

FoodFresh: Multi-Chain Design for an Inter-Institutional Food Supply Chain Network

Philipp Stangl and Christoph P. Neumann 

Department of Electrical Engineering, Media and Computer Science
Ostbayerische Technische Hochschule Amberg-Weiden
Amberg, Germany
e-mail: {p.stangl1 | c.neumann}@oth-aw.de

Abstract—We consider the problem of supply chain data visibility in a blockchain-enabled supply chain network. Existing methods typically record transactions happening in a supply chain on a single blockchain and are limited in their ability to deal with different levels of data visibility. To address this limitation, we present FoodFresh – a multi-chain consortium where organizations store immutable data on their blockchains. A decentralized hub coordinates the cross-chain exchange of digital assets among the heterogeneous blockchains. Mechanisms for enabling blockchain interoperability help to preserve the benefits of independent sovereign blockchains while allowing for data sharing across blockchain boundaries.

Keywords—*blockchain; consortium; supply chain network; controlled transparency; interoperability.*

I. INTRODUCTION

The food industry comprises companies dedicated to manufacturing and processing raw materials and semi-finished products from agriculture, forestry, and fishing. In recent years, food supply chains have progressed from shorter, independent to more unified, coherent relationships among supply chain participants [1]. Developing long-term, and collaborative relationships requires evolutionary technological solutions to simultaneously retain a competitive edge.

Blockchain technology is considered a way to increase supply chain visibility, support fraud detection and provide supply chain optimization. Current applications of blockchain technology in food supply chain management, e.g., IBM Food Trust [2], rely mainly on a single distributed ledger. The implications on supply chain networks are twofold: (i) organizations participating in multiple supply chains must share their data on multiple blockchains, and (ii) participants may see information originally not intended for them because all participants can view every transaction on a distributed ledger. In a single-chain approach with just one ledger, all data would be shared publicly with all other chain participants.

The multi-chain requirement is motivated by achieving controlled transparency, i.e., to enable all parties to control visibility of data based on two levels of chains. Each participant is provided with a chain of type *permissioned*, and sharing data is provided by an additional chain of type *public*. The permissioned chains are subject to a Role Based Access Control (RBAC) mechanism, thus, its information

is hidden from the public and accessible to all users that belong to an organization. Providing organizations each with their own permissioned chain, interconnecting them as a federated ecosystem with a public chain also simplifies the addition or removal of individual organizations from the overall ecosystem with minimal impact.

In this paper, we propose FoodFresh – a multi-chain approach for inter-institutional supply chain networks, allowing organizations to store immutable data on their blockchain. A decentralized hub coordinates the cross-chain communication among the heterogeneous blockchains. The hub further ensures that all parties comply with the overarching rules of the consortium.

The remainder of the paper is organized as follows: in Section II, a selection of related work is presented. Subsequently, an overview of the relevant technology is given in Section III. Next, Section IV discusses our proposal with the design rationale. We conclude the paper in Section V, followed by the references at the end.

II. RELATED WORK

Recently, various solutions for blockchain-enabled supply chains have been proposed. For instance, Longo *et al.* have presented a software connector to connect an Ethereum-like public blockchain with an enterprise information system [3]. The software connector allows companies to share information with their partners with different levels of visibility. Schulz and Freund [4] have proposed a blockchain-enabled distributed supply chain. Their main idea is a network-centric design, which incorporates domain-specific blockchains for handling specific business processes and a hub or main blockchain that connects the blockchains to communicate with each other.

Polkadot uses a hybrid consensus model, separating block production (Blind Assignment of Blockchain Extension (BABE)) from finality (GHOST-based Recursive Ancestor Deriving Prefix Agreement (GRANDPA)). This allows for blocks to be rapidly produced and finalized at a slower pace without risking slower transaction speeds or stalling. Polkadot provides cross-chain communication with arbitrary data. Parachains communicate through the Cross-Chain Message Passing (XCMP) protocol, a queuing communication mechanism based on a Merkle tree. XCMP

is designed to communicate arbitrary messages between parachains. Messages are sent together with the next parachain block (short: parablock), while the relay chain blocks include only the proof of postage. All messages must be processed in proper order, for which a chain of Merkle proofs is used. However, XCMP is still under development. Therefore, the stop-gap protocol is Horizontal Relay-routed Message Passing (HRMP). As soon as XCMP is fully developed, it can replace HRMP. The primary difference between the two is the data stored on the relay chain. In HRMP, the relay chain stores the full message with its payload. XCMP, on the other hand, will only store a reference to the payload. The target parachain will be responsible for decoding the message payload.

