MedTech Chain

Decentralised, Secure and Privacy-preserving Platform for Medical Device Data Research

Petru-Rosu, Alin; Tataru, Tamara; Zelenjak, Jegor; Kromes, Roland; Erkin, Zekeriya

DOI
10.1109/BRAINS63024.2024.10732045

Publication date
2024

Document Version
Final published version

Published in
2024 6th Conference on Blockchain Research and Applications for Innovative Networks and Services, BRAINS 2024

Citation (APA)
Petru-Rosu, A., Tataru, T., Zelenjak, J., Kromes, R., & Erkin, Z. (2024). MedTech Chain: Decentralised, Secure and Privacy-preserving Platform for Medical Device Data Research. In N. Salhab (Ed.), *2024 6th Conference on Blockchain Research and Applications for Innovative Networks and Services, BRAINS 2024* (2024 6th Conference on Blockchain Research and Applications for Innovative Networks and Services, BRAINS 2024). IEEE. https://doi.org/10.1109/BRAINS63024.2024.10732045

2024 6th Conference on Blockchain Research & Applications for Innovative Networks and Services (BRAINS) | 979-8-3503-6784-3/24/$31.00 ©2024 IEEE | DOI: 10.1109/BRAINS63024.2024.10732045

MedTech Chain: Decentralised, Secure and Privacy-preserving Platform for Medical Device Data Research

Alin Petru-Rosu, Tamara Tataru, Jegor Zelenjak, Roland Kromes, and Zekeriya Erkin
Cyber Security Group, Department of Intelligent Systems
Delft University of Technology
Van Mourik Broekmanweg 5, 2628XE, Delft, The Netherlands
a.rou@student.tudelft.nl, t.tataru@student.tudelft.nl, j.zelenjak@student.tudelft.nl,
r.g.kromes@tudelft.nl, z.erkin@tudelft.nl

***Abstract*—Rapid advancements in digital medical technologies have significantly improved patient care but have also raised complex security and privacy challenges. Traditional tools for detecting vulnerabilities in networked medical devices, primarily used by network administrators and security specialists, have become insufficient due to their large-scale use across the entire healthcare network. Aiming to improve security in healthcare, MedTech Chain proposes a way to solve this challenge by leveraging blockchain and privacy-enhancing technologies, offering an authenticated, decentralised, secure, and privacy-preserving environment for the research and monitoring of medical device data. Currently, the framework enables counting, averaging, and grouped counting queries with multiple filtering capabilities like time frame and location. Such functionalities can provide valuable insights not only for threat intelligence but also for medical research and hospital management. MedTech Chain is modular and flexible, designed to seamlessly extend to new device technologies and research demands. To our knowledge, the approach is among the first to employ ϵ-differential privacy in the context of medical device data.**

***Index Terms*—blockchain, healthcare, security, privacy, IoT, networked medical devices, ϵ-differential privacy, Hyperledger Fabric**

I. Introduction

Recent technological advancements have enabled the fast-paced innovation of the healthcare sector. Networked devices, cloud computing, and decentralised architectures are promising new solutions that meet the demands of a modern, paperless society. However, with rapid digitalisation, the threat landscape expands, putting the safety of health providers and their patients in danger. As medical infrastructure is critical, disregarding its security and privacy concerns can bring disastrous consequences. Compromised devices can not only disclose sensitive information, causing financial and reputational damage to the care provider but also put patients in life-threatening conditions [1]. Moreover, the frequency of cyberattacks on healthcare providers has surged, with reports indicating an alarming average of 1,684 attacks per week in the first quarter of 2023 [2]. As such, the relevance of cybersecurity in healthcare becomes substantial.

In the context of networked medical devices, addressing cybersecurity threats is particularly challenging [1], [3]. The interoperability between medical devices and healthcare systems, i.e., seamless communication and functioning, introduces additional security vulnerabilities, as each point of interaction can potentially be exploited. Additionally, the variety of device types makes it difficult to enforce consistent security measures across the entire healthcare network [4]. Consequently, hospitals have to rely instead on conventional tools used by network administrators and security personnel.

In light of the vulnerabilities caused by rapid digitalisation, much of the existing research has centred on safeguarding medical data. Still, the potential of the information regarding medical devices in enhancing threat intelligence is generally neglected. Medical *device* data refers only to information related to the device itself, such as technical specifications; it should not be misunderstood for medical data (e.g., patient vital signs like heart rate). Device data can provide invaluable insights into device behaviour, essential to identifying and mitigating potential cyber threats [5]. While the security of patient data remains a priority, device data can help develop more comprehensive security strategies that protect all aspects of healthcare technology.

Understanding and monitoring the expansive attack surfaces of medical devices remains difficult. The challenge primarily comes from the lack of efficient methods for collecting comprehensive data about medical infrastructure, which is essential for robust cyber threat intelligence [6]. Data like operating system versions of medical devices must be constantly monitored to reduce the attack space. The challenge is compounded by stringent privacy concerns and regulatory frameworks restricting cybersecurity researchers' access to necessary data. Additionally, the heterogeneity of medical infrastructure, which varies widely in composition and technology, adds another layer of complexity to data-gathering.

