

FHIRChain: Applying Blockchain to Securely and Scalably Share Clinical Data

Peng Zhang^a, Jules White^a, Douglas C. Schmidt^a, Gunther Lenz^b,
S. Trent Rosenbloom^c

^a*Department of Electrical Engineering and Computer Science, Vanderbilt University, Nashville, Tennessee, USA*

^b*Varian Medical Systems, Palo Alto, California, USA*

^c*Department of Biomedical Informatics, Vanderbilt University Medical Center, Nashville, Tennessee, USA*

Abstract

Secure and scalable data sharing is essential for collaborative clinical decision making. Conventional clinical data efforts are often siloed, however, which creates barriers to efficient information exchange and impedes effective treatment decision made for patients. This paper provides four contributions to the study of applying blockchain technology to clinical data sharing in the context of technical requirements defined in the “Shared Nationwide Interoperability Roadmap” from the *Office of the National Coordinator for Health Information Technology* (ONC). First, we analyze the ONC requirements and their implications for blockchain-based systems. Second, we present FHIRChain, which is a blockchain-based architecture designed to meet ONC requirements by encapsulating the HL7 *Fast Healthcare Interoperability Resources* (FHIR) standard for shared clinical data. Third, we demonstrate a FHIRChain-based decentralized app using digital health identities to authenticate participants in a case study of collaborative decision making for remote cancer care. Fourth, we highlight key lessons learned from our case study.

Keywords: Blockchain, Smart Contracts, Decentralized App,

Email addresses: peng.zhang@vanderbilt.edu (Peng Zhang), jules.white@vanderbilt.edu (Jules White), d.schmidt@vanderbilt.edu (Douglas C. Schmidt), gunther.lenz@varian.com (Gunther Lenz), trent.rosenbloom@vanderbilt.edu (S. Trent Rosenbloom)

1. Introduction

The importance of data sharing in collaborative decision making. Secure and scalable data sharing is essential to provide effective collaborative treatment and care decisions for patients. Patients visit many different care providers' offices during their lifetime. These providers should be able to exchange health information about their patients in a timely and privacy-sensitive manner to ensure they have the most up-to-date knowledge about patient health conditions.

As another example, in telemedicine practice [1]—where patients are remotely diagnosed and treated—the ability to exchange data securely and scalably is of paramount importance. Data sharing helps improve diagnostic accuracy [2] by gathering confirmations or recommendations from a group of medical experts, as well as preventing inadequacies [3] and errors in treatment plan and medication [4, 5]. Likewise, aggregated intelligence and insights [6, 7, 8] helps clinicians understand patient needs and in turn apply more effective in-person and remote treatments.

Data sharing is also essential in cancer care, where groups of physicians with different specialties form tumor boards. These boards meet on a regular basis to analyze cancer cases, exchange knowledge, and collaboratively create effective treatment and care plans for each patient [9]. Regional virtual tumor boards are also being implemented via telemedicine [10, 11] for institutions that lack inter-specialty cancer care due to limited oncology expertise and resources [12].

Administrative support for coordinating health IT efforts. The Office of the National Coordinator for Health Information Technology (ONC) is a division of the Office of the Secretary within the United States Department of Health and Human Services. ONC is the principal federal entity to oversee and coordinate health IT efforts, including the development of interoperable, privacy-preserving, and secure nationwide health information systems and the promotion of widespread, meaningful use of health IT to improve healthcare.

Data sharing barriers to collaborative decision making. In practice, many barriers exist in the technical infrastructure of health IT systems today that impede the secure and scalable data sharing across institutions,

thereby limiting support for collaborative clinical decision making. Examples of such barriers include the following:

- **Security and privacy concerns.** Despite the need for data sharing, concerns remain regarding protection of patient identity and confidentiality [13]. For instance, virtual medical interactions may increase the risk of clinical data breaches due to electronic transmission of data without highly secure infrastructures in place, which can result in severe financial and legal consequences [14]. Likewise, medical identity theft may occur more frequently, especially in telemedicine [13], where virtual (*i.e.*, networked) interactions are replacing face-to-face interactions between providers and patients.
- **Lack of trust relationships between healthcare entities.** Trust relationships between healthcare entities [15] (*e.g.*, care providers and/or healthcare institutions) are an important precondition to digital communications [16] and data sharing in the absence of custody over shared data. Larger healthcare facilities (such as enterprise hospital systems) may be networked [17], but communications between private or smaller practices may not be established.
- **Scalability concerns.** Large-scale datasets may be hard to transmit electronically due to restrictive firewall settings or limitations in bandwidth (which is still common in rural areas [18]). Lack of scalability can also impact overall system response time and data transaction speed [19].
- **Lack of interoperable data standards enforcement.** Without the enforcement of existing interoperable data standards (such as HL7's *Fast Healthcare Interoperability Resources* (FHIR)[20] for shared data), health data can vary in formats and structures that are hard to interpret and integrate into other systems [21].

What is needed, therefore, is a standards-based architecture that can integrate with existing health IT systems (and related mobile apps) to enable secure and scalable clinical data sharing for improving continuous, collaborative decision support.

