means designing the right node topology, including optimizing validating nodes, observer nodes and whatever other roles are necessary based on the specific use case. Participants should be onboarded properly, with clear processes for identity verification and credential management. The most critical aspect of any solution, security, is also adequately covered, with extensive threat modeling, high-level protection against external malicious attacks, and mitigation against possible insider threats. There are a few challenges/limitations to permission blockchains that need to be accepted. It marks a shift from the fully decentralized utopian ideal of blockchain — and from that perspective, it could create points of failure or a vulnerability. Requiring some form of identity verification adds administrative costs, which in some circumstances, raises concerns about privacy (especially in times of war). However, successful interoperability with other blockchain networks remains a complex technical challenge that has made headway with different cross-chain protocols. While it is true that regulatory uncertainty continues to pose challenges to any emerging technology, many jurisdictions are still evolving frameworks that address the use and deployment of AI broadly. Many developments can be foreseen, going forward for permission block chains. Hybrid models that integrate aspects of both permissioned and permissionless systems are also emerging, providing a more balanced approach that harnesses the advantages of both paradigms. To better protect user privacy while preserving the verifiability that makes blockchain useful, more diverse solutions such as zero-knowledge proofs and secure multi-party computation are being employed. Organizations working to create common protocols that help drive toward increased interoperability and growth are seeing increasing momentum as



Notes

standardization efforts are emerging throughout the industry. The regulatory environment is evolving and there are increasingly formalised requirements, with a growing number of jurisdictions putting in place more sophisticated frameworks that take into account the specific characteristics of blockchain technology. The permission blockchains are a onground modification of the blockchain network to be compatible with the needs of companies and institutions working in the platforms that have complex regulatory structures. Through a combination of the well-known features of immutability and transparency provided by blockchain, together with the relevant controls required for compliance and governance, these systems are delivering practical applications across a number of sectors. As this technology matures and evolves further, permission blockchains occupy a key part in the blockchain estate, especially for situations in which trust, compliance, and effectiveness ought to all focus together.

5.2. Use Cases

These systems are based on blockchain technology where certain participants are granted permission to use the system, allowing for control over who can access and participate in the network. Because participation is the participants with known provenance and verified, these networks become specifically ideal for use situations in which compliance, privacy and high performance are the critical components. Until recently, permission blockchain used cases have multiplied and shown value in real-world implementations as organizations have started to understand the power of blockchain far beyond cryptocurrencies. Common use cases Permission blockchains are making a big impact in the financial services industry, where they are streamlining existing workflows to catch up with the speed of modern banking while still keeping regulatory forces in place. With blockchain networks, interbank settlements, once fraught with delays and reconciliation challenges, can be streamlined, enabling participating banks to collectively validate transactions and shortening the settlement period from days to minutes. Trade finance, which relies on a maze of paperwork and several players, is a perfect use case for blockchain-based platforms which digitize letters of credit, bills of lading and other important documents and create a single source of truth available to everyone with the right permissions. This change has substantially lessened paperwork, decreased

fraud risks, and sped up the funding process for international commerce. The idea of shared blockchain registries has transformed Know Your Customer (KYC) processes — Internet of Value → regulatory compliance, a KYC is necessary to comply but costly and often redundant with many players. These systems enable financial institutions to share with each other verified information about their customers in a secure manner, avoiding duplicate verification and still ensuring strong privacy controls. Blockchain platforms are bringing enhanced efficiency for syndicated lending, which is a loan type comprised of multiple lenders and involves a high-value loan in which multiple lenders work together to provide funds in a collaborative manner, helping to automate elements such as covenant tracking, payment distributions and reporting — cutting operational costs while increasing transparency for all parties involved. Establishing permission blockchains has become a popular subject for insurance companies to delve into to solve some of the long-standing issues plaguing the industry. Shared ledgers that track claim histories, policy details, and supporting documentation have made claims processing smoother and more efficient, while immutable record-keeping makes fraud more difficult to carry out. Blockchain-based smart contracts have improved parametric insurance products that automatically trigger payouts based on predefined criteria by connecting these contracts to trusted data sources and automatically executing payments, leading to a more efficient and consumer-friendly insurance process. Supply chain management is one of the most attractive application areas for permission blockchains, because those complex global networks involve thousands of stakeholders needing to coordinate activities while also keeping their operational autonomy. This system records every transfer of custody in a tamper-proof manner, allowing for an immutable record of provenance across a products journey from factory door to retail shelf. This allows for timely decision making that is critical for industries such as pharmaceuticals, where counterfeiting can risk human lives, or luxury goods, where verification of authenticity directly relates to economic value.

Blockchain solutions enabling transparency over real-time data around inventory levels and movement in distributed supply networks have brought transparency to inventory management processes, reducing both stockouts and overstocking, and significantly cutting down on the administration needed to reconcile inventories between organizations. Supplier credential



Notes

management has been made easier by blockchain-based platforms that store verified supplier certifications, compliance documents, and performance histories, streamlining the supplier onboarding process and guaranteeing that all participants meet required standards. They have realised significant value in utilizing the permission blockchain applications that address the unique requirements of the healthcare industry in terms of data security, patient privacy, and regulatory compliance. Advanced EHR systems built using blockchain infrastructure permit secure sharing of patient data among healthcare providers whilst incorporating rigorous access policies that are sensitive to patient consent preferences. The fact that all data on the blockchain is immutable means that any access to those medical records will be permanently recorded, establishing responsibility and also empowering patients with more control over their health information. Streamlining clinical trials using blockchain platforms that register study protocols, patient consent and research outcomes in a secure but transparent manner helps improve data integrity at all stages of the research process and simplifies regulatory review. Industry leaders have adopted blockchain tracking systems in pharmaceutical supply chains to monitor temperature conditions, handling procedures, and chain of custody of sensitive medications, which helps secure required product efficacy and patient safety while streamlining compliance documentation processes.

The complete energy sector has been adopting permission blockchain technology towards transitioning to more distributed and sustainable energy systems. First, a major development in cutting-edge business models: peer-to-peer energy trading platforms allow consumers that generate energy from renewables to sell excess energy to neighbors, creating local energy marketplaces with minimal intermediation. Production, consumption, and financial settlements are automatically recorded through secure blockchain ledgers, making this more economically viable for communities as a whole in integrating renewable energy. Blockchain-based renewable energy certificate (REC) tracking systems ensure irrefutable verification of green energy generation and usage, thus resolving double-counting issues and offering integrity to corporate sustainability statements. These applications use the decentralized power of blockchain to coordinate between numerous energy sources, such as both traditional power plants and renewable/ battery

storage systems, mission-critical in helping build resilient and efficient energy networks. The real estate industry has witnessed a major revolution through blockchain-based solutions that allow digital records of property titles, ownership histories, and other relevant documents. This reduces the time-consuming task of performing title checks, and minimizes the possibility of ownership disputes, as there is only one official record of property rights, which can be accessed by all the relevant parties (buyers, sellers, banks and government). Automating agreement processes, maintenance history, payment processing — blockchain applications have made property management operations more efficient while offering more transparency to property owners and tenants. The public sector is identifying greater use cases to leverage enhanced government services through permission blockchain technology with the necessary control. Onureacts as an integrated portfolio for a citizen's personal digital identity — eliminating paper documents, sensitive information, and an abundance of unnecessary credentials, while providing them a secure, portable and cryptographically signed identity credential that can be shared across domains such as government agencies and services without the need to expose sensitive information. These approaches handle privacy issues and also mitigate administrative costs, as well as allow for more efficient service delivery for multiple facets of government.

