UniCon: A Scalable and Universal Architecture for Content Management

Pablo Rodrigo[1] , **Martijn de Vos**[1] , **Johan Pouwelse**[1]
[1]Delft University of Technology

Abstract

Creating content in the digital era has never been so easy. However, there exists no framework in which artists can share, cooperate and transfer universal content. We refer to the term *universal* as that allowing our architecture to manage content both from the physical and digital world. Current platforms are either centralized and sector-specific such as Netflix and Spotify, or decentralized. On the one hand, the centralized platforms have control over both the data and the identity of their users. On the other hand, decentralized platforms lack scalability and an identity layer. Our architecture expands current decentralized solutions with a more scalable and generic framework. This solution allows users, and their verified identities, to possess control over their data; share and exchange content, both from the physical and digital world; and get attribution and royalties for their work in any digital coin.

1 Introduction

Content can be described as: ”ideas contained in something written, said, created, or represented” [1]. The digital world has exploded the amount of content we create. In 2018, 2.5 quintillion bytes were created every day [2]. Many platforms allow sharing content. For instance, every day 86 Million pictures and 60.000 tracks are uploaded to Instagram and Spotify respectively. Both of these platforms own the data of their users. In the current digital paradigm, it has become ordinary for users not to own the data they create. Google, Facebook, Microsoft, and Amazon alone store at least 1,200 petabytes of information [3]. The increasing control and power from big corporations and governmental institutions together with the appearance of Blockchain technology has lead to a sharp increase in the number of *decentralized applications*.

Blockchains are decentralized digital ledgers of transactions enabling data to be exchanged within a network without the need for intermediaries [4]. One of the most popular blockchains is Ethereum. Its main characteristic is the possibility to build decentralized applications. In fact, there are currently more than 3.000 decentralized applications running on the Etereum network [5]. These applications span multiple fields, some of the most known applications of Blockchain technology include Cryptocurrencies, digital currencies removing the need for central banks, and the creation and trading of non-fungible tokens (NFTs). NFTs are unique digital certificates representing digital content showcased in marketplaces. As an example, the first tweet ever published on the platform of Twitter was sold for $2.9 million [6]. Even though the volume of NFTs marketplaces in the first quarter of 2021 was $2 billion [7], NFTs are not universal and cannot manage non-digital assets such as the traditional art market or the exchange of furniture. We believe that the possibilities of decentralizing content sharing go far beyond the exchange of NFTs and therefore, our work will aim to manage all types of content, both physical and digital, including those in the music industry, the art world, and the real estate market.

Nonetheless, although current NFTs architectures are a good starting point for our solution, they have some drawbacks. These architectures are based on Ethereum and suffer from scalability issues and high transaction fees [8], deriving in the management and exchange of content being prohibitively expensive for mass usage. Moreover, these architectures lack an identity layer reducing traceability and accountability [9], key aspects to bridge the gap towards mass adoption. In fact, the lack of an identity layer is a well-extended problem. Since the creation of the World Wide Web in 1990, user identities have been limited to the scope of specific websites. Recently, new approaches such a Single SingOn have reduced the total amount of different identities. Nevertheless, the ownership of data remains out of reach for the user.

Self-Sovereign Identity is an identity management system that allows individuals to fully own and manage their digital identity [10]. The appearance of Self-Sovereign Identities together with Blockchain aims to return control over our identity and our data. By decentralizing data management, users can choose what and when to disclose certain information.

In this work, we devise a fully decentralized system architecture for the management, transfer, and attribution of any type of content. Our programming interface enables any artist to quickly link their work to their verified identity and to share the content with others using the BitTorrent protocol. The enabling element of our system is a scalable and lightweight distributed ledger called TrustChain [11]. Our architecture will be integrated together with the current work done by the Delft Blockchain Lab in the SuperApp, application leveraging TrustChain and IPv8 to maintain accountability of transactions across the network for distinct applications. The main research question this paper will try to answer is:

How can we design a universal and scalable content-sharing architecture with verified Self-Sovereign Identities and generic coin transfer?

In this work, we do not only identify the need for a new architecture and design an architecture satisfying those needs; but we also build and evaluate a prototype of the skeleton of our architecture to prove the added capabilities of our system, and we expose the future work needed to build upon the current prototype towards a system ready for mass adoption.

2 The shortcomings of NFTs

The ever-growing amount of content would benefit from a scalable ecosystem where artists, and their respective digital identities, enjoy the freedom of sharing their work while maintaining data ownership and ensuring entitled attribution and royalty payments. However, the majority of the existing content-sharing platforms are sector-specific and the control over user data is managed by big corporations. Instead, decentralized solutions propose content-sharing architectures which allow users to maintain ownership of their data. However, current decentralized architectures solely focus on the transfer of NFTs and are based on the Ethereum blockchain. These architectures have some drawbacks described in this section.

Firstly, current architectures are not *universal*. These manage only digital content, and in most cases, only digital artwork. In this work, we devise a wider scope for the application of decentralizing content-sharing. In fact, we do not restrict ourselves to the digital world but aim to include traditional trade markets from art to furniture and real estate.