Research focus and contributions → Architectural considerations for secure and scalable blockchain-based clinical data sharing

systems. Blockchain technologies have recently been touted [22, 23, 24] as a technical infrastructure to support clinical data sharing that promotes care coordination. A key property of blockchains is their support for "trustless disintermediation." This property enables multiple parties who do not fully trust each other to exchange digital assets (such as the Bitcoin cryptocurrency [25]), while still protecting their sensitive, personal data from each other.

Our prior work [26] provided evaluation recommendations for blockchain-based health IT solutions on a high-level, focusing on common software patterns [27] that can be applied to improve the design of blockchain-based health apps. This paper examines previously unexplored research topics related to alleviating the data sharing barriers described above, namely: *what are the architectural considerations associated with properly leveraging blockchain technologies to securely and scalably share healthcare data for improving collaborative clinical decision support?*

This paper provides the following contributions to using blockchain technologies in clinical data sharing to improve collaborative decision support:

- We summarize key technical requirements defined in the "Shared Nationwide Interoperability Roadmap" [28] drafted by the *Office of the National Coordinator for Health Information Technology* (ONC) for creating an interoperable health IT system and analyze the implications for blockchain-based system design.
- We present the structure and functionality of a blockchain-based architecture called FHIRChain that meets the ONC technical requirements for sharing clinical data between distributed providers. FHIRChain uses HL7's FHIR data elements (which have uniquely identifying tags) in conjunction with a token-based design to exchange data resources in a decentralized and verifiable manner without requiring duplicated efforts of uploading data to a centralized repository.
- We demonstrate a FHIRChain-based *decentralized app* (DApp) that uses digital health identities to more readily authenticate participants and manage data access authorizations in a case study of clinical data sharing in remote cancer care. This DApp enables users to share specific and structured pieces of information (rather than an entire document), thereby increasing the readability of data and flexibility of sharing options.

- We highlight key lessons learned from our case study and discuss how our FHIRChain-based DApp can be further extended to support other technical requirements for improving advanced healthcare interoperability issues, such as coordinating other stakeholders (*e.g.*, insurance companies and pharmacies) across the industry and providing patients with direct and secure access to their own medical records. We also explore the data exchange issues that blockchains cannot yet address effectively, including semantic interoperability, healthcare malpractice, and unethical use of the data, which remain as future research problems in this space.

Paper Organization. The remainder of this paper is organized as follows: Section 2 provides an overview of blockchain technologies and the Ethereum platform, which is an open-source blockchain implementation that supports the development of DApps via “smart contracts;” Section 3 surveys different blockchain-based research approaches in the healthcare domain and compares our research on FHIRChain with related work; Section 4 summarizes ONC’s key technical requirements for sharing clinical data and analyzes their implications for blockchain-based designs; Section 5 describes how the blockchain-based architecture of FHIRChain is designed to meet ONC requirements and motivates why we made certain architectural decisions; Section 6 analyzes the benefits and limitations of a case study that applied a FHIRChain-based DApp to provide collaborative clinical decision support; and Section 7 presents concluding remarks and outlines our key lessons learned and future work on extending the FHIRChain architecture described in this paper.

2. Overview of Blockchain

The most popular application of blockchain is the Bitcoin blockchain [25], which is a public distributed ledger designed to support financial transactions via the Bitcoin cryptocurrency. This blockchain operates in a peer-to-peer fashion with all transactions distributed to each network maintainer node (called a “miner”) for verification and admittance onto the blockchain. These miners validate available transactions and group them into blocks, as shown in Figure 1. Miners then compete in solving a computationally expensive cryptographic puzzle, known as “proof-of-work,” where the first miner to solve this puzzle receiving a reward (*i.e.*, an amount of Bitcoin) and appending their block of validated transactions to the blockchain sequence.

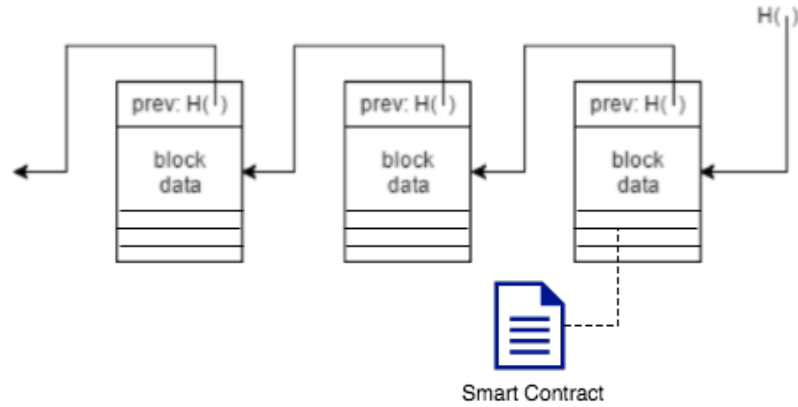


Figure 1: The Blockchain Structure: a Continuously Growing and Immutable List of Ordered and Validated Transactions

The Bitcoin blockchain uses the proof-of-work process outlined above to achieve consensus by

- incentivizing miners to contribute powerful hardware and electricity to the network with small amounts of cryptocurrency as rewards and
- discouraging rogue actors from attempting to manipulate or maliciously control the system.