Blockchain-based voting systems provide a secure and tamper-evident record of ballots without compromising voter privacy, offering unprecedented levels of security and transparency in electoral processes. Though reader implementation challenges remain, especially around accessibility and integration with existing electoral infrastructure — permission blockchain voting platforms guide privacy and integrity issues in controlled environments. For example, blockchain-based systems used in government procurement processes can ensure the highest level of transparency throughout the tendering, bidding, and contract management lifecycle, which reduces the risks of corruption and guarantees fair competition among vendors. Blockchain platforms with immutable timestamps for creative works, patent applications, and trademark registrations have revolutionized the management of intellectual property. These systems offer creators and inventors objectively verifiable proof of priority dates as well as allow for the automatic licensing and distribution of



Notes

royalties through smart contracts. Industries reliant on intellectual property rights, including music, film and publishing, have ushered blockchain solutions for closer ties between creators and consumers while ensuring appropriate compensation for rights holders. The issuance, verification, and sharing of education credentials and professional certifications through permission blockchain networks is almost a natural fit, where academic degrees, professional licenses, and skill certifications can be efficiently managed. By digitalising credentials, these systems end diploma fraud and put credential portfolios in the hands of individuals, making it easier for employers to verify credentials and decreasing administrative overhead for educational institutions.

Implemented Permission Blockchain Platforms to Improve Transparency and Accountability Throughout the Donation Lifecycle These systems that track funds from donor to beneficiary create verifiable records of how resources are allocated and utilized. For international aid organizations working in challenging settings, blockchain-based tracking gives donor the confidence that their funds reach intended beneficiaries, crucial to encouraging donors to support humanitarian initiatives. Some global production networks are using permission blockchain solutions to self swab all the production lines, as well as to address quality control and compliance documentation challenges in the manufacturing sector. For complex products such as cars and jets, component tracking systems keep a complete digital catalog of the provenance of all components used in the product, along with specifications, maximizing regulatory ease and maintenance efficiency throughout the product lifecycle. Access to tamper-proof service records stored on blockchain improve resale value and help ensure regulatory compliance for the critical machines.

Blockchain applications which link farmers, processors, distributors, retailers and consumers through transparent information sharing have aided in agricultural supply chains. Such systems can validate sustainable agricultural practices such as organic certifications and fair trade compliance, generating value for producers who engage in responsible methods while providing consumers with confidence in product claims. In developing regions, some blockchain platforms have provided greater direct market access and more equitable compensation for smallholder farmers,

which serves to mitigate reliance on traditional intermediaries. However, the realization of these different use cases has provided a number of shared benefits across sectors. The advantages of permission blockchains like Hyperledger Fabric are two-fold, first by creating a common and canonical source of data that cannot be altered after a certain point and second, through a reduced need for reconciliation and query systems which only look at their individual datasets to resolve any discrepancies. Compliance documentation has been simplified with immutable ledger-keeping that meets regulatory needs while minimizing administrative burden. Identity has been improved between businesses through records that are transparent, verifiable, and eliminate the need for costly third-party verification providers. However, many factors determine whether a permission blockchain solution can be successfully implemented in different settings. It is crucial to analyze the existing complexity in processes before moving onto adopting blockchain, to avoid digitizing processes that are fundamentally broken or inefficient in the first place. Different industries and jurisdictions impose stringent requirements termed as regulations, which need to be adequately reviewed and documented while designing the solution. This allows organizations to take advantage of each other in creating duplicate infrastructure without other organizations lingering close by to their sensitive data. The best permission blockchains we see being executed in practice did not simply arrive at the playground without some curated best practices during development and implementation. By prioritizing finding clear business problems over technology-driven approaches, blockchain solutions will meet real needs instead of looking for how to apply promising technology. Since blockchains are designed to remain performant compared to other universal systems, prioritizing interoperability from the beginning by utilizing industry standards and open protocols ensures that these blockchains can connect to other systems and networks, as both evolve over time, thus increasing the long-term value of blockchain investments.

When it comes to governance, these elements should all be decided upon, agreed upon, and set in stone long before building out the technical aspects of a project so that everyone understands their rights and responsibilities within the blockchain ecosystem and disputes or operational discontinuity don't take place once the system is deployed. Adopting an agile approach to



Notes

development and deployment, starting with MVPs and gradually increasing the feature set based on user feedback, minimizes project risk and helps keep alignment with stakeholder desire. User experience will be your top priority so that you create interfaces that mask the technical complexities of the blockchain system from end users, which leads to higher adoption rates and user experience leading to higher operational efficiencies. Emergence of Trends With Maturity of Permission Blockchain Technology This will facilitate greater data accuracy and improve the reliability of blockchain records by increasing integration with Internet of Things (IoT) devices for more automated data collection and verification, rather than relying on manual input. New applications crossing traditional industry silos are generating cross-industry value networks connecting previously disconnected business networks. A continued expansion of regulatory frameworks recognizing blockchain-based approaches is also taking place, with jurisdictions such as the EU already adapting their legal structures to acknowledge that records on a blockchain can be a legally binding form of data storage.

Unit 14: Smart Contracts and Their Limitations

5.3. Smart Contracts

Among the most disruptive elements of blockchain technology are smart contracts, which used code to execute agreements. Smart contracts have unique properties for permission blockchains, as these platforms are controlled and identity-aware.[2] Smart contracts deployed on permission networks are working in environments where the participants in the network are known and verified, allowing for smart contracts to perform significantly more advanced, private, and regulatory compliant automation of business policies (as compared to their public blockchain counterparts). In essence, smart contracts are self-executing programs in the blockchain that automatically spell out the terms of an agreement, once some predefined conditions are met. They serve as digital contracts with embedded execution mechanisms, removing the need for third parties to monitor compliance and execute actions. By enforcing contracts seamlessly without an intermediary, this ability dynamically restructures the nature of contractual relationships, as it lessens the dependency on trust between parties, decreases execution latencies, and lowers administrative costs due to less overhead in contract management. Permission blockchains are enterprise platforms, and that is reflected in the architecture of their smart contracts. In addition, smart contracts on a public blockchain are executed in a fully transparent environment where code and data are visible to all the participants in the network, whereas permission blockchain implementations often include complex privacy mechanisms. These include private transactions, which hide execution from parties not involved in the transaction and zero-knowledge proofs, which allow for the verification of contract execution without revealing the original data. This architectural design allows companies to automate sensitive business processes while keeping sensitive data private. On the network layer, the blockchain infrastructure is only accessible to authorized participants, creating a baseline level of security. The contract layer contains further permission mechanisms that specify which participants are allowed to deploy contracts, invoke specific functions or have access to specific data elements. The execution environment layer provides isolation mechanisms that disallow unauthorized access to the contract state and its execution logic, thwarting attack vectors.

