MiColEC Project - “Micro-hubs Colaborativos para a Economia Circular”

Specification of the Blockchain Component

March 2022

Cofinanciado por:

MADEIRA 14-20

REGIÃO AUTÓNOMA DA MADEIRA

PORTUGAL 2020

UNIÃO EUROPEIA
Fundo Europeu
de Desenvolvimento Regional

Document information

Delivery date:
Version: 0.2
Responsible Partner:

Dissemination level

Public

Revision history

Date	Editor	Status	Version	Changes
28 February 2022		Draft	0.1	Initial Draft
15 March 2022		Draft	0.2	Complete Draft

Contributors

Logimade Lda
ARDITI - Agência Regional para o Desenvolvimento da Investigação, Tecnologia e Inovação
Universidade da Madeira
IL Technologies Lda

Summary

In this document we analyze potential blockchain solutions to be used as the support technology for the MiCoLEC micro-hub platform.

The Blockchain is a shared, immutable ledger that facilitates the process of recording transactions and tracking assets in a business network. An asset can be a house, car, cash or a delivery request. But the perfect blockchain doesn't exist. There is not a blockchain technology that solves all problems in any field. Different use cases require different solutions. This report aims to present the best solution based on current technologies available.

We begin by briefly describing the problem MiCoLEC is aiming at tackling and try to map the correspondent technical requirements. From those, we set to identify the most suitable blockchain technology to be used in the project's development efforts. Since this technology space is crowded and not always well categorized and analyzed, we provide an overview of the different commercially available options.

We finalize the document proposing a concrete technological approach for MiCoLEC taking into consideration technology characteristics and also constraints related with the project's scope such as resulting costs. The solution proposed is the one where costs are minimized while still offering sound technological support.

Context

In this project, a micro-hub of logistic companies is expected to participate in a cooperative, while still competitive, package delivery marketplace that will effectively improve their overall performance. This improvement is achieved by gains in efficiency levels (e.g. route sharing), less waste (e.g. deduplication of courier routes), and increased trust in cooperation through the platform’s technology. To this end, an innovative digital platform is proposed to allow trustworthy and accountable information sharing between the participants. In particular, this platform promises to allow, for instance, package handoff between competing companies in a completely transparent and secure way, while ensuring that even if something goes wrong it is possible to trace the package movements and assess liabilities. To support the platform, MiCoLEC proposes a blockchain-based system to manage interactions between participants and allowing most of the processes to be run on smart-contracts that transparently enforce fairness and that cannot be tampered with.

Based on the problem described, it becomes important to choose the right technology to support the platform. In order to do so we must be able to identify the main problem characteristics and map them into technology requirements. In particular, since we are dealing with a platform with multiple participants that do not trust each other, we will be focusing on blockchain technology with its many variants and try to argue which might be the best option for the MiCoLEC project.

Following an analysis of the functional requirement, the supporting technology should at least have the following characteristics:

- low or no fees on transactions. Having to pay a fee per transaction is a problem a delivery service may not support. These fees may be unpredictable and sometimes surpass the service value. The result would be an impairment in system scale. As soon as the system scales to high numbers of transaction throughput, the cost of those transactions would quickly become unbearable for any service provider.

- a scalable solution. A system may need to increase or decrease in performance and cost in response to changes in application and processing demands. It is important how easy these changes can be introduced in the system.

- low barrier entry for new users. No need to set up nodes or incur in high upfront costs. Roadblocks for new users (new logistics companies) to join the system will negatively impact the ability for the project to grow and expand.

Along this document we map the different blockchain technologies available in the market and how they compare with one another. We use the characteristics above in an attempt to assess which technology best fits our objectives while trying to avoid disregarding other characteristics and how they may potentially affect the resulting system. Naturally, we will be focusing our attention on commercially available systems that can be used in practice and are backed by a company or consortium that provides adequate support to platform users and shows evidence of continuous development and improvement of the platform itself.

We begin by providing some background on blockchain technologies and their characteristics. We then try to compare the different approaches on the theoretical level in order to identify the type of blockchain that better fits our purposes. Next, we dive into concrete systems analyzing the state of the art and providing data to allow comparison between the different solutions. Finally, we discuss the solutions we believe better fit MiCoLEC and conclude the document.