After a block is added to the blockchain, its transaction history is secured from tampering via cryptography.

The Bitcoin blockchain is the most widely deployed example of this distributed ledger technology. In recent years, however, other types of blockchain technologies have emerged. For example, the Ethereum blockchain [29] provides a more generalized framework via “smart contracts” [30] that allow programs to run on the blockchain and store/retrieve information.

Smart contracts enable code to execute autonomously when certain conditions are met, as shown in Figure 1. They can also store information as internal state variables and define custom functions to manipulate or update this state. Operations in smart contracts are published as transactions and thus occur in a globally sequential order. These operations are deterministic and verifiable by miners in the Ethereum blockchain to ensure their validity.

The mechanisms described above make a blockchain decentralized and immutable, thereby removing the need for a trusted central authority. These

properties make blockchain technologies attractive to certain communities of health IT researchers and practitioners as means to improve clinical communications while protecting the privacy of healthcare participants. The remainder of this paper examines how to effectively leverage blockchains for securely and scalably sharing clinical data that enables collaborative decision support.

3. Related Work Summary and Comparison

Due to the growing interest in using distribute ledger technologies for health IT systems, related work has explored various blockchain-based design considerations and prototypes. This section summarizes this related work and compares it with our research on FHIRChain and DApps that provide collaborative clinical decision support for remote patients.

3.1. Conceptual Blockchain-Based Design Considerations

Krawiec et al. [31] presented several existing pain points in current health information exchange systems and the corresponding opportunities provided by blockchain technologies. They also discussed how blockchain can be leveraged in the health IT systems so that patients, health providers, and/or health organisations can collaborate. Nichol et al [32] presented an analysis that assembles concepts in blockchain-related technologies and speculates on how blockchain can be used to solve common interoperability problems facing healthcare.

A team at IBM [33] took a broader approach by highlighting the challenges in the healthcare industry and providing concrete use cases to showcase potential applications of blockchain technologies. Our prior work also provided software design recommendations for creating general blockchain-based health IT systems [27] and proposed assessment metrics for blockchain-based health systems [26], which include a subset of the technical requirements defined in the ONC roadmap. This prior work of ours focused on providing more general or high-level recommendations for developers creating blockchain-based health IT systems.

The review paper by Kuo et al. [34] presented several blockchain applications in healthcare, such as improved medical record management and advanced healthcare data ledger, and their benefits for each described application. They then analyzed key challenges associated with using blockchain technology for healthcare, including issues like confidentiality, scalability, and

treat of a 51% attack on the blockchain network. According to the authors, some example implementation techniques that may mitigate the challenges are (1) encryption of sensitive data or dissemination of only meta data and storing sensitive data off-chain to protect confidentiality, (2) keeping only partial, ongoing verified transactions on-chain rather than the entire transaction history to increase scalability of the blockchain network, and (3) the adoption of a virtual private network or HIPAA-compliant components to prevent the 51% attack.

3.2. Blockchain Prototype Designs

Ekblaw et al. [35] created a decentralized record management platform that enables patients to access their medical history across multiple providers. This platform used a so-called “permissioned” blockchain (which is only accessible by authorized users, rather than one that is open to the public) to manage authentication, data sharing, and other security properties in the medical domain. Their blockchain design integrated with existing provider data storage to enable interoperability by curating a representation of patient medical records. Medical researchers were incentivised to contribute to mining of the blockchain by collecting aggregated metadata as mining rewards.

Peterson et al. [36] presented a healthcare blockchain also considers the integration with FHIR standards. They proposed a merkle-tree based blockchain system that introduces “Proof of Interoperability” as the consensus mechanism during block mining. Proof of interoperability is based on conformance to the FHIR protocol, meaning that miners must verify the clinical messages sent to their blockchain to ensure they are interoperable with known structural and semantic standards.

Dubovitskaya et al. [37] also proposed a permissioned blockchain framework on managing and sharing medical records for cancer patient care. Their design employed a membership service to authenticate registered users using a username/password scheme. Patient identity was created via a combination of personally identifying information (including social security number, date of birth, names, and zip code) and encrypted for security. Medical data files were uploaded to a secure cloud server, with their access managed by the blockchain logic.

Unlike other blockchain designs, Gropper’s “HIE of One” system [38] focused on the creation and use of blockchain-based identities to credential physicians and address the patient matching challenge facing health IT

systems. Patients are expected to install a digital wallet on their personal devices to create their blockchain-based IDs, which can then be used to communicate with the rest of the network. Instead of storing patient information, Gropper’s system would consume only the blockchain-based ID and use it to secure and manage access to patient data located in EHR systems.