From the perspective of inter-institutional supply chains, FoodFresh extends our previous work on inter-institutional cooperation [5]–[8] that was focused on healthcare, in which central organizations from primary care and secondary care act as leaders and hubs of cooperation. The FoodFresh scenario extends our perspective to more decentralized and autonomous institutional cooperation in a food supply chain network, without central protagonists.

III. BACKGROUND

This section provides a brief overview of the different relevant technologies: Section III-A describes the characteristics of food supply chain networks, Section III-B presents different types of blockchain technology, and Section III-C different blockchain interoperability approaches.

A. Food Supply Chain Network

A supply chain is an interconnection of organizations, activities, resources, people, and information. Organizations along a food supply chain are dedicated to growing and processing raw materials (e.g., fruits) and semi-finished products (e.g., fruit juices) for delivery to the end customer. Food supply chains are complex and affected by various factors, such as the sociopolitical environment [9]. Regulatory bodies, such as the US Department of Agriculture (USDA), aim to protect consumer health and increase economic viability. Thus, they release frequent updates to ensure their criteria are met by food supply chains.

In a Food Supply Chain Network (FSCN), more than one supply chain and more than one business process can be identified, both parallel and sequential in time. The parties involved in the business processes depend on the type of FSCN. This article considers a FSCN for fresh agricultural products.

Van der Vorst *et al.* have identified farmers, retailers, and their logistics service suppliers as parties involved in a FSCN for fresh agricultural products [9]. Figure 1 depicts such a supply chain at the organization level within the context of a FSCN for fresh agricultural products. Each organization is positioned in a product lifecycle stage and belongs to at least one supply chain. That means an organization can have multiple suppliers and customers

at the same time and over time. Figure 1 visualizes this by showing the perspective of the processor (bold lines), who has multiple connections to distributors and farmers. Other stakeholders, such as nongovernmental organizations, governments, and shareholders, are indirectly involved at each stage of the product lifecycle.




Figure 1. Schematic diagram of an FSCN (based on Van der Vorst *et al.* [9])

B. Blockchain Types

There are three different types of blockchain systems [10]. Public blockchains are considered permissionless because, in principle, everyone can attend the consensus process and read the stored data. The application of public blockchains has several use cases, including cryptocurrencies and document validation. In a consortium blockchain, an elected group of participants is allowed to attend the consensus process. The stored data may be read by selected members or by the public. Supply chain and research environments are two exemplary use cases for this type of blockchain. In a private blockchain, all participants belong to the same organization, and the public cannot access the system. Two use cases for this final blockchain type are banking and asset ownership. Private and consortium blockchains are considered permissioned blockchains because, in both cases, only a limited group can attend the consensus process.

C. Blockchain Interoperability

Blockchain interoperability involves the ability of independent distributed ledger networks to communicate with each other. Various approaches have been established to provide blockchain interoperability, resulting in a highly fragmented market [11]. Belchior *et al.* were the first to conduct a systematic literature review in [11] on blockchain interoperability solutions: Their resulting Blockchain Interoperability Framework categorizes interoperability solutions into three categories: 1) interoperability across public blockchains (public connectors), 2) independent blockchains that interoperate among each other (blockchains of blockchains), and finally, 3) approaches that neither fit into the public connectors nor blockchains of blockchains category (hybrid connector).

IV. FOODFRESH

In this section, we describe a consortium blockchain for a food supply chain network for interoperability and

controlled transparency. Section IV-A, introduces the approach. The three tiers of the system architecture are described in the following sections: the presentation tier in Section IV-B1, the application tier in Section IV-B2, and the relay tier in Section IV-B3. In Section IV-C, we offer a concise introduction to the Substrate Framework. The subsequent Section IV-D, delineates the details concerning the deployment process. Finally, Section IV-E addresses the limitations of our proposed approach.

A. FoodFresh Approach

The FoodFresh approach provides an implementation of the multi-chain approach (Section II). The blockchain consortium comprises a multi-chain ecosystem for organizations. Each organization is allowed to participate in the consensus process. A permanent and shared record of food system data connects participants across the food supply chain network. This is done through the use of a main blockchain, called relay chain. The sole purpose of the relay chain is to coordinate and share appropriate data and ensure all parties are complying with overarching rules. Each organization can set up and manage its own permissioned blockchain, which keeps full control over the data to itself. Within a single permissioned blockchain for an organization, data is shared between different users that belong to that organization, but not to inter-institutional parties. Via the public relay, the FoodFresh approach allows them to share immutable and accurate data with other participants in the inter-institutional supply network. This also allows for the addition or removal of individual organizations from the overall ecosystem with minimal impact.