Both in practice and in theory, development of blockchain smart contracts for permission environments is many formalized than for standard, public blockchain environments, which follows their enterprise context. Requirements gathering: involves detailed consultations with business stakeholders, legal experts and compliance officers to make sure contract logic properly encodes business rules and regulatory constraints. Formal specification techniques are typically used in design phases to help define contract behavior in detail and allow full validation prior to implementing behavior. We use specialised domain-specific languages for the actual implementation that can be optimised to express business logic or you work with well established general-purpose languages with suitable security boundaries. Permission blockchain smart contract testing strategies are broad, with unit tests of individual functions, integration tests of contract interactions, security audits to discover vulnerabilities, and tests through simulation to confirm behavior against range of scenarios. For critical contracts, formal verification, which is resource-intensive, is increasingly being used to mathematically prove correctness and security properties. Governance controls within deployment processes may enforce cross-party approval ahead of contract operationalization, in alignment with organization policies.




Figure 9: Smart Contracts in Blockchain

Governance of smart contracts on permission blockchains doesn't end once the code is deployed, but in fact must start with the code and go through the entire lifecycle of the contract. Contract code changes are tracked in version management systems, which maintain historical records that streamline audit and compliance processes. One area that has been addressed by a variety of upgrade mechanisms is the support of controlled modifications to deployed contracts. As business needs evolve or vulnerabilities are discovered, a changing business environment can lead to the need for contract modifications, but the immutability of contracts on the blockchain can complicate this management on public blockchains. Sources: Monitoring systems track contract operation in real time, creating alerts when unusual patterns that could indicate errors or malicious activity emerge. A practical implementation issue is the integration of the already obtained values from permission blockchain smart contracts with the existing enterprise systems. Oracle services offer secure connections to external data sources and allow contracts to retrieve real-world data including commodity prices, weather data, or shipping status updates. Integrating with enterprise resource planning (ERP) systems allows smart contracts to interface with existing operational processes, bringing alignment between blockchain and conventional systems. Legal integration frameworks define the type of relationship between a smart contract code and a legal agreement, thereby providing clarity on how the automated execution of the smart contract relates to the legal rights and obligations authorized by the legal agreement. In permission blockchain environments, certain types of smart contracts have demonstrated considerable value. Contracts with conditional payments automatically release funds as soon as certain conditions are met, which can simplify financial processes in trade finance, insurance claims, and even service-level agreement enforcement. Supply Chain Contracts: Manage a multi-party workflow; contracts trigger actions based on shipping milestones, quality verification results, or regulatory clearances; these contracts will lead to more responsive and transparent supply networks. Identity and access management contracts are clever because they automatically enforce complex rules for permissions to access resources, which can change depending on who you are, your transaction history, or some outside factors.



Notes

Smart contracts that are focused on compliance will automate the enforcement of regulatory and compliance requirements with business logic that has been programmed into the smart contracts. In highly regulated industries, such as financial services, healthcare, and energy, failures to comply with contractual terms can lead to large penalties, making these contracts especially valuable. These contracts encode compliance requirements directly into execution logic, creating systems that are "compliance by design" and reducing both regulatory risk and the administrative burden of compliance documentation. However, deploying smart contracts on permission blockchains is not without its technical challenges. In order to maintain consensus, all nodes must reach the same computation for a particular contract code execution, motivating strong requirements for deterministic behavior. Although not mentioned here, resource consumption must be controlled in such a way that systems won't slow down, especially in the case of complex contracts that may take considerable computing resources. Yet communication between diverse contract systems is still not as easy as it should be, though standards have begun to emerge to allow for contracts operating on a variety of platforms to communicate with each other. Special care needs to be taken when dealing with error handling in smart contracts because all transactions on the block chain are immutable. While with software it may be possible to roll back the database, a contract on a blockchain requires that the exceptional conditions be explicitly defined within the code — whether in terms of fail-safe protocols, graceful degradation paths, or penalties. Contract developers need to completely understand possible failure modes and the right responses that can keep business operations running while protecting the data integrity. The security of permission blockchain smart contracts requires special attention, as the security of standard software practices is not sufficient to address the blockchain-specific vulnerabilities. These access control bypass analogues create an opportunity for access control exploits—the once evil access control bypass exploits that require a complex layered defence based not only on technical constructs, but also governance controls. One such attack is a re-entrancy attack which allows an attacker to execute functions multiple times and in ways that were not intended, such attacks require certain implementation patterns to be secure around state changes. Oracle manipulation — an attack that can work when

external data sources are compromised — merits defense-in-depth tactics like independent data sources and anomaly detection systems. Smart contracts are still a growing field and different countries are looking for ways to modify current legal frameworks to include blockchain automation. Many jurisdictions have started to acknowledge smart contracts as legally enforceable under certain conditions, albeit specific criteria differ based on electronic signatures, contract formation evidence, and dispute resolution mechanisms among others. Smart contracts are recognized through various Modules in different legal frameworks to be incomplete in themselves and require traditional contracts to cover the decision-making and processes that cannot be entirely automated.

Because smart contracts execute automatically and on a distributed basis, the challenges of dispute resolution for smart contract-based agreements are unique. Cryptocurrency-related issues can often get complicated for traditional courts and, as a result, arbitration services with blockchain know-how are likely to come up in such cases. A few permission blockchain ecosystems have exercised on-chain governance mechanisms to allow disputants to resolve certain types of disputes between them by participating in a renewed consensus process where the majority rules, reserving more complex disputes for the traditional legalisation of the courts, with this latest system allowing for more routine issues to be settled by the herd of peers as opposed to the judiciary of the courts. This interaction of smart contract and traditional legal contracts has led to several modes of implementations. Smart contracts link with traditional contracts as code-plus-contract models, where technical representation and legal representation remain distinct but interrelated — smart contract code automates execution while traditional contracts provide the relevant legal context and governance frameworks. Legal-to-code translation methodologies adopt systematic approaches to transform the legal language into executable code, establishing traceability relationships between the contract clauses and the code components. Natural language processing techniques are progressively being deployed to detect possible inconsistencies between legal text and smart contract deployments, minimizing the chance that the automated execution would escape the bounds of legal intent. While permission blockchains do reduce costs in many cases due to fewer resources needed to validate transactions, the business value gained from the implementation of smart contracts goes



Notes

beyond cost savings. Therefore, after the automation takes over the processing, micro-transactions which only make sense in an automated world can become viable and every single service can become a true commodity for every bit of usage where significantly less transaction cost can make a process economically feasible where it was no longer before. Smart contracts tied to relevant sources of data can be deployed to implement dynamic pricing mechanisms that adjust terms, thereby creating more efficient markets for goods and services based on real-time conditions. When monitoring and settlement are automated via blockchain-based systems, performance-based agreements that trigger compensation based on concrete metrics become actionable.