3.3. Differentiating Our Research Focus of FHIRChain from Related Work

This paper presents our blockchain-based framework, called FHIRChain, whose architectural choices were explicitly designed to meet key technical requirements defined by the ONC interoperability roadmap. Our design differs from related work on blockchain infrastructures and associated consensus mechanisms since it is decoupled from any particular blockchain framework and instead focuses on design decisions of smart contract and other blockchain-interfacing components. FHIRChain is thus compatible with any existing blockchains that support the execution of smart contracts.

In the remainder of this paper we describe how our FHIRChain-based DApp demonstrates the use of digital health identities that do not directly encode private information and can thus be replaced for lost or stolen identities, even in a blockchain system. While our approach is similar to the use of digital IDs in the *HIE of One*[38] system, FHIRChain provides a more streamlined solution. In addition, we incorporate a token-based access exchange mechanism in FHIRChain that conforms with the FHIR clinical data standards. Finally, we leverage public key cryptography to simplify secure authentication and permission authorizations, while simultaneously preventing attackers from obtaining unauthorized data access.

4. Technical Requirements for Blockchain-Based Clinical Data Sharing

The “Shared Nationwide Interoperability Roadmap” defines technical requirements and guiding principles for creating interoperable health IT systems [28]. Based on our experiences to date, we contend that crafting a blockchain architecture to meet these requirements necessitates overcoming significant challenges to utilize blockchain technology in healthcare most effectively.

This section first analyzes five key technical requirements fundamental to clinical data sharing systems and then discusses the implications of these requirements on blockchain-based architectures. Sections 5 and 6 subsequently

describe how we developed and applied our FHIRChain blockchain-based architecture to create a decentralized app (DApp) that meets the ONC requirements in the context of collaborative clinical decision making.

4.1. Requirement 1: Verifying Identity and Authenticating All Participants
ONC requirement summary. The ONC requirements state that an identity ecosystem should be employed to minimize identity theft and provide redress in case of medical identity fraud, while complying with individual privacy regulations. Providers, hospitals, and their health IT systems should be easily identity-proofed and authenticated when exchanging electronic health information. Healthcare systems today, however, lack “consistently applied methods and criteria” for identity proofing and authentication across organizations [28]. For example, different network service providers have different policies or requirements and may not acknowledge the methods applied by other network service providers.

One of the most popular—and least complex—approaches to exchange data is through direct secure messaging [28]. For example, the Direct project [39] was launched to create a standard way for participants to send authenticated, encrypted health information directly to known, trusted recipients over the Internet. Providers or care centers using EHR systems *without* Direct integration, however, cannot benefit from the direct exchange capability.

Implications for blockchain-based system design. For a blockchain-based system, storing identification information (such as personal email) directly on-chain is problematic [40]. In particular, a property of blockchains is information “openness,” *i.e.*, all data and associated modification records are immutably recorded and publicly available to all network participants. In the case of Bitcoin, data is open to everyone with Internet access [25], whereas in a non-public blockchain (such as a consortium blockchain [29]) data access is limited only to authenticated blockchain participants.

To meet the requirement of openness while complying to health privacy regulations [41], a blockchain-based system should thus support user identity-proofing and authentication while encapsulating sensitive personal information. Section 5.2.1 shows how FHIRChain addresses this identifiability and authorization requirement via digital health identities based on public key cryptography [42].

4.2. Requirement 2: Storing and Exchanging Data Securely

ONC requirement summary. The ONC requirements state that data should be shared securely and privately without unauthorized or unintended alteration, while making the information available to authorized parties. Data encryption is recommended both when data is sent over networks (data-in-motion) and when it is stored (data-at-rest). Management and distribution of encryption keys must be "secure and tightly controlled" [28].

Implications for blockchain-based system design. There has been recent interest [43, 44] in using blockchain technologies as decentralized storage for encrypted health data. As discussed in Section 2, however, the open and transparent nature of blockchain raises privacy concerns when attempting to integrate blockchain into the health IT domain. Although sensitive data can be encrypted, flaws in encryption algorithms or software implementations may expose the data contents in the future. To ensure long-term data security, therefore, a data storage design should be "simple" to minimize software bugs [45], *e.g.*, by not storing sensitive data (encrypted or not) on-chain, yet still enable data flow from one user to another [26].

Another implication of storing data on a blockchain is scalability. All blockchain transactions (such as storing data in a smart contract and modifying the data) and data records are distributed as an entire copy to all blockchain nodes. In a public blockchain, moreover, transaction fees are paid to miners to reward their validation efforts, as described in Section 2. As new data is added or modified, each change must be propagated to all nodes, raising scalability challenges and potentially incurring significant long-term operational costs. Section 5.2.2 shows how FHIRChain addresses this requirement via a hybrid on-chain/off-chain storage model.

4.3. Requirement 3: Consistent Permissioned Access to Data Sources

ONC requirement summary. The ONC advocates "computable privacy" that represents and communicates the permission to share and use identifiable health information [28]. Individuals should be able to document their permissions electronically, which are then honored as needed. Permission authorizations to receive or access an individual's clinical data should be accurate and trustworthy, requiring both the data requestor and holder to have a common understanding of what is authorized.