Blockchain technology: public vs private

A distributed ledger is a database that can be accessed across several locations or among multiple participants. However, most companies still use a centralized database with a fixed location. Unlike a centralized database, a distributed ledger is decentralized, which helps to remove the need for a central authority or intermediary for processing, validating, or authenticating transactions.

Furthermore, these records will only be stored in the ledger after the parties involved have reached a consensus.

A blockchain is a form of distributed ledger that has a specific technological underpinning. Blockchain creates an unchangeable ledger of records maintained by a decentralized network after a consensus approves all the records.

The most notable characteristic of a blockchain is the maintenance of a cryptographically signed chain of records. The way this chain is created and maintained allows for anti-tampering capabilities, which are key for cooperation between entities that do not want to rely on other layers of trust.

The content stored on the records of the blockchain—and the activities performed by the various participants—can be controlled depending on how the blockchain is configured. Generally, blockchains are designed for specific purposes, with users receiving multiple types of access or tasks.

One important category to think about blockchain technology is its access level and who runs the infrastructure to support it. In Figure 1, we depict the different types of blockchain according to the type of access they allow.

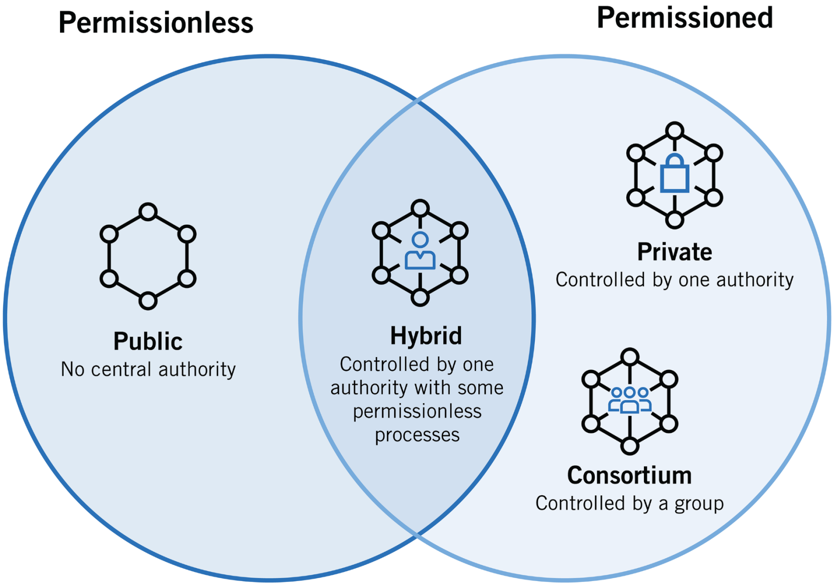

Figure 1: High-level overview of different blockchain technology setups

Both public and private blockchains use consensus algorithms to validate transactions, and both store them on a distributed ledger that every participant has a synchronized copy. The difference is that you need special permission to interact with a private blockchain, while anyone can freely enter a public network and see the history of transactions. There is a third type of blockchain known as permissioned or consortium blockchain. Permissioned infrastructure can come in many different permutations as it is a hybrid between private and public blockchains. As the name suggests, users require permission to use the network or participate in the consensus process. But private and permissioned infrastructure aren't the same. While private blockchains operate in an isolated network, this isn't necessarily the case for permissioned blockchain. A permissioned blockchain can also be a public network that only allows participation based on different access levels.

Each category we have been thinking about comes with compromises and will be more adequate to different use cases. In the following table we group a set of advantages and disadvantages of each category with respect to the MiCoLEC project's context.