B. System Architecture

FoodFresh, as a distributed system, is a composition of three tiers. This section will outline each of the three tiers. The presentation tier in Section IV-B1, the application tier in Section IV-B2, and finally the relay tier in Section IV-B3. Figure 3 depicts the system architecture for two interoperating supply chain organizations.

1) *Presentation Tier*: To provide the user with convenient access to the FoodFresh system, the presentation tier is responsible for interacting with the application tier through a websocket connection. Any websocket-capable client or device can communicate with the endpoints exposed by the application tier. The user interacts with a Graphical User Interface (GUI) to manage the permissions of participating members, register shipments and products, and trace shipments along the supply chain. A browser extension is required to manage blockchain accounts and to sign transactions within those accounts.

2) *Application Tier*: The application tier encompasses application-specific blockchains (the parachains) that allow organizations to join with their blockchain, where they can store immutable data. Through this, organizations can create products and shipments. A shipment’s storage and

transportation conditions can be monitored and tracked through the supply chain. The business logic is decomposed in tightly coupled modules called pallets. Figure 2 depicts the business logic pallets, each with its provided functionality that can be invoked via transactions on the parachain. Additionally, an Off-Chain Worker (OCW) is used to communicate the latest shipment status with the external world. With the subsystem Cumulus, parachains can send and receive cross-chain messages and enable validators to validate their state transitions. RBAC, formalized by Ferraiolo *et al.* [12], has become the predominant model for user access control. RBAC is used in the FoodFresh approach to control the access in terms of who can submit transactions. The *rbac* pallet maintains an on-chain registry of roles and the users to which those roles are assigned. A role is a tuple with the name of a pallet and a permission that qualifies the level of access granted by the role. A permission is an enumeration with the variants *Execute* and *Manage*. The *Execute* permission allows a user to invoke a pallet’s dispatchable functions. The *Manage* permission allows a user to assign and revoke roles for a pallet, and also implies the *Execute* permission. Access control validation is done within the transaction pool of a parachain.

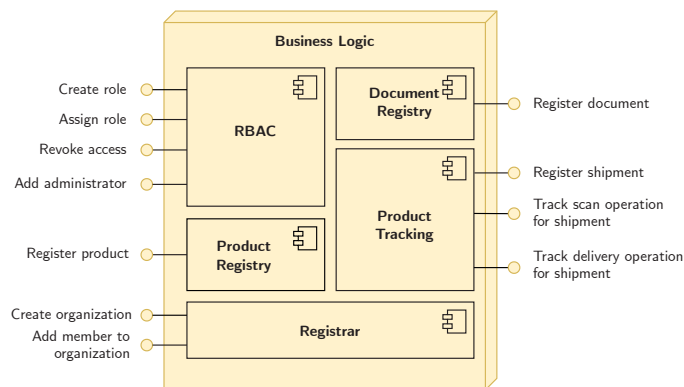


Figure 2. Overview of the business logic, decomposed into five pallets

3) *Relay Tier*: The relay chain, in the relay tier, is the essential hub in the network of heterogeneous blockchains, the parachains. The relay chain provides parachains with parablock validation and allows them to communicate with each other using the Cross-Chain Messaging (XCM) format for cross-chain messaging.

Validators are the actors of the relay chain and have three responsibilities: (1) to verify that the information contained in parablocks is valid, such as the identities of the transacting parties, (2) to participate in the consensus mechanism to produce the relay chain blocks based on validity statements from other validators, and (3) to handle cross-chain messages. For validators to fulfill their responsibilities, they are equipped with six primary runtime modules. The *inclusion* module handles the inclusion and availability of parablocks. In addition, *shared* manages the shared storage and configurations for other validator




Figure 3. Overview of our approach. The architecture is composed of three tiers: presentation, application, and relay.

modules. The *paras* module manages the chain-head and validation code for parachains. The *scheduler* is responsible for parachain scheduling, as well as validator assignments for the consensus mechanism. The *validity* module addresses secondary checks and disputes resolution for available parablocks. Finally, the *XCMP* module handles cross-chain messages and ensures that the messages are relayed to the receiving parachain.

An integral part of cross-chain communication is the establishment of a cross-chain messaging channel between the validators of two communicating parachains. Burdges *et al.* [13] have stated that a messaging channel aims to guarantee four things: “First that messages arrive quickly; second that messages from one parachain arrive to another in order; third that arriving messages were indeed sent in the finalized history of the sending chain; and fourth that recipients will receive messages fairly across senders, helping guarantee that senders never wait indefinitely for their messages to be seen”.