Implications for blockchain-based system design. Unfortunately, smart contract operations only occur in the blockchain space to ensure deterministic outcomes. Services (such as OAuth [46]) that exist off the blockchain therefore cannot be used. Given this constraint, incorporating other alternatives to provide data access permissioning should be a key component of a blockchain-based design. Section 5.2.3 shows how FHIRChain addresses this requirement via a token-based permission model.

4.4. Requirement 4: Applying Consistent Data Formats

ONC requirement summary. To satisfy interoperability needs, the ONC requirements state that health IT systems should be implemented with an “intentional movement and bias” [28] toward a clinical data standard identified by ONCs recently finalized *Interoperability Standards Advisory* [47]. The data exchanged should be structured, standardized, and contain discrete (granular [48]) information. Likewise, standards should use metadata to communicate their context along with pieces of structured data.

Implications for blockchain-based system design. To provide collaborative clinical decision support, health IT systems must present shared data to clinicians in a structured and readable format [49]. This requirement implies the enforcement of existing, commonly accepted clinical data standard(s), rather than introducing new data exchange formats. Section 5.2.4 shows how FHIRChain addresses this requirement by enforcing the FHIR standard.

4.5. Requirement 5: Maintaining Modularity

ONC requirement summary. The ONC requirements state that since technology inevitably changes over time, health IT system designs should be capable of evolving by maintaining modularity. When divided into connected, modular components, health IT systems become more resilient to change with increased flexibility. In turn, these properties enable the adoption of newer, more efficient technologies over time without rebuilding the entire system.

Implications for blockchain-based system design. Modularity requires a carefully crafted design to avoid “information lock-in” due to the immutability of smart contracts. Every change to a smart contract code creates a new contract instance on the blockchain, nullifying previous versions and their data. To minimize dependencies and the need to upgrade, therefore, smart contracts should be loosely coupled with other components in the system.

Section 5.2.5 shows how FHIRChain addresses this requirement by applying the *model-view-controller* (MVC) pattern [50].

5. FHIRChain: a Blockchain-Based Architecture for Clinical Data Sharing

This section first presents an overview of FHIRChain, which is a blockchain-based architecture we designed to meet the ONC requirements for secure and scalable sharing of clinical data described in Section 4. We then explain why we made specific architectural decisions in FHIRChain to address each requirement and how they solve the five challenges facing blockchain technology described in Section 4.

5.1. FHIRChain Overview

Figure 2 shows the FHIRChain architecture we devised to address key ONC technical requirements. This architecture provides a general data shar-




Figure 2: Architectural Components in FHIRChain

ing solution applicable to a wide range of health IT systems. It also serves as the basis for our decentralized app (DApp) prototype describe in Section 6, which customizes FHIRChain to support collaborative clinical decision making using a case study of cancer care in telemedicine.

The dashed ellipse in Figure 2 represents a blockchain component that mediates data sharing between collaborating medical professionals (represented by providers with green check marks). Clinical data silos are represented by heterogeneous database symbols, which we normalized with the FHIR standards to enforce a common structure of shared data. Secure database

connectors (represented as small circles) connect siloed data sources to the blockchain by exposing secure access tokens to data references that can be obtained only by authorized entities. The secure tokens are recorded in a smart contract (represented by linked documents) for decentralized access and also traceability.

In addition to storing secure access tokens, the smart contract also maintains an immutable timestamped transaction log (represented as a keyed file symbol) of all events related to exchanging and actually consuming these tokens. These logs include specific information regarding what access has been granted to which user by whom, who has consumed which token to access what resource, etc. To ensure the validity of shared data, FHIRChain can be configured to only approve participation from certified clinicians and healthcare organizations with a membership registry.

5.2. FHIRChain Architectural Decisions that Address Key ONC Technical Requirements

Below we explain why specific architectural decisions were made to address each ONC requirement presented in Section 4.

5.2.1. Addressing Requirement 1: Verifying Identity and Authenticating All Participants

Context. Blockchains like Ethereum and Bitcoin provide pseudo-anonymous personal accounts (*i.e.*, public addresses composed of random hash values) for users to transact cryptocurrencies. These native identities, however, do not address healthcare requirement for identifiability or authentication of all participants.

Problem. By design, public blockchains are globally accessible to anyone with Internet access and allow users to hold any number of blockchain accounts to minimize the identifiability of account holders. This ONC requirement, however, specifies that all U.S. healthcare participants should be identifiable, implying the need for an entirely separate, traceable user base from blockchains' native identities. A key problem is thus how to properly define identities for healthcare users participating in clinical data sharing, while protecting sensitive personal information on the blockchain.

Design choice → *use of a digital health identity.* Inspired by the success of secure shell (SSH) [51] and blockchain address generation mechanism, FHIRChain employs public key cryptography [42] to create and manage health

identities. In public key cryptography, a pair of mathematically related public and private keys is used to create digital signatures and encrypt data. Since it is computationally infeasible to obtain the private key given its paired public key, these public keys can be shared freely, thereby allowing users to encrypt content and verify digital signatures. In contrast, private keys are kept secret to ensure only their owners can decrypt content and create digital signatures.