A few industries have applied permission blockchain smart contracts in particularly innovative ways. Investment contracts for automated loan syndication have been executed by financial services firms with multiple lending institutions in which capital contributions, interest payments, and covenant monitoring are coordinated across the different lending institutions, facilitating a complex financial arrangement. Case in point: insurance firms have utilized parametric insurance contracts, which automatically handle claims based on weather data or other quantifiable parameters, in many cases cutting claims processing from weeks to minutes. Multi-party quality assurance contracts have been created by supply chain organizations that validate product attributes at each phase of manufacturing, transportation, and distribution, confirming that consistent standards are maintained throughout perished product lifecycles. Automated grid management contracts in the energy sector synchronize production, consumption, and storage of electricity across distributed resources to operate increasingly complex energy systems more efficiently. Healthcare organizations have implemented patient consent management contracts that provide individuals with granular control over how their health information is shared among providers and automatically enforce privacy preferences across provider networks. Automated, transparent procurement contracts allow government agencies to evaluate bids from potential vendors according to predetermined criteria and to select and pay the vendor in accordance with imprinted guidelines, thus significantly decreasing administrative overhead and increasing accountability.

Moving forward, a few trends point to the continued evolution of smart contract functionality on permissioned blockchains. Natural language interfaces are being developed where business users without programming expertise can create and modify their own smart contracts simply by engaging in controlled conversations, democratizing access to contract automation. Artificial intelligence is integrated into smart contracts, enabling them to redefine themselves or their tasks, creating some form of flexible automation that can respond appropriately to complex situations more clearly than before, touching towards the boundaries of human capabilities in relatively undefined situations while still keeping up with the efficiency of automation. Cross-chain contract protocols allow for smart contracts to be configured to interact with contracts deployed on other networks, opening up the opportunities for more complete process automation across the organizational and technological borders.

Smart contracts and their regulation are still in progress and many legal systems are working on niche regimes that could suit smart contracts and give them a new level of priority and enforceability. These regulatory developments are increasingly aimed at striking the right balance between enabling innovation and having adequate safeguards in place, as smart contracts have the potential to increase the efficiency of the business process while having adequate protections in place for market participants. A slew of industry-specific standards bodies are establishing best practices and reference implementations that encourage security, interoperability, and regulatory compliance helping to expedite adoption across sectors. In this context, permission blockchain smart contracts are a potent marriage of the trust mechanisms of blockchain and enterprises' need for control, privacy and compliance. Through enabling the automated, transparent execution of business agreements and contracts, smart contracts are changing the way that organizations collaborate, transact, and manage processes within organizations. As we continue to experience new developments in this space, with growing maturity of the tech and deeper integration into existing frameworks of delivery, it is reasonable to expect that smart contracts on permission blockchains will play a central role in our digital business future, driving efficiency in operations and in some cases, enabling completely new forms of value-creation across sectors.



5.4. Design Limitations

Design limits are one of the limiting factors that affect the mechanics and performance of distributed systems. Such restrictions arise from some material commerce expected in dispersed computing, in which program creators really need to strike a balance between competing aims like persistence, availability and partition tolerance. This phrase describes three properties of distributed systems, and according to the CAP theorem, formalized by Eric Brewer, no distributed system can simultaneously satisfy these three properties during a network partition. This inherent limitation leads system architects to make carefully considered tradeoffs in terms of guarantees they concatenate, based on the application-specific use-cases. There is yet another aspect to limitation arising from performance considerations where we must also contend with network latency and bandwidth constraints, and computational overhead on the various nodes to be connected in the distributed system. As the system scales, these components can have dramatic effects on system throughput and response times. Resource limitations further limit the design for distributed systems such as hardware functionalities, memory limitations and energy requirements. Well behaved modern distributed systems must operate effectively across a wide variety of computing environments from data centers to edge devices with disparate resource profiles. Additional consideration is introduced by security and privacy requirements, as stronger security protocols are inherently more complex to implement and may degrade the performance of the overall system. Designers have to balance the security proposals with performance drawbacks. This is a basic design issue, where systems validated to work at small scale design often experience substantial failures when expanded. This limitation presents itself in forms such as additional coordination overhead, resource contention, and management complexity. Design choices become more challenging in distributed settings due to their heterogeneity: systems must function in various hardware, software, and network environments. This can create inconsistent performance profiles and unexpected failure modes. Another limiting factor is administrative boundaries, as distributed systems often cover numerous organizational domains that have their own policies, priorities, and operational constraints. This can lead to additional

complexity in system management, system update, and consistent policy enforcement. Due to the constraints in time synchronization (i.e., it is impossible to have perfect synchronization of clocks in distributed nodes), distributed algorithms that depend on truths provided by perfect clocks are affected by the system's symptoms. This limitation informs the design of consensus algorithms and coordination mechanisms. Design options are also constrained by deployment considerations like back-compatibility, migration paths, and zero-downtime upgrades. Design systems to evolve through crawl, walk, run phases without disrupting current operations. Recovery capabilities become a critical design constraint, as systems should be able to keep data integrity and service availability even though failures can and will happen. 3 It complicates the recovery mechanisms and can affect the performance of the normal operations. Challenges in Monitoring and observability in distributed systems In a distributed system, obtaining a complete view of system state at any given point in time is difficult. Lack of visibility compromises debugging, performance optimization, and anomaly detection. There are inherent limitations to testing, which makes it difficult to validate distributed systems because it is impractical to reproduce every failure scenario, and every execution path. This restriction hinders full verification and may cause transient bugs to be found later in production when they cause problems. Configuration management adds even more complexity, since many distributed systems have many configuration parameters that interact in non-obvious ways. However, determining the best configurations for specific workloads and environments stays complex. Because programming a distributed system involves dealing with concurrency, partial failures, and asynchronous communication patterns, this is inherently more complicated than programming under sequential programming models — something that has an effect on both developer productivity and system correctness. These design constraints form the basis in the architecture of distributed systems, leading to various design trade-offs and constraints depending on the application requirements and operational context. An important design constraint for distributed systems is consistency model limitation as stricter consistency guarantees usually sacrifice availability or increase latency. The system designers must ensure different models of consistency are chosen according to the semantics of the application and the degree of tolerance to stale data. The consensus algorithms are designed to achieve one of various consistency goals —

from linearizability (strong consistency) to one of various eventual consistency models — with different trade-offs between correctness guarantees and performance characteristics. Network constraints fundamentally limit distributed system design — introducing unpredictable latencies, limited bandwidth, and network partitions. Such constraints can impact communication patterns, data replication strategies, and coordination mechanisms. Designers write systems to work efficiently in these network conditions and here lies a huge part of the complexity — retry algorithms, timeouts, and consensus protocols lead to incredible developer burnout. An inadequate ability to detect failure affects system's reliability since determining whether a remote component has failed is difficult as it can be hindered by network delays and message losses. This leads to unavoidable uncertainty regarding the current state of distributed components, needing algorithms which can work correctly despite imperfect failure detection.