The act of removing an organization from the ecosystem does not necessitate the elimination of its associated parachain. This concept is facilitated by the existence of a systematic protocol, specifically the modification of the relay chain validator registry. Through the processes of registration and deregistration, organizations are added to and removed from this registry. Structurally, this registry is characterized as a hash map, a data structure that comprises paired elements: a unique identifier (ID) and the corresponding parachain ID. It should be noted that the relationship between organizations and their parachains is fundamentally non-destructive, meaning that the alterations in the organization’s status within the ecosystem do not directly impinge on the existence of the related parachain.

C. Substrate Framework

FoodFresh is built with substrate [14], a modular framework for building blockchains. A nontechnical reason for using substrate is its flexibility. Organizations must be able to adapt their blockchain system to meet the supply chain compliance requirements of regulatory bodies. Regulations happen frequently, especially in food supply chains, as shown in Section III-A. Due to the modular nature of substrate-based blockchains, developers have the necessary freedom to swap or add modules to their blockchain runtime.

Technical reasons include the chosen programming language, the software design, and the off-chain abilities. Substrate is implemented in the programming language Rust, which aims to provide performance (comparable to C++), reliability, and better means of productivity. In terms of reliability, Rust manages resources (including memory, files, network, and thread) and avoids problems, such as resource leaks or data races. Finally, for productivity, Rust provides Integrated Development Environment (IDE) support and type inspections. Furthermore, substrate is

generic by design, meaning transactions are abstracted to so-called extrinsics (things that happen outside the chain) and intrinsics (things that happen inside the chain). Transactions are stored as binary large objects. As a result, users can transfer and store any type of data on the blockchain.

Nonetheless, with FoodFresh as a permissioned blockchain, concerns about off-chain processes need to be raised. For instance, Helliar *et al.* have made the assumption that “off-chain processes may become a major barrier for permissioned blockchains” [15]. Using substrate, off-chain data can be queried or processed before it is included in the on-chain state through OCW, a collator node subsystem that allows for the execution of long-running and possibly nondeterministic tasks. Moreover, an OCW does not influence the block production time.

D. Deployment

FoodFresh requires validator nodes for the relay chain and collator nodes for the parachains to be set up by the organizations participating in a supply chain network. Nodes can be deployed locally or remotely via a cloud service provider, such as Amazon Web Services. Before parachains can participate in cross-chain communication, they need to be registered on the relay chain. The following rule is defined in the Collator Protocol [16], which implements the network protocol for the Collator-to-Validator networking: To accept n parachain connections, $n + 1$ validator nodes need to run on the relay chain. For the FoodFresh prototype, two relay chain nodes are started to connect one parachain node. Further, the relay chain needs to obtain the hex-encoded parachain’s genesis state (exported from a collator node) and the WebAssembly runtime validation function to validate parablocks.

E. Limitations

While the FoodFresh approach offers a comprehensive framework for leveraging blockchain technology in food supply chain management, there are several potential limitations and areas of concern, including:

Scalability: Parachains might face scalability challenges, depending on the scale of the organizations involved and the number of transactions. These issues are typically dependent on their specific implementation, the consensus mechanism used, and the volume of transactions they handle. If an organization’s parachain is not optimized to handle large quantities of data or high transaction throughput, it could become a bottleneck that slows down the overall system’s performance.

Complexity of Implementation: The FoodFresh approach, with each organization having its own blockchain and one relay chain for cross-communication, increases the complexity of the system compared to commonly used single blockchains. This presents significant challenges in terms of maintenance and understanding the system for non-technical stakeholders.

Adoption Challenges: Organizations might be reluctant to adopt the proposed approach due to perceived risks, lack of understanding, or the costs involved in implementation and training.

Evaluation: The software prototype is implemented and available on GitHub [17] under Apache License 2.0. A key part of designing a supply-chain network is ensuring the network is versatile enough to cope with future risks. The current solutions to analyze and mitigate endogenous risks lack continuous monitoring, as a result, risks from irregularities (e.g. abnormal order quantities by retailers) remain mostly undetected. Thus, a core part of our future evaluation is to answer whether we can develop an approach to detect abnormal activity in a multi-chain scenario. Our plans will focus on capturing the variability of transfer volume in cross-chain messaging in order to detect abnormal activity in blockchain-enabled inter-institutional supply chain networks.