FHIRChain generates a cryptographic public/private key pair (also used for encryption, as described in Section 5.2.3) for each participating provider, *e.g.*, in-house providers and remote physicians in telemedicine clinics. The public keys represent users' digital health identities. These identities are recorded in the blockchain for both identity- and tamper-proofing, thereby ensuring that users holding the corresponding private keys can be authenticated to use FHIRChain's data sharing service.

FHIRChain's design applies a smart contract to maintain health users' identifiability without exposing personal information on the blockchain. It also replaces the need for a traditional username/password authentication scheme with the use of a public/private cryptographic key pair for authentication. In a general clinical setting, these digital health identities (*i.e.*, their private keys) would be hard to manage for patients. FHIRChain, however, only creates these identities for clinicians to facilitate data sharing, which enables more effective collaborative decision making for patients.

5.2.2. Addressing Requirement 2: Storing and Exchanging Data Securely

Context. A key capability offered by blockchains is their support for "trustless transactions between parties who lack trust relationships established between them. Bitcoin is the most common example of this trustless exchange via its native cryptocurrency. Blockchains are peer-to-peer by nature and thus contribute to the ubiquitousness of digital assets being transacted.

Problem. Health data represented via digital assets are more complex and harder to share *en masse*. There are also privacy and security concerns associated with its storage in an "open" peer-to-peer system (*i.e.*, public blockchains), such as encryption algorithms applied to protect data being decryptable in the future [26]. A key problem is thus how to design a blockchain-based health IT system so that it balances the need for ubiquitous store and exchange and the concerns regarding privacy of the data and scalability of the system.

Design choice → *keeping sensitive data off-chain and exchanging reference pointers on-chain.* Rather than storing encrypted health data in the blockchain, a more scalable and secure alternative is to store and exchange encrypted metadata referencing protected data (*i.e.*, a reference pointer to a data set), which can be combined with an expiration configuration for short-term data sharing. Exchanging encrypted reference pointers allows providers to maintain their data ownership and choose to share data at will. This technique also prevents an attacker who intercepts the encrypted pointers from obtaining unauthorized data access.

FHIRChain attaches a secure connector to each database, as shown in Figure 2. Each connector generates appropriate reference pointers that grant access to the data. These reference pointers are digital health assets that can be transacted ubiquitously with reduced risks of exposing the data.

An added benefit of exchanging metadata *en masse* is more scalability compared to exchanging the original data source. As discussed in Section 4.2, each transaction or operation on the blockchain (*e.g.*, querying a smart contract state variable value or updating it) is associated with a small fee paid to the miner for verification and then included onto the blockchain. Transacting these lightweight reference pointers is more efficient in terms of time and cost in production because small changes to data generally require no modifications to reference pointers.

5.2.3. Addressing Requirement 3: Permission to Access Data Sources

Context. Data references can be stored on the blockchain for ubiquitous access via a smart contract. Access rights, however, must be granted only to authorized providers for viewing the data. As discussed in Section 4.3, OAuth is a popular platform for communicating permissions in web-based apps that are not based on blockchain.

Problem. Smart contracts cannot directly use external services like OAuth since they do not produce deterministic outcomes that can be verified by blockchain miners. A key problem is thus how to design a mechanism that balances the need of permission authorization for clinical data and blockchain requirements for deterministic outcomes.

Design choice → *token-based permission model.* To overcome the limitation with public blockchains, FHIRChain protects the shared content via a secure cryptographic mechanism called “sign then encrypt” [52]. This design

employs the users’ digital health identities to encrypt content so that only users holding the correct digital identity private keys can decrypt the content. FHIRChain also generates a new pair of signing keys for each participant and registers the public portion of signing keys alongside users’ digital identities.

To concretely demonstrate this workflow, Figure 3 provides an example of using FHIRChain to create and retrieve an access token. Suppose provider

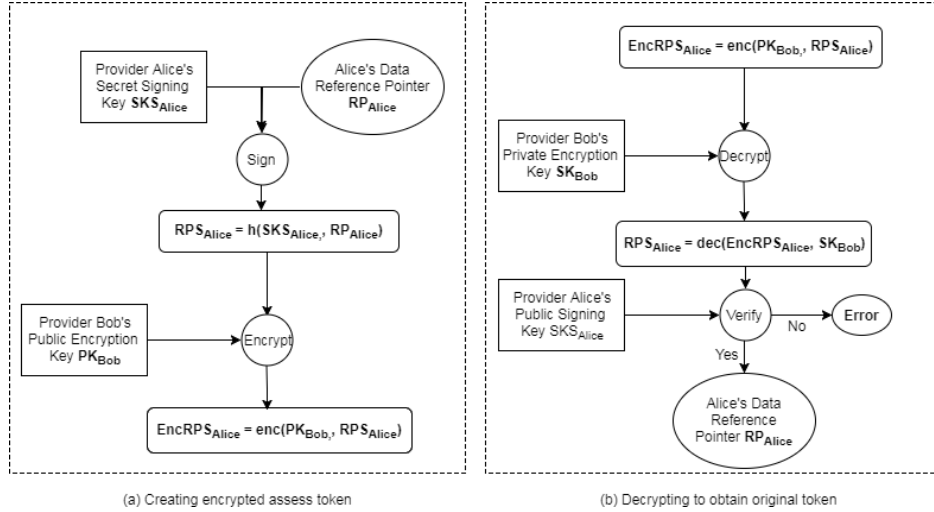


Figure 3: Example of the Creation and Retrieval of an Access Token Using FHIRChain.