Secondly, the Ethereum network has a fraud prevention rather than a fraud detection approach which limits the *scalability* of the network. The fraud prevention approach is achieved by enforcing global consensus across all nodes. The scalability limitations of the current architectures pose obstacles towards the mass adoption of current marketplaces. In numbers, the Ethereum network is only able to process around 20 transactions per second. These numbers are dangerously small compared to electronic payment giants like Paypal and Visa, which are able to verify around 193 and 1670 transactions per second respectively [12]. Although Ethereum 2.0 promises to drastically improve the scalability of its network, its deployment is still to be seen and is not expected to be completed in the near future. It is therefore clear that there exists a challenge to increase the number of transactions per second of the Ethereum-based architectures. The use of the Ethereum network as the blockchain to store the transactions of the management of content also leads to high transaction fees, adding hurdles in the transformation of current NFT marketplaces towards frameworks to share and transfer all types of collaborative content.

Thirdly, the use of Ethereum blockchain in the architectures of NFTs marketplaces creates a *dependency* with the Ether cryptocurrency for the management of such content. This dependency leads to the value of the showcased content being highly volatile due to the inherited volatility of the Ether coin. Furthermore, in the current fragmented digital currency paradigm, the requirement to manage content with a specific coin, such as Ether, creates a higher barrier entry for those users not using this coin. Therefore, we believe it is of the utmost importance to offer the possibility to remunerate artists with any digital currency.

Lastly, a key threat of current architectures is that a user could pretend to be an author and sell stolen content. Any piece of content can be currently sold without verifying the identity of the seller nor buyer. It is estimated that in NFT marketplaces, more than $150.000 worth of content has been published under a false identity [13]. This problem originates from the *lack of verified identities*, which provides a getaway for illegal activities, creating a major flaw for an architecture aiming to become the framework of universal content sharing. Along with verifying the identity, we foresee the importance of maintaining the ownership attribution and the payment of royalties to the original content creators when its work is used as the base for other content. This enhances content sharing and the development of collaborative projects. For instance, in the case of an author releasing a song, a remix can be done using the original song as the main raw material and it is therefore clear that the original author should be rewarded for any income the remix author obtains.

In conclusion, the four main problems our architecture will try to tackle are the management of universal content, the lack of scalability and high transaction fees of the Ethereum network, the dependency on the volatile coin Ether, and the lack of verified identities.

3 Solution outline

The four main shortcomings of current architectures lead to each one of the four main pillars of our solution. These four main pillars are universal content, accountability-based blockchain, generic digital coin, and digital identities. The cooperation of the four pillars creates the core of the outlined solution. The main contribution of this solution is to provide a framework, never seen before, enhancing the sharing, distribution, and cooperation of all types of content by artists in distinct fields while maintaining attribution and royalties payment in any digital coin.

These four main pillars do not work in isolation but highly cooperate together in order to achieve the goal of creating a universal, scalable architecture with verified identities and generic coin transfer. The interaction amongst them is briefly depicted in Figure 1. The Universal content (pillar I) is the core of the architecture. This content is stored in a decentralized manner with the accountability-based blockchain (pillar II), transferred by any digital coin (pillar III), and owned by digital identities (pillar IV). Moreover, the accountability-based blockchain (pillar II) records and

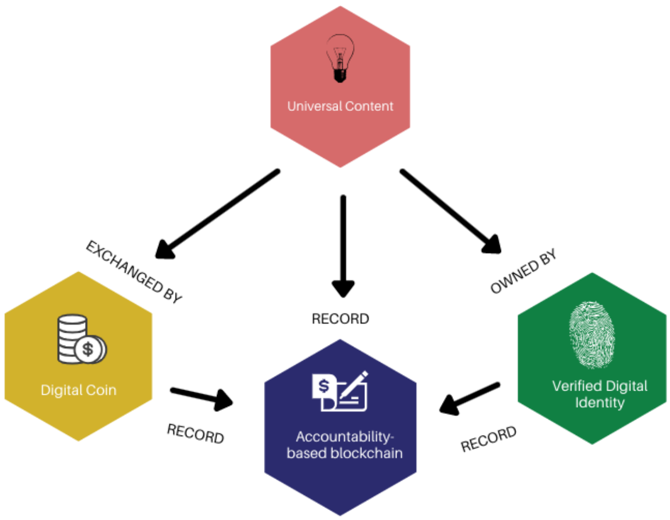

Figure 1: Main pillars of the architecture

maintains accountability of the transactions concerning the exchange of content by digital coins (pillar III). Additionally, all stored transactions are signed by verified identities.

I. Universal content

The first pillar, Universal Content, is the raw material of the solution. Current architectures are based solely on the management of digital content, do not enhance the creation of collaborative content, and do not allow to get attribution in the shape of royalties for incremental content such as a song on which a remix is based. Our work ambitiously expands the scope from digital content, the only type of content showcased in current NFTs marketplaces, to all types of content. Other types of non-digital content include the art market, estimated to be valued over $64 billion, consumer goods such as the furniture market, whose market is estimated to be worth over $564.7 billion, and the real state market, valued around $2687.35 billion.