Previous efforts to address the challenges of monitoring and securing medical device infrastructures. The current approaches involve manual data collection and monitoring methods, which are inefficient for modern medical environments' large-scale and diverse nature [1]. Attempts have been made to implement centralised data repositories for storing and managing medical device information. However, these approaches suffer from single points of failure and increased vulnerability to cyberattacks. Additionally, centralised systems often face regulatory compliance and privacy issues, requiring extensive access to sensitive information, which is hard to manage and secure [6]. Despite these efforts, the lack of scalable, efficient, and privacy-preserving solutions has limited the effectiveness of these methods, stressing the need for more advanced approaches.

We propose MedTech Chain, a solution that leverages blockchain and privacy-enhancing technologies to facilitate the research of networked medical devices. Our solution uses blockchain's inherent security and transparency to ensure that device data is securely recorded, verifiable, and traceable while maintaining the privacy of healthcare providers' infrastructure through ϵ-differential privacy. ϵ-differential privacy is key in preventing the inference of sensitive information about healthcare providers or patients using the given device. For instance, the number of patients within a particular hospital can be deduced by counting the number of active devices. The MedTech Chain architecture is designed to support a robust platform where medical device data can be queried, analysed, and monitored without compromising data privacy or integrity.

MedTech Chain's objective is to provide a secure, accessible, and privacy-preserving environment for investigating medical device data. The platform facilitates the seamless sharing of medical device data across a diverse network of users, including medical professionals, researchers, and health administrators, all while upholding strict privacy standards. Our ultimate goal is to empower the healthcare community by providing reliable and secure data insights and stimulating advancements in networked medical device research.

The platform's vision is supported and aligns with initiatives like the SEPTON[1] project, which aims to develop a secure data-sharing platform for cyber threat intelligence and statistical analysis on medical devices. Integrating MedTech Chain with SEPTON to securely manage medical device data will support the latter's efforts to enhance cybersecurity and privacy in healthcare. This collaboration ensures that both projects can achieve their shared objective of safeguarding healthcare infrastructure by enabling an efficient research environment of networked medical devices.

II. Background

A. Blockchain Technology in Healthcare

Blockchain technology is increasingly recognised for its potential to address critical security and privacy challenges in the healthcare sector. Its core attributes—decentralisation, immutability, and transparency—make it a robust technology for managing sensitive healthcare data. Blockchain operates as a distributed ledger, eliminating the need for a central authority. This decentralisation enhances security by reducing single points of failure, making it harder for cyber attackers to compromise the system. Once data is recorded on a blockchain, it cannot be altered or deleted. The immutability ensures the integrity and authenticity of the data, which is essential for maintaining accurate records [7].

In healthcare, it is crucial to maintain precise device histories and prevent counterfeit or compromised devices from tampering with the data. Moreover, blockchain provides a transparent ledger where all transactions are visible to all participants. This transparency can enhance trust among stakeholders by ensuring accountability and traceability. Blockchain's decentralised architecture supports interoperability, allowing seamless and secure data exchange between healthcare systems and devices [8]. Together, these make blockchain suitable for MedTech Chain, offering a safe, verifiable, and tamper-resistant environment for managing medical device data.

B. Hyperledger Fabric: A Permissioned Blockchain

Hyperledger Fabric[2] is an open-source blockchain framework designed to support the development of enterprise-grade applications [9]. It provides a modular architecture with high flexibility, scalability, and security, making it suitable for healthcare applications. Unlike public blockchains, Hyperledger Fabric operates on a permissioned network where participating organisations are known and authenticated. This feature is critical for healthcare environments where data access needs to be controlled and monitored. Hyperledger Fabric provides different services to handle blockchain transaction verification and their storage on the ledger. A Hyperledger Fabric network comprises three kinds of nodes: endorser peers, committer peers and orderers. The endorsers' role is to verify the submitted transaction's validity, while the orderers' objective is to place the valid transactions to a new block in order. Finally, the committer peer verifies block validity and commits the transaction on the ledger within a block. Hyperledger Fabric allows customising various components, such as consensus mechanisms, membership services, and data privacy protocols. These features make Hyperledger Fabric a good candidate for implementing our platform.

C. Differential Privacy: Ensuring Data Privacy

ϵ-differential privacy is a technique that ensures the privacy of individual data entries while allowing aggregate data analysis [10], [11]. The method involves adding controlled noise to the results of queries on sensitive data. This noise prevents identifying individual data points and ensures that privacy is maintained even when aggregated data is analysed. The strength of ϵ-differential privacy is quantified by a privacy parameter ϵ. A lower ϵ value indicates more robust privacy protection, making inferring information from the query results difficult. ϵ-differential privacy can be applied to query results, ensuring that device data analysis does not compromise the privacy of healthcare providers or patients while enabling aggregate data analysis for research purposes.

[1] https://septon-project.eu/

[2] https://www.hyperledger.org/projects/fabric

III. Related Work

Integrating big data into healthcare, driven by advancements in machine learning and artificial intelligence, offers notable potential for improving medical practices. However, it also introduces constant privacy and security challenges that lack comprehensive solutions. Price et al. highlight the legal and ethical issues these developments raise, particularly concerning patient privacy. They stress the need for robust data handling practices and strong governance frameworks to prevent discrimination and data breaches, yet practical implementations are limited [12]. Apte et al. examine the practical challenges of compiling comprehensive health records, including coordinating with various administrative entities, securing a centralised data repository, defining relevant health metrics, and integrating data for research. They point out that managing health data effectively is challenging due to these multifaceted obstacles, and solutions remain elusive [13]. While big data holds promise for transforming healthcare, the persistent lack of solutions to privacy and security challenges and the complexity of data management pose significant problems for realising its full potential.