V. CONCLUSION AND FUTURE WORK

Developing long-term and increasingly collaborative relationships among supply chain participants requires advanced technological solutions to retain a competitive edge. Blockchain is presented as a promising technology that might increase supply chain visibility and improve efficiency. We have presented FoodFresh – a multi-chain consortium for an inter-institutional food supply chain network. This approach overcomes the challenges associated with current approaches (e.g., IBM Food Trust), such as lack of controlled transparency and restricted interoperability among supply chain participants. By implementing a multi-chain consortium with an overseeing decentralized hub, FoodFresh allows organizations to maintain their independent blockchains, thereby preserving data sovereignty and enabling effective data exchange across blockchain boundaries. The design approach used for FoodFresh could apply to other networks that require the distribution or transfer of sensitive data. Future work could apply the approach to other industries, for instance, healthcare. The safe and secure transfer of patient health records or other sensitive information between healthcare providers, insurance companies, and the patients themselves is a major concern in the healthcare industry. The presented approach could allow each party to maintain control over their data while enabling necessary data sharing.

REFERENCES

- [1] M. A. Bourlakis and P. W. H. Weightman, *Food supply chain management*. John Wiley & Sons, 2008.
- [2] IBM Corporation, “About IBM food trust,” Jun. 2019, [Online]. Available: <https://www.ibm.com/downloads/cas/8QABQBDR> (visited on 05/23/2023).
- [3] F. Longo, L. Nicoletti, A. Padovano, G. d’Atri, and M. Forte, “Blockchain-enabled supply chain: An experimental study,” *Computers & Industrial Engineering*, vol. 136, pp. 57–69, 2019.
- [4] K. F. Schulz and D. Freund, “A Multichain Architecture for Distributed Supply Chain Design in Industry 4.0,” in *International Conference on Business Information Systems*, Springer, 2018, pp. 277–288.
- [5] C. P. Neumann, *Distributed Case Handling*. München: Verlag Dr. Hut, 2013, ISBN: 9783843909198.
- [6] C. P. Neumann, F. Rampp, M. Daum, and R. Lenz, “A Mediated Publish-Subscribe System for Inter-Institutional Process Support in Healthcare,” in *Proc of the 3rd ACM Int’l Conf on Distributed Event-Based Systems (DEBS 2009)*, Nashville, TN, USA, Jul. 2009, 14:1–14:4.
- [7] C. P. Neumann and R. Lenz, “The alpha-Flow Approach to Inter-Institutional Process Support in Healthcare,” *International Journal of Knowledge-Based Organizations (IJKBO)*, vol. 2, no. 4, pp. 52–68, 2012.
- [8] C. P. Neumann and R. Lenz, “A Light-Weight System Extension Supporting Document-based Processes in Healthcare,” in *Proc of the 3rd Int’l Workshop on Process-oriented Information Systems in Healthcare (ProHealth’09) in conjunction with the 7th Int’l Conf on Business Process Management (BPM’09)*, Ulm, DE, Sep. 2009, pp. 557–568.
- [9] J. Van der Vorst, A. Beulens, and T. van Beek, “Innovations in logistics and ICT in food supply chain networks,” *Innovation in Agri-Food systems*, pp. 245–291, Jan. 2005.
- [10] Z. Zheng, S. Xie, H. Dai, X. Chen, and H. Wang, “An overview of blockchain technology: Architecture, consensus, and future trends,” in *2017 IEEE International Congress on Big Data*, IEEE, 2017, pp. 557–564.
- [11] R. Belchior, A. Vasconcelos, S. Guerreiro, and M. Correia, “A Survey on Blockchain Interoperability: Past, Present, and Future Trends,” *ACM Computing Surveys (CSUR)*, vol. 54, no. 8, pp. 1–41, 2021.
- [12] D. Ferraiolo, D. R. Kuhn, and R. Chandramouli, *Role-based access control*. Artech house, 2003.
- [13] J. Burdges *et al.*, “Overview of polkadot and its design considerations,” *ArXiv*, vol. abs/2005.13456, 2020.
- [14] Parity Technologies, “Substrate - The Blockchain Framework for a Multichain Future,” 2020, [Online]. Available: <https://substrate.io/> (visited on 05/23/2023).
- [15] C. V. Helliar, L. Crawford, L. Rocca, C. Teodori, and M. Veneziani, “Permissionless and permissioned blockchain diffusion,” *International Journal of Information Management*, vol. 54, pp. 102–136, 2020, ISSN: 0268-4012.
- [16] Parity Technologies, “Collator Protocol,” Jun. 2021, [Online]. Available: <https://paritytech.github.io/polkadot/book/node/collators/collator-protocol.html> (visited on 05/23/2023).
- [17] P. Stangl and C. P. Neumann, “FoodFresh,” May 2023, [Online]. Available: <https://github.com/cyberlytics/FoodFresh> (visited on 05/23/2023).