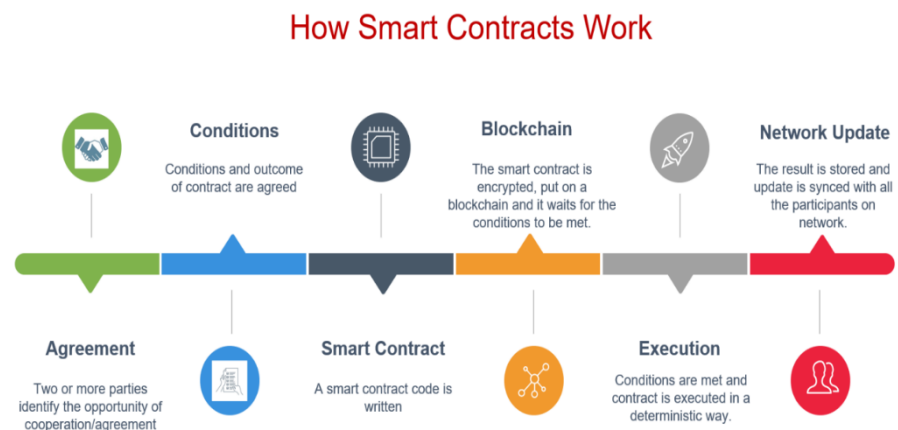


Figure 10: Working Procedure of Smart Contract

Composable state management limitations affect the design of distributed systems, especially when stateful services should share consistent data among multiple nodes. There are a lot of ways state can be modified in a distributed system (concurrent, partitioned, or general network, or node failure, etc), and coordinating all of that, whilst ensuring consistency is a huge and complex task. While there are approaches like CRDTs conflict-free replicated data types and consensus algorithms to address these challenges, each has its own limitations and trade-offs. Versioning and

history management create additional design problems, particularly in systems that allow offline operation or implement optimistic concurrency control. Logging the divergent operation histories and reconciling that over time is storage overhead and added complexity. Periodic or per-committed deployment and upgrades can impact the evolution of a system, as few distributed systems can be upgraded atomically. This requires that you plan managed at various levels of compatibility between each version of software that may be running at the same time during rolling upgrades. The distributed nature and non-determinism of distributed systems complicate debugging and troubleshooting, often making them difficult if not impossible to reproduce. These limitations compound the trouble of root cause analysis, because this kind of analysis is more complex and time-intensive than single-node systems. Especially for those that handle sensitive data or exist in regulated spaces, regulatory and compliance constraints are becoming a significant factor in disaggregating/designing a distributed system. Data localization, access controls, audit trails and retention policies all impose constraints on architectural choices and implementation methodologies. Transformability: Physical infrastructure limits how systems can be deployed and operated – power availability, cooling requirements, and physical security are all factors, for example. These limitations apply specifically to edge computing use cases and geographic distribution. Similarly, human factors constrain the design and operation of systems because the complexity of distributed systems consistently reaches beyond the level of understanding and management of operators and developers. This human cognitive limitation creates the need to invest in automation, monitoring, and self-healing capabilities, leading to reduced operational burden. Aspects such as energy efficiency, carbon footprints, and sustainable resource usage are becoming progressively important in the design of distributed systems. These factors can limit decisions about data replication, processing locality, and hardware selection. This design freedom comes with a price; the design constraints exhibit the complexity of distributed systems and the multitude of subtleties involved that an architect needs to balance. It is critical to understand these inherent limitations of distributed systems and making deliberate decisions that follow the principles of your app needs, operational boundaries and business goals.



Unit 15: State Machine Replication in Blockchain

5.5. State Machine Replication

State Machine Replication (SMR) is an old but fundamental building block of distributed systems that provides fault-tolerant, highly available systems. SMR itself is centred around abstracting a system as a deterministic state machine and replicating this state machine across several nodes with the ability to continue operating in the event of node failure. This is based on the concept that two deterministic state machines with the same initial state and responding to the same input sequence will go through the same state transitions and produce the same outputs. Atobh, SMR offers a convenient abstraction for implementing reliable distributed services under a variety of failure scenarios while preserving consistency. SMR's theoretical underpinnings were laid down by Leslie Lamport in his groundbreaking work on distributed systems, which introduced the key components for keeping state across replicated processes consistent. Since then, these foundations have been extended and honed to tackle practical implementation challenges and to maximize performance for a wide range of deployment scenarios. A state machine replication system is comprised of many replicas, where each replica holds a copy of the application state and has the logic to process client requests. Clients interact with the system via a replication protocol that guarantees that every non-faulty replica processes the same sequence of operations. Usually this protocol is based on total order broadcast (also called atomic broadcast) to reach agreement on the ordering of operations between all replicas. These ordered operations are applied deterministically, meaning that all replicas are guaranteed to pass through identical states, ensuring the system is consistent. In SMR, the processing of requests is performed as follows: clients send requests to the system, then the replication protocol orders and delivers these requests to all replicas, subsequently each replica executes the requests in this total order, and then the system replays the responses to the clients. This approach guarantees that all replicas perform identical operations in strictly the same order, achieving synchronized state in the system, even in the presence of failures or perceived network latencies. This will vary depending on the implementation of the state machine replication protocol used for consistency guarantees. Most SMR systems implement strong consistency

semantics, often linearizability, giving the illusion that each operation emits as an atomic operation, as well as in real-time order from the clients' perspective. This strong consistency model eases the burden of developing applications on top of a distributed system by exposing an interface that resembles a single-node system, freeing developers from reasoning about the distributed execution and not having to worry about the details of partial failures or replication. But this consistency has trade-offs on performance and availability, especially during network partitions or when it encounters high network latencies between replicas. The main motivation of state machine replication is fault tolerance, meaning that even if some copies of the system fail, the system can still work correctly. The fault model that an SMR implementation targets dictates its resiliency properties. Crash fault-tolerant (CFT) systems tolerate up to a certain number of replicas failing by crashing (ceasing operation), while Byzantine fault-tolerant (BFT) systems further tolerate arbitrary (potentially malicious) behavior from up to some fixed number of replicas. The fault tolerance threshold usually mandates that a majority of replicas must stay alive and connected, and BFT systems, in general, require a higher replication factor for the same degree of failure tolerance than the CFT systems. State machine replication involves practical considerations of implementation (e.g., performance, correctness). Moreover, state transfer mechanisms allow newly added or recovered replicas to catch up on the system state before they start processing requests. State transfer protocols are designed as efficiently as possible to minimize the performance hit due to system recovery or reconfiguration. Checkpointing strategies periodically take snapshots of the system state to restrict log growth and to speed up the recovery of replicas. These checkpoints should be the same across the replicas and incorporated into the replication protocol for correctness reasons. voted, or elected if in a leader-based protocol, in case the current leader fails or is not reachable. Such reconfiguration procedures need to maintain the ordering guarantees of the replication protocol while keeping the disruption to the running applications at a minimum. Such optimizations can be done using request batching and pipelining by amortizing the overhead of the consensus protocol across multiple client requests thereby improving system throughput. It is important that the optimizations are introduced in a controlled way, as we cannot compromise the safety properties of the replication protocol, and we would like to improve performance.