Blockchain offers a promising approach to addressing these challenges by leveraging its inherent decentralisation, immutability, and transparency features. Moreover, blockchain's distributed nature removes the need for a central authority, ensuring data integrity by linking records in a tamper-resistant chain [14], [15]. The following research articles detail potential blockchain systems available for medical data, focusing primarily on security and privacy concerns.

The Ancile framework [16], proposed by Dagher et al., represents a significant advancement in electronic health record (EHR) security and interoperability by integrating blockchain technology. Ancile's decentralised nature of blockchain reduces risks associated with centralised data storage, enhancing data security and accessibility. Utilising Ethereum blockchain-based smart contracts [17], Ancile enforces stringent access controls, ensuring medical records are accessible only to authorised individuals and automating access policy enforcement. Advanced cryptographic techniques, including data obfuscation and encryption, protect data integrity and confidentiality, making patient information private and resistant to unauthorised access. These features collectively position Ancile as a comprehensive solution to persistent issues in healthcare data management, such as data breaches, significantly contributing to the literature on secure and interoperable health information systems.

In a related study, Huang et al. propose a blockchain-based scheme that balances patient privacy concerns with the needs of research institutions and commercial entities. Their approach allows for the secure sharing of medical data among stakeholders, including patients and researchers, utilising zero-knowledge proofs to enable institutions to verify data relevance without compromising privacy. Additionally, proxy re-encryption ensures that data passed through semi-trusted intermediaries remains secure and private. This method emphasises the importance of maintaining data availability and privacy within intelligent healthcare environments [18].

In the realm of decentralised privacy-preserving healthcare systems, Dwivedi et al. introduced a pioneering framework that integrates blockchain with IoT to address the pervasive security and privacy concerns in managing healthcare data. Their solution emphasises a decentralised architecture, ensuring data integrity and confidentiality by leveraging the blockchain's immutable and transparent ledger while using cryptographic techniques to protect patient data. This framework facilitates secure and efficient data sharing among healthcare stakeholders, reducing the risks associated with centralised data storage and enabling a more resilient and trustworthy healthcare information system. The framework safeguards sensitive information and ensures compliance with regulatory standards using smart contracts and encryption. The robust security measures implemented in this framework make it a significant advancement in healthcare IoT, providing a scalable and secure solution for modern medical data management challenges [19].

Current literature [20], [21] on healthcare technologies often highlights the critical need for enhanced medical data security and privacy. However, the focus predominantly remains on medical data, with almost no mention of the equally crucial medical device data encapsulating connected medical devices' operational characteristics and vulnerabilities.

To address these gaps, recent studies have begun to explore the integration of blockchain and IoT technologies for healthcare applications. For instance, Karunarathne et al. examine the current state of security and privacy of the Internet of Things in the healthcare system to highlight challenges and advocate for comprehensive security solutions [3]. Furthermore, Christidis et al. analyse how blockchains and smart contracts work to identify the possible advantages their introduction brings to a system and highlight how blockchains and IoT can be used together [22].

Still, to the best of our knowledge, no specific tools currently exist tailored to facilitate the large-scale monitoring research of networked medical devices, apart from general tools used by network administrators and security specialists. While many proposals advocate for decentralised data storage using blockchain technology, they typically do not address the entire lifecycle of data privacy, from collection to analysis. These gaps underscore the importance of developing solutions that improve security at all healthcare system levels while preserving stakeholders' privacy.

IV. The MedTech Chain Framework

A. System Architecture

Figure 1 shows the architecture of the MedTech Chain. There are two types of participant organisations, i.e., multiple healthcare facilities/hospitals providing device data and one semi-trusted organisation enabling researchers to query the data—referred to as the MedTech Chain Organisation.

1) Hospitals: The current use case considers two types of medical devices, i.e., portable and wearable, whose data must be stored on the ledger. Each hospital deploys an

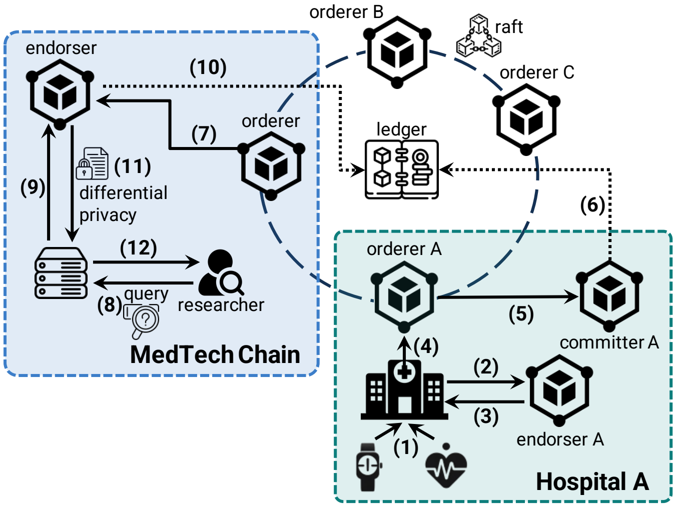

Fig. 1: MedTech Chain Architecture

application that bridges the devices and the blockchain, namely *collectors*, as displayed in Figure 1. The collectors gather and process device data, as illustrated in step (1). Subsequently, the collectors will regularly submit blockchain transactions with new/updated device data. The transaction passes the following flow before being committed on the ledger: the transaction is validated by an endorser, depicted by steps (2) and (3). Next, if valid[3], it has to be submitted to the ordering layer, represented by step (4). Currently, the orderers are configured to use the Raft consensus [23]. Then, after the transactions are ordered into blocks and consistency is assured, the transaction is submitted within a block to the committer, which appends it to the blockchain, as shown by steps (5) and (6). Finally, step (7) indicates that the ledger state modification will become visible to the entire MedTech Chain network.