Alice would like to initiate sharing of her patient’s data, denoted as D_{Alice} (with a reference pointer, denoted as RP_{Alice}) with another provider *Bob*. FHIRChain creates a digital signature on the shared content RP_{Alice} , with *Alice*’s private signing key SKS_{Alice} for tamper-proofing as a first step. With *Bob*’s public encryption key, PK_{Bob} , FHIRChain encrypts the signed RPS_{Alice} to obtain an encrypted token $EncRPS_{Alice}$, and then stores $EncRPS_{Alice}$ in a smart contract for ubiquitous access.

When *Bob* wants to obtain the content *Alice* sent, he must use his corresponding private encryption key SK_{Bob} to decipher the real content of $EncRPS_{Alice}$. *Bob* also verifies that this content was indeed provided by *Alice* with her public signing key PKS_{Alice} . This authentication process is automated by the DApp server component interfacing the smart contract, as discussed in Section 5.2.5.

Digital signing ensures that a resource is indeed shared by the sender and is not tampered with. Likewise, encryption protects the information against

unauthorized access and spoofing. The data requestor’s access to a resource can be approved or revoked at any time via a state update in the smart contract by the data holder where all permissions are logged.

Role-based or attribute-based permissions can also be implemented off-chain in the same manner as in a traditional centralized system (*e.g.*, via Active Directory). In this case, a meta-cryptographic key pair would be created for each role or type of attribute and securely stored within the systems database. The system can then be configured so that only allows users meeting certain permission criteria to use the key for data access, while shielding users from unessential details.

5.2.4. Addressing Requirement 4: Consistent Data Formats

Context. Clinical research data can exist in various formats and structures, which may or may not be meaningful when shared with other providers from different organizations.

Problem. Blockchain-based health IT systems should facilitate data sharing, while adhering to some existing standard(s) for representing the clinical data. A key problem is thus how to design a blockchain-based architecture to enforce the application of existing clinical data standard(s).

Design choice → enforcing FHIR standards. HL7’s FHIR standards use JSON [53], which is a popular format for exchanging clinical information. JSON is more compact and readable compared to the XML format used by other data formatting standards, thereby enabling more efficient transmission of JSON-encoded data. It is also compatible with many software libraries and packages. As more health IT systems upgrade their data exchange protocols to comply to FHIR standards, FHIRChain enforces the use of FHIR to shared clinical data by validating whether the generated reference pointers follow the FHIR API standards [20].

5.2.5. Addressing Requirement 5: Maintaining Modularity

Context. Health IT system updates and/or upgrades are necessary to adopt more efficient, secure, or prevalent technology as it advances.

Problem. If functions in a smart contract have too many dependencies on the rest of a health IT system, then each upgrade to the system must deploy a new contract, which requires restoring data from previous versions to prevent loss. A key problem is thus how to design a modular data sharing system

that minimizes the need to create new versions of existing contracts when the system is upgraded. For example, when more user friendly features are needed, a good design should separate those updates from the underlying back-end services so that a change in the user interface does not require modifications of the server or blockchain component.

Design choice → *applying the model-view-controller (MVC) pattern.* The *MVC* pattern [50] separates a system into three components: (1) the *model*, which manages the behavior and data of a system and responds to requests for information about its state and instructions to change state, (2) the *view*, which manages the display of information, and (3) the *controller*, which interprets user inputs into appropriate messages to pass onto the *view* or *model*.

The FHIRChain architecture applies the *MVC* pattern to separate concerns with individually testable modules as follows: (1) a model in the form of an immutable *blockchain component* is used to store necessary meta data via smart contracts; (2) a view provides a front-end *user interface* that accepts user inputs and presents data; (3) a controller is a *server* component with control logic that facilitates interactions with data between the *user interface* and *blockchain component*, such as queries, updates, encrypting and decrypting contents; and (4) a controller-invoked *data connector* service is used to validate the implementation of FHIR standards and create reference pointers for the data sources upon requests from the server.