Notes

Application considerations — we generally want our application logic to be modeled as a deterministic state machine. This means anything that is non-deterministic (e.g. generating a random number, creating a timestamp, scheduling a thread to run) has to be handled in a specific way to guarantee that all replicas will execute the same command under the same conditions. The most common ways are to either externalize the non-deterministic choice using the replication protocol (e.g. Activiti, BPMN) or using pseudo-random generators over the same seeds by all replicas. As the state model is materialized, it has to over time grow, which needs to be handled using state compaction, or garbage collection (GC) — or hierarchical storage similar — to ensure the system works efficiently. Optimizing performance for SMR systems requires tuning a variety of parameters (e.g., batching size, timeout values, threading model) to yield the best throughput and latency characteristics for the workload at hand. However, in order to monitor and debug SMR systems, specialized tools and techniques are needed to track distributed execution, detect divergences, and investigate performance problems across many replicas.

State machine replication is used in a lot of fields that demand strong consistency guarantees and high availability. Distributed coordination such as Apache ZooKeeper and etcd is an implementation of SMR which provides reliable metadata storage and distributed synchronization primitives for large-scale distributed systems. SMR is used by distributed databases (such as Apache Cassandra, etc.) for replicating transaction logs or controlling configuration metadata, enabling the domain of data management to remain consistent in the event of failures. These systems often need to reach agreement between untrusting parties, which is where blockchain technology, a particular form of state machine replication using consensus, was developed to help maintain a consistent distributed ledger. SMR is employed in critical infrastructure services, such as air traffic control systems [7], financial transaction processing [8], and telecommunication control planes [9], to meet the high reliability required of these mission-critical systems. Examples of cloud computing providers leveraging SMR for internal use include managing cluster state, for service discovery, and for configuration management. As new challenges and application requirements arise, state machine replication continues to be

evolved. The early SMR implementations prioritized mostly correctness and fault tolerance, usually sacrificing performance and scalability. Recent trends focus on whether state-of-the-art performance can be achieved with more or less strong consistency guarantees, new techniques such as speculative execution, locality-aware request routing, and adaptive batching strategies. Over time, Byzantine fault-tolerant protocols have made advances in decreasing communication complexity and resource use, leading to the development of protocols that are more practical for real-world usage. Recent directions of SMR exploration include geo-distributed deployments optimized for low-latency operation across wide-area networks, resource-efficient implementations suitable for the constrained environments of edge computing, and composite (hybrid) architectures that tune replication strategy based on operation properties (e.g. writes vs reads) or consistency requirements. Finally, the combined use of SMR with emerging technologies such as trusted execution environments (TEEs) and formal verification tools holds great promise for further strengthening the security, correctness, and performance of replicated state machines in the distributed systems of tomorrow.

5.6 Replication of a distributed state machine

But Distributed State Machine Replication (DSMR) expands the classical state machine replication paradigm to solve the problems inherited from microservices systems of even larger scale, geographically distributed. Traditional SMR works only within a single data center or a single cluster, whereas DSMR includes techniques for replicating state machines across geographic regions, administrative domains, or heterogeneous environments. This extension adds new complexities in terms of wide-area network latency, partial connectivity, and heterogeneous infrastructure capabilities. Similar to traditional SMR, the primary goal of DSMR is to preserve consistent, fault-tolerant service via replicated deterministic state machines across multiple nodes. As a result, they also require significant innovation in consensus protocols, state management, and client interaction models to deploy at sizes with both acceptable performance and reliability characteristics. DSMR combines theoretical foundations of distributed systems such as eventual consistency, causal consistency, and network partitioning graphs, with elements from consensus algorithms, quorum



Notes

systems, and distributed databases to produce real-world building blocks for replication at a global-scale. Distributed state machine replication systems adopt a hierarchical or partitioned architecture to abstract complexity and improve performance across geographic bounds. Hierarchical architectures typically arrange replicas into local clusters that reach internal consensus, while coordinating between clusters at a higher level with possibly relaxed consistency guarantees, and they often implement some of these appearing consistency guarantees at much higher rates than possible in a flat architecture.

This solution minimizes the operation latency due to wide-area network latency, yet still preserves the strong consistency inside the region. In partitioned architectures, the state space is partitioned between multiple groups of replicas where each group is responsible for some subset of the whole state. This sharding concept increases scalability by making it possible to process independently the operations that do not influence the same state partition (although this creates some complexity for operations that span several partitions). Hybrid methods issued accordingly to the target applications and environment constraints, mixing hierarchical and partitioned approaches. The different architectural styles of DSMR systems allow them to operate at a large scale and deal with the limitation described by CAP Theorem in order to satisfy the needs of different applications. In distributed state machine replication, consistency models take on different forms, with most being some version of linearizability, but softer in terms of guarantees or design constraints, particularly for wide-area deployments. Most multi-level consistency models provide separate guarantees for local and global operations, and a guarantee that is stronger than the other for operations that fall in the regional constraints and relaxed for those that cross a region. This way acknowledges that the cost of strong consistency scales rapidly with geographic distance. Session consistency guarantees retain causal ordering only in the context of the session while allowing reordering globally (that is across sessions), making it a pragmatic solution to give away between strong consistency guarantees and eventual consistency models. Strong consistency is not always achievable due to scaling, so conflict resolution mechanisms address concurrent modifications, utilizing methods such as CRDTs, operational transformation, or

application-specific merge functions to reconcile differing versions of the state. Such approaches enable systems to make progress during network partitions or reasonably high latencies while offering well-defined semantics for eventual state convergence.

Replication protocols for distributed state machine replication build on traditional consensus algorithms to operate effectively in the wide-area with heterogeneous latency and reliability characteristics. Such geo-aware consensus protocols are optimized for wide-area deployment, by minimizing cross-region communication, using regional leaders, or fast local decision paths whenever possible. For availability, only a subset of regions need to participate with consensus in order for the region to make progress, potentially exchanging some of this data consistency with progress. Coordinate users across partitions or regions: Multi-phase commit protocols coordinate distributed transactions across partitions or regions to ensure atomicity and consistency for operations spanning multiple state segments. Such protocols generally consist of prepare, validate, and commit phases to guarantee all members eventually apply or abort a transaction consistently. These leaderless replication strategies spread participant requests over several replicas FT without a single coordinator to manage them, reducing latency but complicating request conflict resolution. This would often combine with some causality tracking that keeps track of operational dependencies in order ensure that state resolution occurs in a coherent manner. In distributed state machine replication, state management coping with issues concerning the responsibility of storing and replicating large volumes of states across geographically distributed replicas. Partial replication strategies enable replicas to hold only a portion of the full state according to access patterns, locality requirements, or resource constraints. This leads to reduced overhead in terms of storage and synchronization, but potentially introduces latency when accessing non-local state. Hierarchical state organization sub-divides the state across hierarchical tiers with different replication factors, consistency guarantees, and access patterns. As an example, metadata that is frequently accessed could be fully replicated across all regions, while larger data items could use more targeted replication strategies. Dynamic state migration mechanisms relocate state partitions for access patterns, network conditions, or administrative policies. This allows the systems to adapt to workload variability while ensuring the



Notes

consistency guarantees for any time during the transition. To mitigate this latency, cache coherence protocols place and manage copies of remote state at local caches, which can potentially take on invalidation or update propagation schemes in order to keep copies of remote state consistent across distributed caches.