In order to include these types of non-digital assets in our architecture, new identifiers need to be used to ensure that each item can be uniquely identified and its ownership chain can be traced. In the case of art pieces, an ownership certificate awarded by a trusted institution such as UNESCO could be used to uniquely identify the item. This certificate can then be stored and traded in our architecture. Similarly, serial numbers and property certificates could be used to uniquely identify furniture items and real state properties respectively. For digital items, Torrent files are created using the BitTorrent protocol and are stored and traded in a similar fashion to non-digital items.

In essence, our architecture will cooperate with organizations to allow for seamless integration of the digital representations of physical items in the platform. The benefits for such markets to be integrated into our architecture are increased traceability, impermeability to illegal activities by adding a layer of verified identities, and standardization of transactions. This last argument refers to the fact that different countries may have different transaction protocols which can derive in the interaction amongst different countries, and therefore between different transaction protocols, to become problematic. Instead, our architecture proposes a general solution enabling any type of content to be transferred using the same procedure enabling the user to have an intuitive user experience.

II. Digital coin

The third pillar of the architecture is the digital coin. In the current fragmented composition of the field of digital currencies, with over 4,000 different cryptocurrencies [14], we believe the use of a specific coin should not be a constraint making the entry barriers to our system any higher. Therefore, the architecture has been designed to allow for the management of content with any digital coin. This includes the usage of stablecoins such as the DAI, Tether, or the Eurotoken, recently introduced by Delft Blockchain Lab [15]. Stablecoins are digital currencies that attempt to peg their market value to some external reference, such as the U.S. dollar, or the Euro. These external references have significantly less volatility than current non-stable digital coins. The volatility of the market is a key factor hindering the mass adoption of decentralized solutions using non-stable coins [16]. Market players within distinct fields such as the real state market are used to small value fluctuations and the sharp increase of volatility can negatively impact the adoption of the architecture by such market players.

The main action of this pillar is to provide the possibility for extension with any digital coin. The implementation should be able to manage the balance and exchange the content for the stipulated value in the stipulated coin. Moreover, this implementation should include a protocol for monetary exchange. In the current version of the prototype developed, the monetary exchange has been left as future work.

III. Accountability-based blockchain

The second pillar of the architecture is the accountability-based blockchain. The main task of this pillar of the architecture is to store and maintain accountability of every transaction, both creation, and transfer of content. Each user has a grow-only personal ledger in which the transactions are stored [17].

The accountability-based blockchain uses a fraud detection approach rather than the fraud prevention approach used in current architectures. Detecting fraud, such as forking of person ledgers, is obtained by: "having users continuously requesting random transactions from other users and sharing their transactions with other users upon request. The consistency of incoming transactions is checked against known transactions and illegitimate modifications of the personal ledger can quickly be revealed by the collective effort of users in the network" [17]. By using the fraud detection approach, we remove the need for global consensus and mimic the dynamics of the real-world trade in which not everyone reaches an agreement on the occurrence and validity of transactions. Instead, we record the transactions and in the case of an illegal activity being detected, the transaction can be traced back and the responsible can be held accountable. This approach trades the ability to prevent fraud for added scalability.

There are five types of blocks that are recorded in the accountability-based blockchain. These can be observed in Figure 1. The user is able to create content, offer it for sale (in a future developed marketplace), transfer the content, prove its verified identity and list its wallet for content exchanges, such as an Ethereum wallet. These blocks create the dynamics of our system and are stored in the personal ledgers. These ledgers are then crawled to build the user experience of the application. For instance, when the user wants to see his content, the personal ledger is crawled and all created content with the public key of the user which has not been transferred is listed. Similarly, in order to display the available items of the marketplace, the ledger is crawled for elements created which are offered to sale but are not transferred. These can be items of the user itself (available to change settings such as price) or of other peers which the user can purchase.

Table 1: Types of blocks in the accountability-based blockchain

Block type	**Description**
create_content	This block shows the creation of content and contains the torrent Hash.
sellable_content	This block indicates a item is for sale and links to the its *create_content* block.
transfer_content	This block is an offer from a buyer to an owner for an item which is open for sale (links to *sellable_content* block)
indentity	This block is the verified identity of the user, contains the claim indicating its identity has been verified.
wallet	This block stores wallet addresses

This pillar interacts with all other pillars in the following way: it stores the content of pillar I, Universal Content, both from the physical and digital world; it stores the transactions that users have sign through verified public keys, linked to identities by a third-party (pillar IV); and stores the transactions containing monetary exchange (pillar III), a key aspect in the fraud detection approach.