Note that handling the interconnection between the devices and the collectors is out of the scope of our work.

2) The MedTech Chain Organisation: The system responsible for query functionalities is maintained by the MedTech Chain Organisation. Researchers can submit queries to the *user application*. Query submission is represented as step (8). The application translates the query into a transaction further evaluated by an endorser, as depicted by step (9). It should be noted that this transaction is read-only and, therefore, does not need to be submitted to the ordering layer. During evaluation, the invoked smart contract will read the ledger's state, indicated by step (10), run the query, apply ϵ-differential privacy and return the result to the application, as in step (11). Lastly, the result returns to the researcher, shown in step (12).

B. Functionality

1) User Application: The user application is hosted as a distinct semi-trusted service where only registered users, such as researchers or administrators, can gain access. It has an interface facilitating efficient interactions with the data stored on the blockchain. Researchers interact with a dedicated interface for data analysis, supporting complex queries that provide detailed insights into the available medical device data. Currently, in our prototype, user management is also enabled by this application, with the administrators having access to a dedicated interface that allows the creation, modification and revocation of researcher accounts. Nevertheless, we encourage migrating towards a more modular architecture to extract user management as a separate service, allowing for better integration of more complex features.

2) Data Schema: Listing 1 outlines a simplified version of our experimental medical device data schema. They include common fields such as medical speciality, manufacturer, model and firmware version. From a cyber threat perspective, an interesting use case is the frequency of communication with another external service for synchronisation. Such information provides insights into possible anomalies indicating security breaches or attempted intrusions. Unusual communication frequencies or unexpected data transfer times might signal that a device has been compromised or is being tampered with. Regarding hospital management, another use case could be monitoring the average age of the infrastructure, enabling hospitals to identify if they are using older devices that may need to be replaced soon. Furthermore, the schema's flexible design allows for incorporating additional use cases, such as maintenance updates or usage patterns, ensuring that it can adapt to future advancements and the evolving needs of the medical technology landscape.

Listing 1: Schema for Medical Device Data

```
message DeviceData {
  string udi = 1;
  string manufacturer = 2;
  string model = 3;
  string firmware_version = 4;
  string device_type = 5;
  string device_category = 6;
  string production_date = 7;
  string last_service_date = 8;
  string warranty_expiry_date = 9;
  string usage_hours = 10;
  string battery_level = 11;
  string hospital = 12;
  string speciality = 13;
  string active_status = 14;
  string last_sync_time = 15;
  string sync_frequency = 16;
  string mac_address = 17;
}
```

Nevertheless, this schema is a prototype. Designing a real-world schema requires collaborative efforts from multiple parties and thus is outside the scope of the paper. While the schema must contain data that provides useful insights, it should also be compliant with legal frameworks such as GDPR [24] and HIPAA [25]. Therefore, researchers, privacy experts, and medical device specialists must carefully design and agree upon the schema to capture each device's relevant

[3] The transaction is considered valid when a registered blockchain member signs it with respect to certain agreed policies.

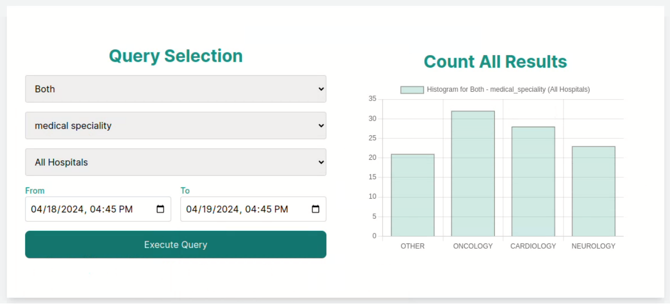

Query Selection

Both

medical speciality

All Hospitals

From 04/18/2024, 04:45 PM

To 04/19/2024, 04:45 PM

Fig. 2: Graphical Interface for the *Grouped Counting* Query. The figure illustrates the result of a query counting the number of both wearable and portable devices in all hospitals grouped by medical speciality over a specified time interval.

aspects and particularities, enabling comprehensive data analysis while maintaining regulatory compliance. Additionally, the success of this schema depends on ongoing collaboration, rigorous validation, and continuous adaptation to the evolving landscape of medical technology and regulatory requirements.

The schema must be easily extensible, agreed upon, and distributed among all participating stakeholders, which is achieved by implementing it using Protocol Buffers[4]. Protocol Buffers provide platform-neutral, extensible mechanisms for serialising structured data, facilitating easy updates and extensions for new device types and research requirements. The proposed device data structure is modular, allowing for easy extension and removal of unused fields. Additionally, all participating parties must respect this data structure, meaning all stakeholders must approve any modifications. Schema agreement and distribution are managed through the blockchain. Once a modification is approved, the updated version of the device data structure is stored in the blockchain ledger, ensuring accessibility for all participants in the network.