Category	Advantages	Disadvantages
Public	**Transparency** All transactions are visible on a public network, meaning that anyone (even outside of the network) can view the entire record of transactions. Each network participant gets a copy of the distributed ledger containing all previous transactions, which is updated as transactions are executed on the network. **Censorship resistance** Public blockchains are censorship-resistant, meaning that no central party or authority can shut the network down or alter a transaction on the ledger. **High accessibility** Since there's no permission required to participate, public blockchains are some of the most accessible networks, even more so than traditional banking services. All you need is a smartphone or laptop with internet access.	**Energy inefficiency** One of the biggest downsides of PoW-powered public blockchains is their high energy consumption, which critics say is environmentally unsustainable. Newer blockchain networks are building on a proof-of-stake (PoS) consensus mechanism, which is more energy-efficient than proof of work. **Transaction traceability** While the identity of a public blockchain's participants is anonymous, transactions are technically traceable. If for instance, the wallet address of a network participant gets linked to the user, others will be able to trace the amount of cryptocurrency and past transactions of the participant, since the distributed ledger is publicly available.

Private	**Increased security** All private network participants require an invite by a central entity, which reduces the number of people with potentially malicious intent on the network. Combined with the fact that the main ledger is in a protected state, private networks are usually more secure. **Higher throughput** Private networks have limited access, hence they are usually much smaller than public blockchains. This leads to higher throughput and faster transactions due to the type of consensus algorithms required in larger networks, which tend to impair scalability. **Increased trust** As opposed to public blockchains, users on private networks aren’t anonymous, which increases the level of trust in these limited access blockchains. Every network participant can be identified.	**Lack of decentralization** One of the main disadvantages of private networks is that they aren’t decentralized. The shared ledger keeping track of transactions operates as a closed, central database, run by a single entity or organization. **Lack of immutability** Due to the inherent centralization of private networks, on-chain data and transactions can be altered by the network operator.
Permissioned	**Better performance** Since permissioned blockchains are not open to the public, they are usually much “lighter” than public blockchains — which means that there is much less on-chain data clogging the network. And with less on-chain data there’s less strain on the network, which leads to faster transactions and improved overall performance. **Varying levels of decentralization** The network operator(s) of permissioned blockchains can choose the desired level of decentralization. They can be partly decentralized or fully centralized as well. **Peak customizability** Out of the three blockchain categories, permissioned blockchains provide the most customizable infrastructure. The permission management feature enables the network operator to invite and give different roles to the participants. **Governance** Since they are operated by a central entity, permissioned blockchains usually don’t require community approval for hard forks.	**External data storage** Permissioned blockchains often require external storage space, but the decentralized storage methods used by public networks can’t be employed by some permissioned chains, depending on their degree of decentralization. This can put the integrity of on-chain data at risk. **Inconsistent level of security** The security of permissioned blockchains relies entirely on the chosen consensus algorithm and participants, which in case of bad actors, can compromise the entire network. Combined with the fact that these networks also require some type of central regulation, the potential for manipulation increases, in comparison to public infrastructure.

	Meaning that updates can be implemented quickly and easily, according to the needs of the respective entity.	

Cost analysis

The cost of implementing a blockchain solution depends on various factors such as, features, complexity and type of blockchain. Using a public blockchain such as Ethereum introduces transaction fees known as *gas*, while private blockchains present its cost in the form of infrastructure necessary to operate the cluster.

On Ethereum each transaction consumes a given *gas* amount, a value that oscillates according to the transaction backlog, transaction data and required computing power. At the time of writing the fees for a simple ERC20 Token transfer is ~5€. As previously stated this amount may raise or become lower depending on the type of transaction.

Private blockchains require an infrastructure properly dimensioned to the expected performance. Maintaining an infrastructure implemented leveraging cloud services or hosted on-premises introduces monthly and upfront costs that vary greatly dependending on the cluster capacity. Running a blockchain cluster requires multiple nodes, RPC, block explorer, monitoring (store and plot) and storage, which can easily reach 1500€ to 2000€ for a modest setup.

The main advantage of private infrastructures on-premises regarding cost of ownership is predictability. While the on-demand nature of cloud services makes it more elastic, it also introduces price variability. Nevertheless, a public blockchain that requires fees to process transactions can present high variability during periods of high network usage, moreover, as fees are incentives for processing a transaction, transactions with higher incentive are processed first.