IV. Digital identities

The problem this pillar aims to tackle is the possibility for users to publish stolen content, use the platform for money laundering, and other illegal activities that can be hidden under the anonymity of the network. Earlier in *pillar I*, it was explained how fraud detection is used instead of fraud prevention. The goal of this approach differs from fraud prevention in a shift towards maintaining accountability in order to be able to trace back the transaction chain in case of illegal conduct, rather than preventing this conduct in the first place. This pillar builds upon this concept and enables this traceability to be done one step further, up to a real identity. The result of this addition is that any misuse of the architecture will be able to be traced back not only to a transaction but to an identity that can therefore be held accountable. Furthermore, the digital identity pillar of this architecture is based on the concept of Self-Sovereign identities. As defined by A. Mühle et al. in their survey on essential components of a self-sovereign identity [10], Self-Sovereign Identity is an identity management system that allows individuals to fully own and manage their digital identity. In addition, it states that users exist independently from services. This concept of identity ownership diverges from current website-centered identities in which big corporations own the data of their users. Furthermore, the attestation of the digital identities is done by a trusted third party who verifies the public key of the user and links it to a real identity. This allows maintaining anonymity for the user within the network while maintaining traceability and accountability.

4 System Architecture and Implementation

All the above-mentioned elements compose the core elements of the proposed architecture. We have built a prototype to prove the added capabilities outlined in the previous section. The system architecture of our system is shown in Figure 2, and it provides one more level of granularity depicting how the different individual components of the proposed solution work together. Some of the components of the system are not included in the prototype and will be briefly discussed, and for those elements included in the prototype, relevant implementation details will be given. The architecture is mainly structured in three big layers: the client wallet, the content management layer, and the distributed storage. The client wallet is the layer interacting with the user and providing the user with the experience of creating and exchanging content, the content management layer is the behind-the-curtains transaction logic which allows the user to experience the mentioned capabilities of the system, and the distributed storage is the foundation of the system, which allows to distribute and maintain accountability of the information stored in blocks.

I. Client wallet

The first main layer of the proposed architecture is that concerning the user, the client wallet. This wallet contains three elements: digital coins, collaborative content, and verified digital identity.

As previously explained, the architecture allows using any digital coin, including stablecoins, reducing the dependencies on single volatile coins current architectures suffer from. To provide this generality, we have decoupled the coin logic from the wallet and the transaction logic. All the logic related to the management of coins has been isolated in an abstract coinAPI class. This class contains the following methods: *addCoins(), subtractCoins(), checkBalance(), buyContent(), sellContent()*. Any implementation of the digital coin system will extend this class. Moreover, this implementation will

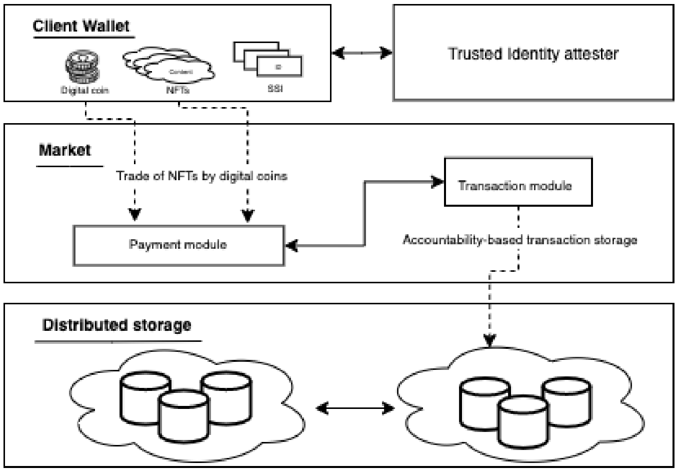

Figure 2: System architecture of UniCon

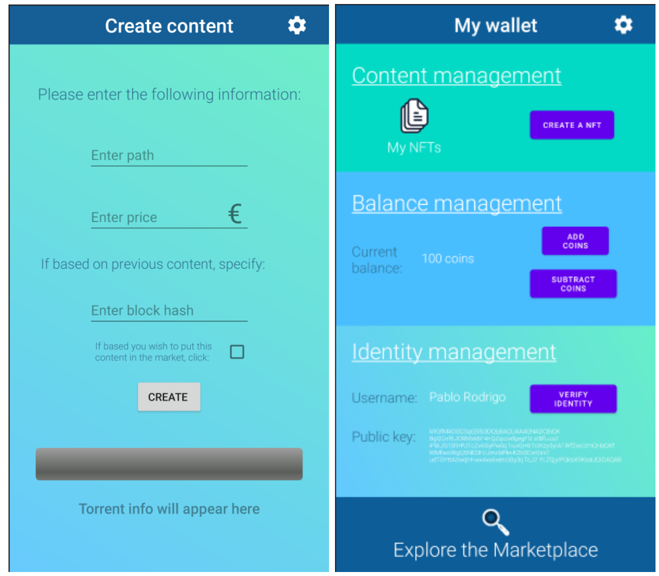

Figure 3: User Interfaces of the Client Wallet screen and the Creation of Content screen

contain the protocol which will be used to ensure the content is exchanged securely for the right amount.