3) Queries: Researchers can perform three types of queries:

- **Count**: Researchers can retrieve the count of wearable or portable devices and filter based on a specified hospital over a defined period. Additional filters can be selected, and the dynamic interface is tailored to capture the dependencies of the queried data. For instance, the researchers can further specify a firmware version filter when querying information related to a specific device model and manufacturer.
- **Grouped Counting**: This query allows researchers to count the number of wearable or portable devices grouped by a specified property. As depicted in Figure 2, researchers could determine how many devices are used in each medical speciality. Similarly to counting, filters can be applied, such as selecting a hospital or a time frame. The result is currently visualised as a histogram, yet the design allows easy extension to different plot types.
- **Average**: Similar to counting, the query can calculate the average of particular properties of wearable or portable devices filtered on various specifications.

Due to the modular design, new queries can be easily added, and the existing ones can be extended. Importantly, all queries apply ϵ-differential privacy to ensure that the privacy of individual data entries is preserved.

C. Scalability

The solution's scalability is directly dependent on the number of participating hospitals and devices. These impact many aspects, including network transaction load and query response time. The current assumption is that device data is infrequently stored on the ledger and updated based on events rather than periodically. Precisely, a device would submit a transaction to update its data asset when there is a state modification, such as a firmware version update or the device switched on/off. This would enable a decreased network load. Regarding query response time, it scales linearly with the number of medical devices across all hospitals. In the worst case, the query has to read all device data assets in the world state of the blockchain.

D. Security and Privacy

MedTech Chain runs on a permissioned blockchain network built using Hyperledger Fabric. The network's permissioned nature is reinforced by Membership Service Providers (MSPs), which manage identities and authenticate members, granting access only to verified individuals or organisations. Additionally, Hyperledger Fabric utilises a unique channel architecture, allowing transactions and communications to occur within isolated channels among participants, further securing data exchange and limiting access to authorised stakeholders only.

One authorised participant must be the MedTech Chain Organisation, surrounded by the blue frame in Figure 1, our approach's main assumption. This organisation is responsible for facilitating research of medical devices and is assumed to be a semi-trusted entity, potentially coordinated by a regulatory body, government agency, or non-governmental organisation.

On the one hand, the framework tries to preserve the privacy of the hospitals by employing ϵ-differentially private queries. All currently implemented queries are anonymised using the Laplace mechanism, and the noise addition is performed within smart contracts. Applying ϵ-differential privacy at this level is necessary as no third party can be trusted to obtain the exact results of the queries, not even the MedTech Chain Organisation. This technique consequently improves the privacy of all the participant hospitals.

Privacy and utility are balanced based on the ϵ parameter, which controls the amount of noise added to query results. Normally, the parameter is determined for a fixed dataset representing the best trade-off between privacy and utility. Since the dataset continuously expands as more device data is stored on the ledger in time, the value of ϵ should be dynamically configurable. Note that the parameter's value depends on the queried data, and since the data schema currently used is experimental, the parameter value was not our priority. Rather, the focus was on implementing differential privacy and enabling specialists to configure the privacy parameters further.

[4] https://protobuf.dev/

On the other hand, since the entire network operates on a permissioned blockchain, the MedTech Chain Organisation plays a crucial role in integrating healthcare providers into the network. This involves regulating and verifying the legitimacy of participants, ensuring they meet the necessary standards. Additionally, the organisation must be vigilant in monitoring for any potential collusion or fraudulent activities to maintain the integrity and trustworthiness of the network. Nevertheless, even the organisation's actions and decisions must be verifiable through independent checks and audits.

Another assumption is secure and reliable communication channels, which are currently implemented in the MedTech Chain prototype. The communication channels within the system employ Transport Layer Security (TLS) to secure interactions between nodes, such as peers, orderers, and client applications, ensuring data integrity and confidentiality. Furthermore, TLS is also used to secure communication between the graphical interface and the backend of the user application.

Regarding user authentication, researchers are issued credentials upon initial registration and subsequent administrative approval. The authentication process is safeguarded by a system-generated password of configurable length, which users must change upon their first login. A strict password policy is enforced to prevent unauthorised access. Additionally, JSON Web Tokens (JWT)[5] are utilised for session management and to facilitate authenticated communication between administrators/researchers and the user application.

Finally, not even the participants in the network should have direct access to individual device data in plaintext, as this represents a significant privacy concern. While our prototype does not currently address this issue, this topic will be discussed thoroughly in section VI.

V. Experiments and Results

The following experiment aims to highlight the time overhead caused by applying ϵ-differential privacy on top of the query. In our setting, 10,000 assets containing medical device data were stored in the ledger to mimic a solid database. To compare query computation between the two scenarios, 1000 query transactions were sent when ϵ-differential privacy was applied and 1000 queries when not. Elapsed time is measured as the average of the 1000 transaction evaluations.

Regarding the Hyperledger Fabric network deployment, the network contains 3 hospitals and the MedTech Chain Organisation, each containing one orderer and a peer acting as endorsers and committers (i.e., 4 orderers and 4 peers). The network applies the default configuration with a single channel and a batch timeout of 2 seconds.

The entire MedTech Chain architecture, including the blockchain network and the user application, was deployed on an AMD EPYC 9334 32-core server with 128 CPUs operating at a minimum 1.5 GHz frequency, with 32 cores per socket. The network overhead is negligible as the complete MedTech Chain infrastructure was deployed locally on the above-mentioned server. Table I shows the time to process a query with and without applying ϵ-differential privacy.