Performance optimization for distributed state machine replication includes a wide range of techniques to reduce latency, achieve high throughput, and effectively utilize system resources across geographic regions. In particular, locality-aware request routing, when support by the underlying consistency and replication mechanism, routes client operations to the closest replica when possible, reducing access latency and minimizing the network utilization, but often at the cost of added complexity in strong consistency guarantees. Non-critical updates can be replicated along asynchronous replication paths. As such, many execution techniques sought to predictively apply operations in advance of reaching full consensus, providing client temporary results until consensus is reached where they could be confirmed or rolled back. For read-heavy workloads, this approach can drastically minimize callers' perception of latency. Dynamic batching and pipelining strategies adapt how operations are grouped based on current load, present network conditions, and consistency guarantees, whether needed or desired, thus optimizing throughput and balancing latency tradeoffs. Workload-aware optimization adapts replication policies to different types of operations, which may result in use of separate consistency models or separate replication factors according to operation semantics, application importance, or performance guarantees. Distributed state machine replication tends to work on failure scenarios all over the world, from region-wide outages, between regions, regional- and sub-regional partitions, asymmetric connectivity, etc., as opposed to traditional SMR. Regional fail-over mechanisms, for example, make use of region-agnostic disaster recovery procedures to keep the system available even when the whole region is down, by promoting a secondary replica or reconfiguring the consensus group. A multi-region consensus configuration while providing the semantics of each regions are built on its own network towards the consensus could still be able to achieve system progress if a quorum of the regains connectivity once again, often at least a quorum of the regions

required and sufficient partition tolerance. This principle allows for availability but may lead to donkey conformance and operation on an inconsistent state until the partition is over. The Partial connectivity handling describes how the system behavior adapts if some (but not all) inter-region links fail or perform poorly. Load balancing and scaling: smart load balancer could dynamically adjust quorum configuration, operation routing, or consistency level based on the statistical features of the current cluster network condition.

Asynchronous state reconciliation mechanisms are used to repair splits in state after connectivity has resumed after network partitions by systematically extracting and reconciling differences with the goal of minimizing impact on current run-time operations. Abstractions for the usability of client interaction models for distributed state machine replication that mitigate the complexity of geographically — not necessarily temporally — distributed systems. Regional affinity models direct clients towards specific operational regions, regardless of regional locality, client localization, or administrative policies, improving performance but at the expense of flexibility during regional failure scenarios. The ability of consistency-level selection enables the clients to specify the desired consistency guarantees for individual operations, which in turn allows application-level trade-offs between performance and consistency as a function of the operation semantics. Dynamic routing of requests to all replicas allows operation management at any latency-sensitive level of shared state, or even transparently through several existing state replicas with a possible acquisition of an interim result, as response generated upon mediation to avoid violating strict consistency conditions. Session-based interaction models allow the client to maintain their context across several operations, and offer guarantees such as read-your-writes consistency, as well as sequential consistency with respect to one client session to another, while allowing operations from different clients to be reordered among themselves.

This led to the existence of practical implementations of distributed state machine replication found in systems that are aimed for global wide deployment. One of the first global distributed databases is Google Spanner, which achieves external consistency (linearizability) across many regions



Notes

using True Time, GPS, and atomic clocks. The solution focuses on coordination cost reduction but stays within the theme of strong guarantees by overlaying this with time uncertainty bounds. Cosmos DB is a globally distributed multi-model database with tunable consistency levels provided by Microsoft, enabling data to be served from the closest data center, granting applications the ability to make appropriate trade-offs between consistency, availability, and performance for different workloads. TAO at Facebook provides a distributed to the social graph, with multi-tier caching and region replication, die for the global infrastructure and access patterns of the company. CockroachDB is a geographically distributed SQL database that uses a multi-layer architecture that combines Raft consensus between nodes within regions with asynchronous replication between regions to afford strong consistency for operations local to a region, and eventual global convergence as the global state is updated by replicas in the cluster distributed across regions. These patterns illustrate varying ways of navigating the core trade-offs that exist in distributed state machine replication towards meeting particular application requirements and operational constraints in practice. All these challenges and new research directions for distributed state machine replication are still evolving as the systems grow to billions of virtual machines and cloud resources and dealing with the heterogeneous environment. Very large-scale replication focuses on the new protocols needed to handle systems with thousands or millions of replicas, which have different requirements for consensus and state management than do traditional SMR systems.

The need for heterogeneous infrastructure support stems from the diverse hardware, network conditions, and administrative domains found in modern distributed systems, which necessitate adaptive protocols that remain functional despite the heterogeneity of capabilities and constraints. Cross-cutting between various domains, Energy-efficient replication analyzes techniques for power consumption and carbon footprint minimization while delivering reliability and performance targets, which has become an ever-more relevant challenge for global-scale systems. To address these threats to geographically distributed systems, DSMR provides security and privacy enhancements, including regulatory compliance across multiple jurisdictions, protection against sophisticated adversaries, and

confidentiality guarantees for replicated state. It means self-adaptive replication protocols that adapt their behavior according to the actual workload characteristics, network conditions and failure patterns, achieving better performance and reliability without the need for manual tuning or reconfiguration. Some of the future directions for this area involve new algorithms, new techniques, new applications. Users of data and machine learning techniques with DSMR systems could create predictive designs, identify outliers, and adapt based on real-world input. If you believe that quantum computers will be available for use to some extent in 2050, quantum consensus protocols can use quantum computing capabilities to obtain different fundamental trade-offs in consistency, security, and performance to classical. Basic edge-cloud integration will scale distributed state machine replication to include cloud data centers and edge devices (that is, IoT), creating uniform replication domains that span highly heterogeneous environments (with respect to communications and resource constraints). This convergence may enable the development of hybrid systems that could leverage the best of both worlds that the Blockchain and DSMR offer. Understanding distributed state machine replication at a deeper level is critical to tackling the unique challenges these evolving architectures will face, as distributed systems become incredibly complicated but increasingly promising.

Multiple Choice Questions (MCQs)

1. **What is a key characteristic of permissioned blockchains?**
 - a) Open participation for everyone
 - b) Restricted access controlled by an authority
 - c) Anonymous transactions only
 - d) No consensus mechanism required
2. **Which of the following is NOT a use case for permissioned blockchains?**
 - a) Enterprise supply chain management
 - b) Decentralized public voting
 - c) Private financial settlements
 - d) Healthcare record management
3. **Which blockchain platform is widely used for permissioned blockchain applications?**