Moreover, the universal content is the raw material of the architecture. To store this content in the decentralized storage, we need to do some priors steps depending on the type of content. In the previous section, it was stated that we store ownership certificates for art items, property certificates for real state items, serial numbers for furniture. We create torrent files from these certificates in the same way we can create a torrent file from any digital asset. The main takeaway is that once we can uniquely identify an item, this can be stored in our system. The Torrent files are created by using the BitTorrent protocol with the jlibtorrent library. These files are then stored in the decentralized storage.

Lastly, the verified identities are obtained by verifying the identity and linking such identity to a public key. An external trusted party, such as a governmental agency, is responsible to link the public key to the identity of a user. Once a public key is verified, the user can sign any block, and any transaction involving the user will be traceable not only throughout the transaction chain but up to a verified identity that can be held accountable in case of illegal conduct. This allows the architecture to ensure both accountability and traceability. Both of these are key elements of the architecture due to our design choice of fraud-detection, instead of fraud prevention, blockchain protocol. In essence, it ensures that fraud can be not only detected but also linked to an identity.

The wallet containing the collaborative content, the digital coins, and the verified identity is the main contact point of the user with the architecture. This creates the need for an intuitive user interface allowing the user to create content, manage its balance and verify its identity. The user interface is visualized in Figure 3.

II. Content management

The second layer of the proposed architecture concerns content management. In this layer, we discern two main modules: the payment and the transaction module. These two modules work together to provide the core functionality to the architecture, the creation and transfer of collaborative content.

The transaction module is responsible for the management of content and maintaining accountability of the executed transactions. The two main types of transactions are the creation of content and the transfer of content for generic coins. The creation of content allows for many shapes, it can be either individual or collaborative, and it can be either building upon an existing content or not. Once the user inputs the mentioned information, a proposal block is created. The proposal block is of type *"create_nft"*, and contains a transaction (a map), and *"ANY_PK"* as recipient. The latter allows to broadcast the block rather than send it to a specific recipient and the transaction consists of a map of the following keys with their corresponding value: the price of the content, the Torrent Hash, and the ownership chain. The ownership chain is a list of public keys of previous authors. In the case of new independent content, the author's public key will be the only element of the chain. In the case of the content being based upon existing content, the ownership chain of the previous item must be retrieved and the new author's public key will be prepended. To retrieve the user's public key, we do the following: "*TrustChainCommunity.myPeer.publicKey*". Lastly, in the case of collaborative content, a list of authors will be added to the ownership chain.

The transactions related to the transfer of content also play a key role in the proposed architecture since they are responsible for the transfer of the unique certificate to the right public key in exchange for the stipulated value in the stipulated coin. To execute this type of transaction, the transaction module cooperates with the payment module.

For an item to be offered for sale, the user needs to select this when creating the content (in a future version, it will be possible to add to a marketplace at any time), then a "*sellable_content*" block is created. In the first iteration of this architecture, it is assumed that the user knows the block hash of the block where the item intended to buy is stored. In the current version, after the user inputs the block hash, the block is retrieved with the following call: *TrustChainCommunity.database.getBlockWithHash(blockHash).* Once retrieved, the transaction is extracted, and the price of the item is obtained. At this stage, we observed the interaction of the transaction module with the payment module for the first time in this subroutine. The payment module peeks the balance of the buyer to ensure it has enough coins to execute the transaction. In case the balance is enough, a proposal block is created with the type "*transfer_content*"; the same transaction as the existing block storing the item to be exchange plus a link to the block containing the sellable item; and the public key of the owner of the block as the counterparty. If the block is signed, the payment module is required to execute the payment and pay the corresponding royalties to the given previous authors by the transaction module. In order to provide the mentioned capabilities, the payment module makes use of the CoinAPI explained in the previous subsection. The inner workings of the monetary exchange protocol are left as future work of the implementation of the abstract CoinAPI. The cooperation of the transaction and the payment module yield the necessary content management capabilities to the proposed architecture, making it possible to create and interchange collaborative content while maintaining attribution in a scalable and generic approach.

III. Decentralized storage

This architecture leverages the scalability of IPv8 and Trustchain to have decentralized storage with a shift of focus from safety to accountability.

TrustChain is a distributed ledger protocol used to maintain accountability of transactions [11]. TrustChain was developed in the Delft Blockchain lab and provides the medium for all the content storing and content transfer capabilities. TrustChain relies on guaranteed eventual consistency to solve the double-spending problem [17]. However, since time is not a critical factor in our solution, this problem does not pose a major drawback. Additionally, IPv8 is defined as a Peer-to-Peer protocol providing authenticated communication in which peers in the network are identified by public keys, and physical IP addresses are abstracted away.

The use of TrustChain on top of IPv8 allows us to obtain our accountability-based blockchain with added scalability. In fact, in other literature, it has been shown that TrustChain has no problem scaling up to 10.000 users and beyond [18]. Furthermore, in this evaluation 1000 transactions per second are executed successfully. Comparing the throughput of TrustChain to the 30 transactions per second of the current version of the Ethereum network, we can observe an improvement of over 30x.