TABLE I: Query execution time

	Query Response Time
DP not applied	243.103 ms
DP applied	243.275 ms

The results highlight that query processing involves a negligible, constant time overhead of 0.172 ms when ϵ-differential privacy is applied. This experiment specifically measures the impact of ϵ-differential privacy on query execution time, demonstrating that it incurs minimal performance penalty. This negligible time overhead indicates that ϵ-differential privacy can be applied effectively in a real-time context without significantly impacting data processing speed, thus ensuring both data privacy and operational efficiency are maintained. It should be noted that the query processing time, which is embedded in the blockchain transaction, might vary depending on the network connection and blockchain settings (e.g., number of peers, transaction validation time), number of assets, etc. However, processing the ϵ-differential privacy will be close to constant for a given ϵ value and data size.

VI. Discussions and Limitations

As previously discussed, no participant should be able to access plain device data of individual hospitals directly, as this could lead to significant privacy and security issues. Access to plain device data can expose vulnerabilities within a hospital's technological infrastructure, which might be exploited by malicious actors participating in the network. Additionally, aggregated device data can reveal patterns that indirectly expose sensitive information, such as the number of patients being treated or specific treatments being administered, thus violating regulatory requirements like GDPR [24] and HIPAA [25]. Furthermore, sharing device data among participants could place hospitals at a competitive disadvantage, allowing competitors to influence their decisions based on this information.

Consequently, more robust cryptographic techniques are needed to ensure data privacy. Device data should be encrypted before storage on the blockchain, a homomorphic cryptosystem [26] being a suitable option. Homomorphic encryption allows computations on encrypted data without decryption, enabling basic data operations while preserving privacy (when differential privacy is applied after the encryption of the data). Additionally, a threshold scheme necessitates the collaboration of multiple parties during decryption, preventing any single participant from unilaterally accessing or manipulating the data. This protects against insider threats and unauthorised access, as the network requires agreement from all participants to decrypt device data or query results. These measures collectively ensure that MedTech Chain can evolve as a highly secure and trustworthy platform.

Nevertheless, while critical from the legal perspective, employing encryption increases the system's complexity and

[5] https://jwt.io/introduction

raises implementation concerns like key distribution and management. Encryption can also impact various aspects of the current solution.

The current implementation involves on-chain storage, with each device data asset being estimated to have an average size of 300 bytes considering the message schema from Listing 1. Once encryption is employed, the storage size of each asset depends on the encryption scheme. For instance, the Paillier homomorphic encryption scheme [27] with a key size of 2048 bits leads to ciphertexts of size 4096 bits (512 bytes). Since each field must be encrypted individually to perform homomorphic operations (e.g., counting the number of active devices), this would lead to an asset requiring 512*17=8704 bytes (8.5 kilobytes). These storage demands might force MedTech Chain to migrate towards off-chain storage.

Alternatively, each hospital could store its device data in a self-managed database, with each asset referenced via a hash stored in the ledger. This would provide better storage scaling and enable the storage of large volumes of data. However, security and availability depend on the off-chain storage solution, and ensuring data integrity becomes more complex, aspects crucial for MedTech Chain.

In another direction, encryption directly impacts query execution time due to induced communication and time complexity. The queries are implemented based on homomorphic operations, which are computationally expensive. Additionally, different schemes can require multi-party computation to decrypt the intermediary or final query results.

While storing unencrypted data on the ledger can be considered a system limitation regarding privacy, encryption comes with multi-faced challenges. Still, assessing the feasibility of employing encryption within the system requires some prototypical implementation. This will enable experiments that quantify the encryption's impact properly, which is future work for MedTech Chain.

VII. Future Work

Along with the platform's evolution, key areas of development are identified to further improve MedTech Chain.

Our current approach involves a semi-trusted third party, i.e., the MedTech Chain Organisation. An alternative approach could be a fully decentralised query management system, where each participating hospital independently manages its querying application. This decentralisation would distribute the responsibility and control across multiple entities, potentially increasing the system's resilience to single points of failure and reducing opportunities to misuse centralised power. However, this approach may lead to increased administrative and technical challenges, including inconsistency in data management standards and increased operational costs. Each hospital would need to independently ensure robust data security and privacy measures, requiring substantial investment to maintain system integrity across the network.

Several additional features could be integrated, improving research capabilities and overall system management. Extending existing queries would allow researchers to gain deeper insights into usage patterns of portable and wearable devices. With sufficient device data, researchers could possibly identify usage patterns by investigating when portable or wearable devices function, aiming to predict the frame for possible attacks. Additionally, implementing user action tracking is essential for auditing, ensuring regulatory compliance, and identifying unauthorised activities carried out by the researchers. Furthermore, integrating an encryption scheme management system could facilitate the seamless modification and migration towards secure encryption schemes over time. These enhancements might significantly improve the research capabilities, security, and usability of the MedTech Chain, making it a more robust and comprehensive solution for managing medical device data in healthcare environments.

VIII. Conclusion

In this paper, we introduce MedTech Chain, a blockchain-based platform for secure and privacy-preserving research on networked medical device data. The platform leverages blockchain technology to create a robust, scalable, decentralised infrastructure, ensuring data integrity and transparency. A novel aspect of MedTech Chain is its use of ϵ-differential privacy in the context of medical device data, which guarantees data privacy during analysis.

MedTech Chain offers valuable functionalities, such as counting, averaging, and grouped counting, with multiple filtering options, including time frame and medical specialities. These functionalities are embedded in query smart contracts, which apply ϵ-differential privacy with a negligible overhead of approximately 0.2 ms to the query execution time. This ensures that query results do not leak information, thereby enhancing the privacy and security of the data while still allowing meaningful analysis. Consequently, researchers and hospital administrators can obtain valuable insights from aggregated data without risking exposure to sensitive information.