The cost analysis is a sensible subject for the project since it can become very variable depending on the chosen path. Blockchain infrastructures are required to be highly distributed and scalable given the high computational requirements, moreover, for a project of this nature is expected a high number of transactions per operation or interaction, not only that, but also, high throughput. Another requirement to take into account is storage, since it's something that will increase over time and thus, hard to be dimensioned without initial usage samples, for that reason, a resilient distributed file system able to grow overtime is recommended.

Given the requirements for the project and the previous cost overview it becomes clear that paying fees for each transaction or operation is not feasible long term due to fee volatility and transaction requirements. Also, requiring every participant to pay fees for each interaction may be an adoption deterrent.

The cloud or on-demand/managed route offers the least resistance to entry along with more predictable costs and without the upfront cost of an on-premises infrastructure. Having a permissioned/private blockchain removes the need for fees and opens the path for more creative practices regarding token incentives and penalties. We believe that using cloud or on-premises are the most attractive solutions regarding cost predictability, although there certainly is some up-front cost, it lowers the barrier to entry and the infrastructure investment is gradually reduced over time, with the benefit of not being affected by a volatile fee structure. Also, there is more control and ownership over the infrastructure with much lower vendor lock-in.

State of the art: commercial solutions

From the different systems available we have identified a group of commercial solutions that could be used as the blockchain component for MiCoLEC. In this section, we briefly describe each one in order to document the following proposal of choice of technology. The information listed here is mostly retrieved from each technology's official documentation. All the sources can be found in Section Bibliography.

HYPERLEDGER

Hyperledger is an open source project created to support the development of blockchain-based distributed ledgers. Hyperledger consists of a collaborative effort to create the needed frameworks, standards, tools and libraries to build blockchains and related applications. Hyperledger Fabric was founded by the Linux Foundation, which is Hyperledger's framework with the most use cases and support. This component implements complex permissioning by, besides having validators, allowing each user to have a defined role, restricting the actions they can perform on the blockchain. All participants have known identity which is validated against the organizations‘ identity management system. There are no anonymous or pseudonymous users. There is no PoW algorithm and crypto mining in Fabric, which allows for high scalability and fast transactions. One of the most interesting characteristics of Hyperledger is its modular architecture that allows for the development of custom plug-in components. Hyperledger is supported by one of the richest development communities in the space.

QUORUM

Quorum is an open source blockchain protocol specially designed for use in a private blockchain network, where there is only a single member owning all the nodes, or, a consortium blockchain network, where multiple members each own a portion of the network. Quorum Smart contracts are written in the Solidity language and the Raft protocol is used as the consensus mechanism in Quorum. This choice is targeted at higher transaction throughput rates when compared with Proof-of-work approaches. However, Quorum's channel-based approach to privacy presents challenges for privacy and scalability as use cases become more complex. In addition, Quorum does not require a built-in cryptocurrency because consensus is not reached via mining. As a consequence, it is not possible to develop a native currency or a digital token with Quorum, which for our own scenario is a deal breaker.

MULTICHAIN

Multichain was developed by Coin Sciences and is a fork of the Bitcoin blockchain. However, unlike Bitcoin, MultiChain allows users to configure several parameters such as the permissions to access the network, the privacy of the chain, the maximum block size, and the mining incentive. MultiChain supports a variety of programming languages such as Python, C#, PHP, Ruby or JavaScript. It Focuses on the two strong use cases: the asset ownership lifecycle (issuance, payment, exchange, escrow, retirement) and General immutable data storage. It is very easy to install, configure and create a network MultiChain since you don't need to write any code and can get started immediately through easy-to-use APIs. Multichain simplicity may attract developers and organizations but the inability to write custom smart-contracts prevents us from considering it as an option.