As stated, the time of an individual transaction is not a critical factor. Therefore, we conclude our architecture aims for high throughput and not necessarily low latency. These characteristics fit well the capabilities of TrustChain which with higher latency, reduces the possibility of misuse by reaching eventual consistency, but higher throughput compared to the Ethereum network.

Integration with the Super-App

The TrustChain SuperApp introduced earlier, is a mobile application under development by the Delft Blockchain Lab. Many applications have been integrated within this SuperApp. These applications leverage the capabilities of IPv8 and the TrustChain protocol to decentralize processes in many different domains. Some examples are the Euro-token (decentralize digital currency transfer), the Peerchat (decentralize messaging app), the MusicDAO (decentralize music platform) ...

In order to integrate the architecture proposed with the TrustChain SuperApp, the following steps need to be followed. In the first place, the architecture must be built as an Android module with the project *"common"* as a dependency. Furthermore, the module should be added as a dependency on the build.gradle file of the SuperApp, the activities of our module should be added as so, in the AndroidManifest.xml of the SuperApp and in the AppDefinition.kt file of the SuperApp our new application should be added with the logo, color, and starting class name of our application. These steps will include our architecture as a sub-module of the application. Furthermore, the SuperApp starts an instance of the IPv8 and TrustChain community which can be retrieved to use in each of the applications integrated into the SuperApp. A community (or an overlay) represents a service in the IPv8 network. Every peer can choose which communities to join when starting the protocol stack. In our architecture, we will stick to the TrustChain community.

5 Experiments and results

In this section, we have evaluated the performance of our architecture concerning an individual peer of the network. This evaluation serves as a starting point for the complete assessment needed to ensure the architecture is ready for mass adoption. The setup of the experiment was the following: one peer and 10.000 *create_content* blocks created sequentially with no multi-threading. The metrics to monitor were total time and network requirements.

The results yielded a total time of 171824 ms, leading to almost 60 blocks per second, and a network usage oscillating with maximum peaks of 131 KB/s of send data. Since the evaluation focused on the capabilities of the network to handle one peer creating content, the relevant data concerns the throughput sent by the user to the network. The behavior of the network usage can be visualized in Figure 4.

From this evaluation, we have concluded that the user will not encounter a bottleneck in the creation of multiple pieces of content sequentially. This is mainly due to the lightweight ledger used in our architecture TrustChain. Nevertheless, as stated in Section 8, the performance evaluation will be extended in future iterations with a focus on the global performance of multiple peers managing content simultaneously.

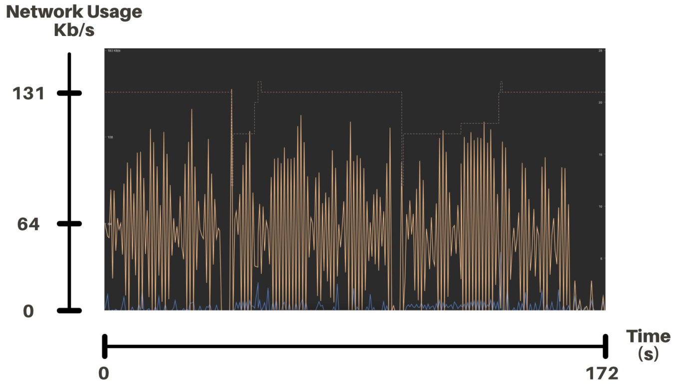

Figure 4: Network usage of the conducted experiment

Therefore, at this moment, we can not make any claims about the global throughput of the architecture until a multi-user assessment is carried out.

6 Related work

The decentralization of data management is not new. Already in 1999, Napster, a peer-to-peer music download system, proposed a semi-decentralized approach to manage musical digital content. In the analysis of the rise and fall of Napster executed by Bengt Carlsson et al. [19], it is explained how the lack of protected ownership (copyright) and royalty system ended up in a court order enforcing these elements to be introduced in their architecture, which lead to the downfall of Napster. In our system, verified ownership and royalty payments are key elements that have been taken into account from the design phase. This ensured the inclusion of all players within the market, including both governmental organizations and artists.

Other existing literature from Zehui Xiong et al. [20], focuses on the creation of an architecture with Blockchain for the Data Management of Internet-of-Things. Although this work is not focused on the exchange of goods but on the data management of IoT devices, some similarities can be drawn. The architecture proposed in this literature makes use of an accountability-based blockchain, the same type of blockchain used in our architecture (second pillar). This shows the importance of accountability in the use of blockchain for data management.

In the work of Zhaofeng Ma et al. [21], a system for Digital Rights Management is described. In this work, the focus is done on the management and it is expressed that in the future the possibility trade this content with the Ethereum network will be added. Our system provides two added advantages compared to this work: its universality, the content is not restricted to the digital world; and its lack of dependency on the Ethereum network for the exchange of this content.