Differential privacy is a widely adopted technique, crucial in the health sector and many other domains. In the Industrial IoT (IIoT) domain, the need to protect production efficiency and production type is significant [28]. In the public sector, web-based mapping applications and recommendation systems should also apply this technique to protect individuals' privacy [29]. It is, therefore, trivial that the application of differential privacy is domain-independent, which has enabled its integration into our framework.

The device data within MedTech Chain includes comprehensive information such as device type, manufacturer, firmware versions, and medical speciality, as well as operational metrics like usage patterns and communication frequencies. This device data scheme can facilitate the identification of anomalies indicative of security breaches or operational issues, thereby enabling threat intelligence and optimising device management strategies.

Moreover, MedTech Chain's user interface provides a seamless and intuitive experience and allows users to access and analyse medical device data securely and efficiently. This user-friendly interface broadens the platform's accessibility, making it valuable for many users, from cyber threat intelligence services to hospital management, thus fostering a collaborative and efficient research environment.

In summary, MedTech Chain balances data utility and privacy, ensuring compliance with stringent privacy regulations while fostering a trustworthy collaborative research and management environment. Its comprehensive design and implementation make it a significant advancement in the secure and private analysis of medical device data.

Acknowledgements

This work is supported by the European Union's Horizon Europe research and innovation programme, grant No. 101094901 (SEPTON). The source code is available here[6].

References

[1] L. Coventry and D. Branley, "Cybersecurity in healthcare: A narrative review of trends, threats and ways forward," *Maturitas*, vol. 113, pp. 48–52, 2018. [Online]. Available: https://doi.org/10.1016/j.maturitas.2018.04.008

[2] Check Point Research, "Global Cyberattacks Continue to Rise with Africa and APAC Suffering Most," *Check Point Blog*, Apr 2023, visited on May 17, 2024. [Online]. Available: https://blog.checkpoint.com/research/global-cyberattacks-continue-to-rise/

[3] S. M. Karunarathne, N. Saxena, and M. K. Khan, "Security and Privacy in IoT Smart Healthcare," *IEEE Internet Computing*, vol. 25, no. 4, pp. 37–48, 2021. [Online]. Available: https://doi.org/10.1109/MIC.2021.3051675

[4] P. A. H. Williams and V. McCauley, "Always connected: The security challenges of the healthcare Internet of Things," in *3rd IEEE World Forum on Internet of Things, WF-IoT 2016, Reston, VA, USA, December 12-14, 2016.* IEEE Computer Society, 2016, pp. 30–35. [Online]. Available: https://doi.org/10.1109/WF-IoT.2016.7845455

[5] B. Hodges, J. T. McDonald, W. Glisson, M. Jacobs, M. V. Devender, and J. H. Pardue, "Attack Modeling and Mitigation Strategies for Risk-Based Analysis of Networked Medical Devices," in *53rd Hawaii International Conference on System Sciences, HICSS 2020, Maui, Hawaii, USA, January 7-10, 2020.* ScholarSpace, 2020, pp. 1–10. [Online]. Available: https://hdl.handle.net/10125/64538

[6] I. Lee, O. Sokolsky, S. Chen, J. Hatcliff, E. Jee, B. Kim, A. L. King, M. Mullen-Fortino, S. Park, A. Roederer, and K. K. Venkatasubramanian, "Challenges and Research Directions in Medical Cyber-Physical Systems," *Proc. IEEE*, vol. 100, no. 1, pp. 75–90, 2012. [Online]. Available: https://doi.org/10.1109/JPROC.2011.2165270

[7] K. Wüst and A. Gervais, "Do you Need a Blockchain?" in *Crypto Valley Conference on Blockchain Technology, CVCBT 2018, Zug, Switzerland, June 20-22, 2018.* IEEE, 2018, pp. 45–54. [Online]. Available: https://doi.org/10.1109/CVCBT.2018.00011

[8] A. Hasselgren, K. Kralevska, D. Gligoroski, S. A. Pedersen, and A. Faxvaag, "Blockchain in healthcare and health sciences - A scoping review," *Int. J. Medical Informatics*, vol. 134, 2020. [Online]. Available: https://doi.org/10.1016/j.ijmedinf.2019.104040

[9] E. Androulaki *et al.*, "Hyperledger Fabric: A Distributed Operating System for Permissioned Blockchains," in *Proceedings of the Thirteenth EuroSys Conference*, ser. EuroSys '18. New York, NY, USA: ACM, 2018, pp. 30:1–30:15. [Online]. Available: http://doi.acm.org/10.1145/3190508.3190538

[10] C. Dwork, "Differential Privacy: A Survey of Results," in *Theory and Applications of Models of Computation, 5th International Conference, TAMC 2008, Xi'an, China, April 25-29, 2008. Proceedings*, ser. Lecture Notes in Computer Science, M. Agrawal, D. Du, Z. Duan, and A. Li, Eds., vol. 4978. Springer, 2008, pp. 1–19. [Online]. Available: https://doi.org/10.1007/978-3-540-79228-4_1

[11] J. Ficek, W. Wang, H. Chen, G. Dagne, and E. Daley, "Differential privacy in health research: A scoping review," pp. 2269–2276, 2021. [Online]. Available: https://doi.org/10.1093/jamia/ocab135

[12] W. N. Price and I. G. Cohen, "Privacy in the age of medical big data," *Nature Medicine*, vol. 25, 2019. [Online]. Available: https://doi.org/10.1038/s41591-018-0272-7

[13] M. Apte, M. Neidell, E. Y. Furuya, D. Caplan, S. Glied, and E. Larson, "Using electronically available inpatient hospital data for research." *Clinical and translational science*, vol. 4, 2011.