CORDA

Corda is written in the Kotlin programming language and supports development both in Kotlin and Java. Corda exists in two main editions. There is an open-source edition that is free for personal and commercial use called Corda, and the enhanced paid edition called Corda Enterprise. Corda Enterprise offers additional performance enhancements, such as higher computational capacity for large-volume transactions. While Corda is used in a variety of industries, the majority of its customers come from the finance, banking, insurance, and capital markets sectors. Corda’s private blockchain features are particularly relevant for the companies in these sectors, as data confidentiality is highly important for their operations. The biggest advantage of Corda for businesses is the ability to protect the privacy of transactions. Corda’s most common use cases include inter-organizational cooperation. By creating a blockchain-based network on Corda, businesses can significantly improve cooperation efficiency and cut down on the cost of interacting. For example, a Corda network of insurance companies, brokers, and re-insurers can streamline claims processing, data verification, mutual payments, and other business processes. Corda brings the benefits of blockchain to finance-related industries while ensuring that confidentiality and privacy, so much heralded by companies in these industries, are not compromised. Considering our use case is in the logistics domain, it is hard to assess the feasibility of using Corda since there is not much previous experience and use cases to draw from. Moreover, the opaque pricing from R3 (the company behind Corda) makes it really hard to take into consideration at such an early design stage due to the increased difficulty in predicting what the system may demand from the blockchain component.

ETHEREUM

Ethereum was created to address some of the shortfalls of Bitcoin. While Bitcoin is great for storing wealth (BTC is the most secure cryptocurrency in the world) it lacks complex functionality. You can send and receive transactions and execute some other essential functions, but smart contracts are not supported. That’s where Ethereum comes in. Ethereum offers a high level of customization so that developers can create custom products. Ethereum has been developed as a permissionless, public blockchain, in which every smart contract can be programmed in connection with decentralized applications (dApps). For this, a virtual machine (VM) is provided on the blockchain, for which a fee must be paid depending on the effort required to execute the programming code. The most used programming language for Ethereum is Solidity. Ethereum is the second most decentralized cryptocurrency in the world, after Bitcoin. It has the largest developer community in the world, even larger than Bitcoin’s. This gives Ethereum a tremendous advantage over other protocols. When you build an app on Ethereum you can instantly connect it to hundreds of other protocols that already exist. In the Ethereum community, this is known as money legos.

As great as Ethereum is, the platform certainly is not perfect. As we can see with Bitcoin and Ethereum, decentralized protocols tend to be slow. Bitcoin has average speeds of 7 TPS (Transactions Per Second), while Ethereum has a speed of 15 TPS. That's double Bitcoin's speed, but it's not nearly enough. The Ethereum coin that powers the network is officially named Ether and popularly known by its ticker symbol — ETH. The gas fee on the Ethereum network remains one of the biggest challenges, and it is also affecting the Ethereum blockchain scaling. In fact, when the gas fee hits high, the Ethereum network has been redirecting users to other platforms, which shows how serious of a problem those fees can be.

HEDERA

Hashgraph is a distributed ledger technology that uses a specific transaction handling and voting protocol to make it faster and more energy-efficient when compared with Bitcoin or Ethereum. That approach is called Hashgraph and instead of grouping data into blocks, a consensus protocol works for each transaction determining if such particular transaction is added to the ledger or not. This approach speeds up transaction times, making a Hashgraph network capable of handling up to 250,000 transactions per second. This speed is currently throttled to 10,000 TPS on the Hedera Hashgraph, but it can get lifted if the need arises. Another benefit of its consensus protocol is that transactions get confirmed in about 3 to 5 seconds. This puts Hedera way beyond the 10 to 60-minute blockchain confirmation time-frames and sets it on par with credit card companies. Hashgraph uses what is called asynchronous Byzantine fault tolerance to maintain a secure network. Byzantine fault tolerance takes the potential unreliability of the network's nodes into consideration when reaching a consensus, to avoid a damaging system collapse. The system is also protected against DDoS and Sybil attacks. Hashgraph uses just about 0.0002 KWh per transaction, making it extremely more energy-efficient and environmentally friendly than most blockchains. Hedera Hashgraph's transaction fees are also very low and start from $0.0001, depending on exactly what you need to get done on the system. The costs are significantly lower than the $15+ that popular blockchains charge per transaction.