Lastly, in the analysis of current NFT architectures performed by Qin Wang et al. [22], some of the problems we elicited in Section 2 are also listed. This paper explains that the lack of scalability of NFTs Ethereum-based marketplaces can only be solved by a redesign of the blockchain topology and existing blockchain systems cannot fulfill such requirements. Therefore, we can conclude that the decision of substituting the Ethereum network for our accountability-based blockchain, Trustchain, is the correct approach. The second main issue mentioned with regards to the high transaction prices due to the use of the Ethereum network is also listed in the aforementioned literature. They estimate the transaction fees for an exchange of an NFT is around between USD 60 and USD 100. Furthermore, they add those expensive fees caused by complex operations and high congestion greatly limit the mass adoption of these architectures. This last argument reinforces our thesis of the need for a new scalable, universal architecture for the exchange of all types of content.

7 Responsible Research

The ethical aspects are inherent to any research project. In this paper, the main relevant ethical aspects reproducibility, and privacy.

Reproducibility is a key component in academics. It aims to ensure the academic community can reproduce the scientific contribution presented, in the same circumstances, and verify the claimed conclusions. Several steps have been taken towards improving reproducibility. In order to use the architecture presented in Section 3, the code has been released in [23]. A readMe has been included to help reproduce the settings of the environment in which the architecture was developed. Throughout the paper, we have claimed that the added value of the presented solution is: its increased scalability with regards to current architectures, the existence of an identity layer, and the possibility to exchange content with generic coin transfer. The increase in scalability is, as explained in Section 4, inherited from the use of TrustChain. In the previously mentioned literature [18], the scalability of the architecture is assessed. To facilitate the academic community to verify that the performance exposed in the aforementioned paper also applies to the presented architecture, we have made our own assessment exposed in 5. Besides this assessment, we encourage the reader to conduct a performance evaluation with the same settings as described in the literature [18] and substituting the type of transactions for content-exchange transactions (create and transfer of content). The second contribution is the identity layer. In Section 4, it is exposed how the identity layer relies on a third party to verify identities. This trusted third party is therefore able to trace back any transaction to the identity of the original owner, adding traceability and impermeability towards illegal activities to the architecture. The academic community can verify such claims by checking the ownership chain stored in the block of each piece of content. In the readMe uploaded, a section explaining how to retrieve such information is described. Lastly, the generic coinAPI allows being extended with any coin system. To facilitate the academic community to explore extending this API, a detailed guide is present in the readMe of the released codebase explaining the steps to follow to integrate a coin transfer system with the architecture. We have therefore taken significant steps to increase reproducibility and enable the academic community to easily verify our conclusions.

The final aspect of the responsible research section concerns privacy. The main privacy aspect of the architecture is related to identities. As mentioned in the previous paragraph, an identity layer differentiates the outline solution from current existing architectures. In order to ensure the highest level of privacy, the identity layer is composed of Self-Sovereign Identities attested by a trusted third-part who verifies the public key of the user, linking it to an identity. Once attested, the user can sign transactions with its public key, maintaining anonymity in the network.

8 Future Work

The architecture presented in this paper lays the foundation for a framework in which artists can share, cooperate and transfer all types of content. However, there are some steps the solution outlined in section 3 needs to take to close the gap from a proof-of-concept architecture to a practical framework ready for deployment. Such steps will be elicited in this section together with future improvements over the current architecture and extra features which were excluded from this first iteration due to time constraints.

In the first place, the development of the marketplace is a key element of the mentioned gap between a research project and the practical use case. Although developing a marketplace has no research foundation, it is a necessary step to allow users to intuitively search, discover and trade content. In future iterations, the public keys will be eliminated from the user journey and the user will be able to choose based on previews of the content similar to how users buy content in Amazon.

In the second place, one of the problems this architecture tries to tackle is the lack of scalability of the Ethereum network. This architecture is the most used for decentralized content transfer. We have discussed how by design, substituting a fraud prevention approach for a fraud detection approach, yields better scalability. For this reason, TrustChain is more scalable than the Ethereum network. We have evaluated the performance of a single peer in the network in Section 5. However, a scalability assessment with multiple peers managing content simultaneously is yet to be performed. A similar assessment has been done in the previously cited literature [18], in which it was shown TrustChain has no problem scaling up to 10.000 users and beyond. At this point, we can conclude that the literature performance evaluation together with our evaluation is a good starting point for future performance and scalability assessments.

Thirdly, the initial concept of the architecture was to use the Euro-token for the transfer of collaborative content. This was later refined to allow for generic digital coin transfer to reduce entry barriers. Nevertheless, in future versions, an implementation of the CoinAPI allowing for the use of the Euro-token could be developed to showcase the potential of UniCon. The Euro-token is the most suited coin for demonstration purposes because it is a stablecoin, pegged to the value of the Euro, and therefore with low volatility; the support of the European organizations would increase its reputation, a key aspect to consider when aiming for mass adoption; and, it has already been integrated into the SuperApp developed in the Delft Blockchain Lab, to which our system has been integrated.

Lastly, in earlier sections, we mentioned the need to create unique identifiers for physical assets to be able to manage them in our system. In order to do this successfully, cooperation with relevant authorities such as UNESCO, real state organizations, and many others is needed to jointly digitalize the management of such assets.