[14] S. Angraal, H. M. Krumholz, and W. L. Schulz, "Blockchain technology: Applications in health care," vol. 10. Lippincott Williams and Wilkins, June 2017. [Online]. Available: https://doi.org/10.1161/CIRCOUTCOMES.117.003800

[15] T. Kuo, H. Kim, and L. Ohno-Machado, "Blockchain distributed ledger technologies for biomedical and health care applications," *J. Am. Medical Informatics Assoc.*, vol. 24, no. 6, pp. 1211–1220, 2017. [Online]. Available: https://doi.org/10.1093/jamia/ocx068

[16] G. G. Dagher, J. Mohler, M. Milojkovic, and P. B. Marella, "Ancile: Privacy-preserving framework for access control and interoperability of electronic health records using blockchain technology," *Sustainable Cities and Society*, vol. 39, pp. 283–297, 5 2018. [Online]. Available: https://doi.org/10.1016/j.scs.2018.02.014

[17] G. Wood *et al.*, "Ethereum: A secure decentralised generalised transaction ledger," *Ethereum project yellow paper*, vol. 151, 2014.

[18] H. Huang, P. Zhu, F. Xiao, X. Sun, and Q. Huang, "A blockchain-based scheme for privacy-preserving and secure sharing of medical data," *Comput. Secur.*, vol. 99, p. 102010, 2020. [Online]. Available: https://doi.org/10.1016/j.cose.2020.102010

[19] A. D. Dwivedi, G. Srivastava, S. Dhar, and R. Singh, "A Decentralized Privacy-Preserving Healthcare Blockchain for IoT," *Sensors*, vol. 19, no. 2, p. 326, 2019. [Online]. Available: https://doi.org/10.3390/s19020326

[20] K. Abouelmehdi, A. B. Hssane, H. Khaloufi, and M. Saadi, "Big data security and privacy in healthcare: A review," in *The 8th International Conference on Emerging Ubiquitous Systems and Pervasive Networks (EUSPN 2017)*, ser. Procedia Computer Science, E. M. Shakshuki, Ed., vol. 113. Elsevier, 2017, pp. 73–80. [Online]. Available: https://doi.org/10.1016/j.procs.2017.08.292

[21] J. J. Hathaliya and S. Tanwar, "An exhaustive survey on security and privacy issues in Healthcare 4.0," pp. 311–335, 2020. [Online]. Available: https://doi.org/10.1016/j.comcom.2020.02.018

[22] K. Christidis and M. Devetsikiotis, "Blockchains and Smart Contracts for the Internet of Things," pp. 2292–2303, 2016. [Online]. Available: https://doi.org/10.1109/ACCESS.2016.2566339

[23] D. Ongaro and J. Ousterhout, "In Search of an Understandable Consenus Algorithm (Extended Version)," *Proceedings of USENIX ATC '14*, 2014.

[24] K. Ider, "Assessment of the quality of user awareness of GDPR in healthcare IOT," in *2021 International Conference on Biomedical Innovations and Applications (BIA)*, vol. 1, 2022, pp. 25–28.

[25] S. Mbonihankuye, A. Nkunzimana, and A. Ndagijimana, "Healthcare Data Security Technology: HIPAA Compliance," *Wirel. Commun. Mob. Comput.*, vol. 2019, pp. 1 927 495:1–1 927 495:7, 2019. [Online]. Available: https://doi.org/10.1155/2019/1927495

[26] I. Damgård and M. Jurik, "A Length-Flexible Threshold Cryptosystem with Applications," in *Information Security and Privacy, 8th Australasian Conference, ACISP 2003, Wollongong, Australia, July 9-11, 2003, Proceedings*, ser. Lecture Notes in Computer Science, R. Safavi-Naini and J. Seberry, Eds., vol. 2727. Springer, 2003, pp. 350–364. [Online]. Available: https://doi.org/10.1007/3-540-45067-X_30

[27] P. Paillier, "Public-Key Cryptosystems Based on Composite Degree Residuosity Classes," in *Advances in Cryptology - EUROCRYPT '99, International Conference on the Theory and Application of Cryptographic Techniques, Prague, Czech Republic, May 2-6, 1999, Proceeding*, ser. Lecture Notes in Computer Science, J. Stern, Ed., vol. 1592. Springer, 1999, pp. 223–238. [Online]. Available: https://doi.org/10.1007/3-540-48910-X_16

[28] B. Jiang, J. Li, G. Yue, and H. Song, "Differential Privacy for Industrial Internet of Things: Opportunities, Applications, and Challenges," *IEEE Internet Things J.*, vol. 8, no. 13, pp. 10 430–10 451, 2021. [Online]. Available: https://doi.org/10.1109/JIOT.2021.3057419

[29] C. Dwork, N. Kohli, and D. K. Mulligan, "Differential Privacy in Practice: Expose your Epsilons!" *J. Priv. Confidentiality*, vol. 9, no. 2, 2019. [Online]. Available: https://doi.org/10.29012/jpc.689

[6] https://github.com/orgs/MedTechChain