Notes

- a) Bitcoin
 - b) Ethereum
 - c) Hyperledger Fabric
 - d) Dogecoin
4. **Smart contracts are primarily used for:**
- a) Automating agreements without intermediaries
 - b) Mining cryptocurrencies
 - c) Encrypting blockchain transactions
 - d) Creating non-fungible tokens (NFTs)
5. **What is a major design limitation of permissioned blockchains?**
- a) Lack of privacy
 - b) Scalability challenges due to high decentralization
 - c) Dependence on a central authority for governance
 - d) No need for identity verification
6. **State Machine Replication (SMR) in blockchain ensures:**
- a) All nodes execute transactions differently
 - b) The blockchain state remains consistent across all nodes
 - c) Each node follows a unique ledger
 - d) Transactions are confirmed without verification
7. **Which of the following best describes Distributed State Machine Replication?**
- a) A method for centralizing blockchain networks
 - b) A technique to synchronize the state across multiple nodes
 - c) A way to replace consensus mechanisms
 - d) A form of encryption for blockchain transactions
8. **What type of consensus mechanism is often used in permissioned blockchains?**
- a) Proof of Work (PoW)
 - b) Proof of Stake (PoS)
 - c) Byzantine Fault Tolerance (BFT)
 - d) Mining-based validation
9. **Which of the following is an advantage of permissioned blockchains over permissionless blockchains?**
- a) Faster transaction processing and scalability

- b) Completely anonymous transactions
 - c) Higher energy consumption
 - d) Decentralized governance
10. **What role do smart contracts play in permissioned blockchains?**
- a) Encrypting all transactions for security
 - b) Automating processes based on predefined conditions
 - c) Eliminating the need for block validation
 - d) Replacing the consensus algorithm

Short Answer Questions

1. What is a permissioned blockchain?
2. Explain the difference between permissioned and permissionless blockchains.
3. Name two industries that use permissioned blockchains and explain why.
4. What is a smart contract, and how does it function in blockchain networks?
5. What are the major design limitations of permissioned blockchains?
6. How does State Machine Replication maintain consistency in blockchain networks?
7. What is Distributed State Machine Replication, and how does it work?
8. Why are permissioned blockchains preferred in enterprise applications?
9. Which consensus mechanisms are commonly used in permissioned blockchains?
10. How do permissioned blockchains ensure security without mining?

Long Answer Questions

1. What are permissioned blockchains, and how do they differ from permissionless blockchains?



Notes

2. Explain the benefits and challenges of using permissioned blockchains in enterprises.
3. Describe the role of smart contracts in permissioned blockchain networks and provide real-world examples.
4. What are the key design limitations of permissioned blockchains, and how can they be addressed?
5. How does State Machine Replication contribute to blockchain consistency and fault tolerance?
6. Explain the concept of Distributed State Machine Replication and its importance in blockchain systems.
7. Compare Hyperledger Fabric and Ethereum in terms of their suitability for permissioned blockchains.
8. Discuss different consensus mechanisms used in permissioned blockchains and their advantages.
9. What are the security implications of using permissioned blockchains in financial transactions?
10. How can permissioned blockchains be integrated with existing IT infrastructure in businesses?

References

Module 1: BLOCKCHAIN TECHNOLOGY

1. Swan, M. (2023). *Blockchain: Blueprint for a New Economy, Second Edition*. O'Reilly Media.
2. Zheng, Z., Xie, S., & Dai, H. (2022). *Comprehensive Guide to Blockchain Architecture and Implementation*. Springer International Publishing.
3. Antonopoulos, A. M., & Wood, G. (2024). *Mastering Blockchain: Distributed Ledger Technology, Decentralization, and Smart Contracts Explained*. Packt Publishing.
4. Tapscott, D., & Tapscott, A. (2022). *Blockchain Revolution: How the Technology Behind Bitcoin and Cryptocurrency is Changing the World*. Portfolio.
5. Prusty, N., & Mukherjee, S. (2023). *Hyperledger Fabric: Enterprise Solutions for Blockchain Technology*. Wiley.

Module 2: BASIC CRYPTOGRAPHIC PRIMITIVES BEHIND THE BLOCKCHAIN

1. Katz, J., & Lindell, Y. (2024). *Introduction to Modern Cryptography: Principles and Protocols*. Chapman and Hall/CRC.
2. Ferguson, N., Schneier, B., & Kohno, T. (2022). *Cryptography Engineering: Design Principles and Practical Applications*. Wiley.
3. Narayanan, A., Bonneau, J., & Felten, E. (2023). *Cryptographic Primitives for Blockchain Systems*. Princeton University Press.
4. Menezes, A. J., van Oorschot, P. C., & Vanstone, S. A. (2022). *Handbook of Applied Cryptography: Blockchain Edition*. CRC Press.
5. Smart, N. P. (2024). *Cryptography: An Introduction to Secure Communications in the Digital Age*. McGraw-Hill Education.

Module 3: BITCOIN BASICS

1. Antonopoulos, A. M. (2023). *Mastering Bitcoin: Programming the Open Blockchain*. O'Reilly Media.



Notes

2. Popper, N. (2022). *Digital Gold: Bitcoin and the Inside Story of the Misfits and Millionaires Trying to Reinvent Money*. Harper Business.
3. Song, J. (2024). *Programming Bitcoin: Learn How to Program Bitcoin from Scratch*. O'Reilly Media.
4. Rosenbaum, E., & Vigna, P. (2023). *The Age of Cryptocurrency: How Bitcoin and the Blockchain Are Challenging the Global Economic Order*. St. Martin's Press.
5. Burniske, C., & Tatar, J. (2022). *Cryptoassets: The Innovative Investor's Guide to Bitcoin and Beyond*. McGraw-Hill Education.

Module 4: CONSENSUS

1. Fischer, M. J., Lynch, N. A., & Paterson, M. S. (2023). *Distributed Consensus: Theory and Practice*. MIT Press.
2. Wang, W., Hoang, D. T., & Niyato, D. (2022). *Blockchain Consensus Algorithms: Design and Applications*. Cambridge University Press.
3. Pass, R., & Shi, E. (2024). *Consensus in the Age of Blockchains*. Morgan & Claypool Publishers.
4. Castro, M., & Liskov, B. (2023). *Practical Byzantine Fault Tolerance and Proactive Recovery*. Academic Press.
5. Daian, P., Pass, R., & Shi, E. (2022). *Snow White: Consensus from the Ground Up*. Springer.

Module 5: PERMISSIONED BLOCKCHAIN

1. Dhillon, V., Metcalf, D., & Hooper, M. (2023). *Enterprise Blockchain: A Definitive Handbook*. Apress.
2. Androulaki, E., & Cachin, C. (2022). *Hyperledger Fabric: A Distributed Operating System for Permissioned Blockchains*. IBM Press.
3. Gorenflo, C., Lee, S., & Golab, L. (2024). *Permissioned Blockchain Systems: Foundations and Applications*. Wiley.
4. Brown, R. G. (2023). *The Technical Foundations of Enterprise Blockchain Systems*. Packt Publishing.

5. Vukolić, M. (2022). *The Quest for Scalable Blockchain Fabric: Proof-of-Work vs. BFT Replication*. Springer International Publishing.

MATS UNIVERSITY

MATS CENTER FOR OPEN & DISTANCE EDUCATION

UNIVERSITY CAMPUS : Aarang Kharora Highway, Aarang, Raipur, CG, 493 441

RAIPUR CAMPUS: MATS Tower, Pandri, Raipur, CG, 492 002

T : 0771 4078994, 95, 96, 98 M : 9109951184, 9755199381 Toll Free : 1800 123 819999

eMail : admissions@matsuniversity.ac.in Website : www.matsodl.com