9 Conclusions

The research question posed in this paper was: "*How can we design a universal, scalable architecture with verified Self-Sovereign Identities and generic coin transfer?*". We have defended that the answer to this question is UniCon.

Early in the paper, we identified the drawbacks of current solutions, both centralized, such as Spotify or Netflix, and decentralized such as NFTs marketplaces. Four main problems arose from this analysis: lack of content-sharing platforms for universal content, lack of verified Self-Sovereign identities, strong dependency on the Ether coin, and scalability issues inherited from the Ethereum network. Each of these problems leads to each one of the main pillars of the solution: Universal content, generic digital coin, verified identities, and the accountability-based blockchain. These pillars are the foundation of UniCon and aim to solve all four previously mentioned issues.

A prototype of the system has been built to prove the capabilities of the system. The relevant implementation details were described together with the system architecture in Section 3. Furthermore, an initial performance evaluation has been added proving that a single peer can produce almost 60 pieces of content a second without multi-threading. This experiment and its results are detailed in Section 5. In addition, our architecture benefits from the added capabilities in terms of scalability of TrustChain detailed in the work of Johan Pouwelse et al. [18]. Moreover, the integration of Self-Sovereign Identities allows tracing back any transaction up to a verified identity which can be held accountable while allowing the user to maintain control over its data. The ability of UniCon to manage content in any digital coin solves the dependency on Ether of current decentralized solutions, reducing volatility and the entry barriers of the system. Lastly, we widened the scope to make our system more universal without restricting its content to the digital world.

It is important to highlight that the focus of this paper was on the design; and solely for demonstration purposes, a prototype has been built. However, to fully develop the architecture and advance towards mass adoption, many more aspects need to be polished, we have mentioned the most important ones in Section 8. Nevertheless, we can conclude that we have successfully designed a universal, scalable architecture with verified Self-Sovereign Identities and generic coin transfer. Lastly, we have decoupled the creation of content from the transfer of ownership, easing the user experience and increasing the scalability of the system.

References

[1] Cambridge University Press. Definition of content.

[2] Branka Vuleta. How much data is created every day, 2021.

[3] Gareth Mitchell. How much data is on the internet?

[4] Nils; Regner, Ferdinand; Urbach and André Schweizer. Nfts in practice: Non-fungible tokens as core component of a blockchain-based event ticketing application, 2019.

[5] Alyssa Hertig. Which crypto projects are based on ethereum?, 2021.

[6] Ledger Insights. Jack dorsey sells his first tweet as an nft for $2.9 million, 2021.

[7] Robert Frank. Nft sales top $2 billion in first quarter, with twice as many buyers as sellers, 2021.

[8] Jordan Lyanchev. High ethereum transaction fees, 2021.

[9] Caspar Roelofs Jack Tanner. Nfts and the need for self-sovereign identity, 2021.

[10] Alexander Mühle, Andreas Grüner, Tatiana Gayvoronskaya, and Christoph Meinel. A survey on essential components of a self-sovereign identity, 2018.

[11] Johan Pouwelse. Trustchain protocol: draft-pouwelse-trustchain-01, 2018.

[12] Om; Singh Mor Tejinder; Singh Mor Tejinder Chauhan, Anamika; Malviya. Blockchain and scalability, 2018.

[13] Valentina Di Liscia. Reports of stolen art on nft marketplace raise issues for crypto collectors, 2021.

[14] Raynor de Best. Number of cryptocurrencies worldwide from 2013 to 2021, 2021.

[15] R. W. Blokzijl. A central bank digital currency (cbdc) with off-line transfers, 2021, Unpublished Master Thesis.

[16] Rainer Lasch Jacob Lohmer. Blockchain in operations management and manufacturing: Potential and barriers, computers industrial engineering, 2020.

[17] Olsthoorn M. Pouwelse J. de Vos, M. Devid: Blockchain-based portfolios for software developers, 2019.

[18] Johan Pouwelse Martijn de Vos. Contrib: Maintaining fairness in decentralized big tech alternatives by accounting work. computer networks, volume 192, 2021.

[19] Bengt Carlsson and Rune Gustavsson. The rise and fall of napster - an evolutionary approach. 12 2001.

[20] Zehui Xiong, Yang Zhang, Cong Nguyen, Dusit Niyato, Ping Wang, and Nadra Guizani. The best of both worlds: A general architecture for data management in blockchain-enabled internet-of-things. *IEEE Network*, 01 2020.

[21] Hongmin Gao and Zhen Wang Zhaofeng Ma, Ming Jiang. Blockchain for digital rights management. *Future Generation Computer Systems*, 89:746–764, 2018.

[22] Qin Wang, Rujia Li, Qi Wang, and Shiping Chen. Non-fungible token (nft): Overview, evaluation, opportunities and challenges, 2021.

[23] Pablo Rodrigo Valero. Release of unicon, 2021, https://doi.org/10.5281/zenodo.5032008.