Proposed solutions and brief experimental evaluation

From the analysis and research leading to this document, we concluded that the two most suitable systems to be used for the MiCoLEC blockchain component are Hyperledger and Hedera. Both support customized smart-contracts and a controlled environment. This will be very important to lower the entry barrier for newcomers to the MiCoLEC platform. Only a subset of participants will be required to provide infrastructure to support the platform and newcomers may try the system without incurring in high upfront costs. Later on, each participant will be able to increase their stake in the system by providing infrastructure and contributing to the decentralization and robustness of the solution. The rejection of public blockchain technology is grounded on the assumption that a system such as MiCoLEC will scale in the number of transactions (packages) running in the system. As a consequence, any solution that has costs per transaction will not be viable in this scenario. Another important input used for this proposal is the fact that both Hyperledger and Hedera are supported by strong commercial companies and consortiums, which typically assures continuous development and bug fixing as well as good developer support. These are critical for a project such as MiCoLEC that is expected to run for 2 years and that has still a number of open challenges to address. Having a customizable, extendable, and mature system as the foundation for such endeavor is essential to avoid discovering technological limitations too late in the development efforts.

The decision on which technology to adopt for the implementation of the MiCoLEC platform naturally depends the most on the scope of the project and its goals. However, it is also important to take into account additional characteristics of the technology such as its maturity and the availability of a support community or services. In order to assess this, we have done a set of brief hands-on tests with the two most promising systems: Hedera and Hyperledger.

Since the goal was to quickly assess potential major differences in maturity of both systems, we devised a simple test where we wanted to create a custom token from scratch and get a sense of the difficulty of such a task.

We have locally deployed both systems and used their testnet setup for the experiments. Both systems were fairly easy to set up and both had good documentation. We then proceeded to create a custom token on both platforms. The main difference between Hyperledger and Hedera were the SDKs provided by the latter, which provided a slightly smoother experience. However, Hyperledger seems to better support customisation, which for the MiCoLEC project is paramount. Additionally, both these technologies are supported by an extensive community of developers and are accompanied by sound documentation.

Although these were very limited tests, they gave us confidence that choosing any of these platforms will allow us to design and develop MiCoLEC without major roadblocks. The edge pends towards Hyperledger due to the broad scope of use cases already running there and for the highly customized design.

Conclusion

Following the different data points and analysis presented along this document, our decision is to move forward with Hyperledger as the blockchain/distributed ledger component for MiCoLEC. The goal is to provide an extendable and customizable platform that will allow the implementation of MiCoLEC but also leave room for improvements and platform evolution in the future. Additionally, the maturity of Hyperledger, which is in fact one of the most mature technologies in the space, provides the necessary support for innovation.

In a different deliverable, we focus on the architecture and design of the MiCoLEC platform and, to the best of our knowledge, Hyperledger will be able to accommodate all the requirements and goals for that design.

Bibliography

Nnamdi Okeke. (2021, July 12). Hashgraph: Meaning, Advantages, Disadvantages. TargetTrend. https://www.exodus.com/news/ethereum-review/

Sam Klemens. (2020, October 16). Ethereum Review: Ethereum Use Cases, Advantages & Disadvantages. Exodus. https://www.exodus.com/news/ethereum-review/

Aran Davies. (2018, August). Pros and Cons of Hyperledger Fabric for Blockchain Networks. DevTeam.Space. https://www.devteam.space/blog/pros-and-cons-of-hyperledger-fabric-for-blockchain-networks/

Shobhit Seth. (2022, July 28). Public, Private, Permissioned Blockchains Compared. Investopedia. https://www.investopedia.com/news/public-private-permissioned-blockchains-compared/

Gideon Greenspan. (*2018, May 18).* R3 Corda: Deep dive and technical review. Multichain. https://www.multichain.com/blog/2018/05/r3-corda-deep-dive-and-technical-review/

Coin Sciences. (2015, Jul. 21). Advantages and Disadvantages of Multichain. Multichain. https://www.multichain.com/qa/6277/advantages-and-disadvantages-of-multichain

ImmuneBytes. (2021, August 27). An Introduction to hedera Hashgraph. Immunebytes. https://www.immunebytes.com/blog/an-introduction-to-hedera-hashgraph/

Demetrios Zamboglou. (2019, Mar 18). Hedera Hashgraph Explained. Medium. https://medium.datadriveninvestor.com/hedera-hashgraph-explained-c5d8ce4730a6