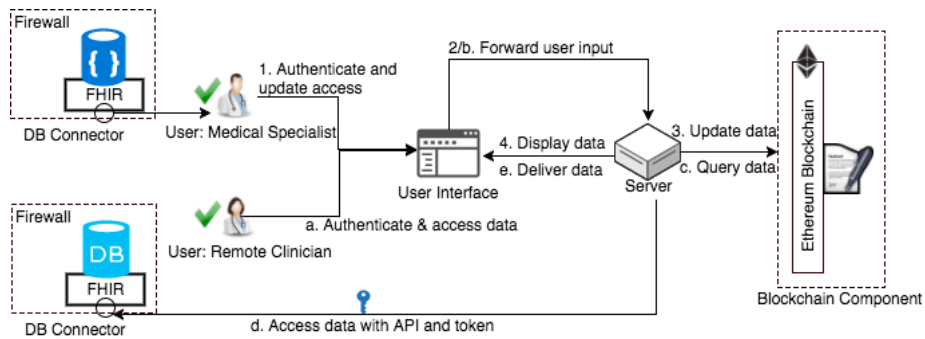


Figure 4: Composition and Structure of the FHIRChain Architecture with Modular Components.

The workflow for updating data access is shown in Figure 4 by the following steps 1-4:

1. A user first authenticates through the user interface (UI), and when successfully authenticated, data access permission request can be input to the system;
2. The UI forwards user’s request to the server;
3. The server logs permissioned or revoked access in the blockchain component (BC); and
4. The server updates UI with proper response to notify the user.

Likewise, the workflow for accessing a data source is outlined in the following steps a-e:

- a) The user first authenticates via the UI, and when successfully authenticated data access request can be input to the system;
- b) UI forwards users request to the server;
- c) The server queries BC for current user’s access token(s);
- d) When permission is valid, the server decodes the access token(s) with correct keys supplied by user and uses the decrypted reference pointer to obtain actual data from the DB connector to the proper database;
- e) When data has been retrieved from the data source via DB connector, the server updates UI to display data in a readable format.

FHIRChain stores all relevant information in smart contracts, decoupling data store from the rest of the system. This decoupling enables future upgrades to all other components without losing access to—or locking out—existing users or their permission information.

6. Case Study: Applying FHIRChain to Create a Prototype DApp

This section first describes the structure and functionality of a *decentralized app* (DApp) that customizes the FHIRChain architecture described in Section 5 to support collaborative clinical decision making via a remote tumor board case study. We then analyze the benefits and limitations of our DApp case study.

6.1. Overview of the FHIRChain DApp Case Study

The FHIRChain DApp is written in Javascript. It consists of ~1,000 lines of core app code that interacts with a private testnet of the Ethereum blockchain and three Solidity smart contracts, each containing ~50 lines of code. Our DApp customizes the FHIRChain architecture in a private

Ethereum testnet to address the various ONC requirements described in Section 4.

This DApp has an intuitive user interfacing portal that facilitates the sharing and viewing of patient cancer data for a remote tumor board to collaboratively create treatment plan for cancer patients. In addition, the DApp implements a notification service [27] that broadcasts events to appropriate event subscribers. The FHIRChain DApp notification service is used to alert collaborative tumor board members when new data access is available for review.

Verifying identity and authenticating participants with digital identities, as discussed in Section 5.2.1. Our DApp contains a *Registry* smart contract that maintains the digital health identities of providers who registered with our app. The registry maps provider email addresses (or phone numbers) from a public provider directory to both their public encryption (used as digital identity) and signing keys, which are generated automatically at user registration time. Figure 5 demonstrates the user registration and authentication workflow.

Storing and exchanging data securely with FHIR-based reference pointers, as discussed in Sections 5.2.2 and 5.2.4. Our DApp defines two cancer patient databases and referencing paths to patient data entries using the open-source HapiFHIR [54] public test server. Validation of the FHIR implementation is performed via regular expression parsing of the paths against the FHIR APIs [20].

Permissioning data access with token-based exchange, as discussed in Section 5.2.3. Our DApp also contains an *Access* smart contract that logs all user interactions and requests on the portal, *e.g.*, what resource is shared or no longer shared with which provider by whom and when. These access logs are structured as a mapping between user digital health identities (public encryption keys) and authorizations to custom-named access tokens (represented as a nested object associated with a *true/false* boolean value indicating if an access token access is granted for a provider). If an access revocation occurs, authorization is set to *false* and the associated token is set to an empty value. The workflow of this process is shown in Figure 6.

Maintaining modularity with the MVC pattern, as discussed in Section 5.2.5. The *view* component is a user interfacing portal that accepts provider user input, including registration and authentication credentials (corresponding keys) and data access information (*e.g.*, tumor board member email to query, a reference pointer to securely access data, and ap-




Figure 5: Workflow of the User Registration and Authentication Process in the FHIRChain DApp.

proval/revocation of access). Figure 7 is a screenshot of our DApp, presenting the following features (1) display recent sharing events related to the user, (2) display reference pointer APIs created by logged in user and available actions, and (3) display all references shared with logged in user and the option to view data.

The portal then forwards the user requests along with data input to the *server* component, where all the complex logic is encapsulated.

Our FHIRChain DApp *server* performs all functions and control logic, including verifying provider user email account, generating cryptographic keys, token creation via signing and encryption, token retrieval via decryption and signature verification, forwarding requests and delegating tasks between

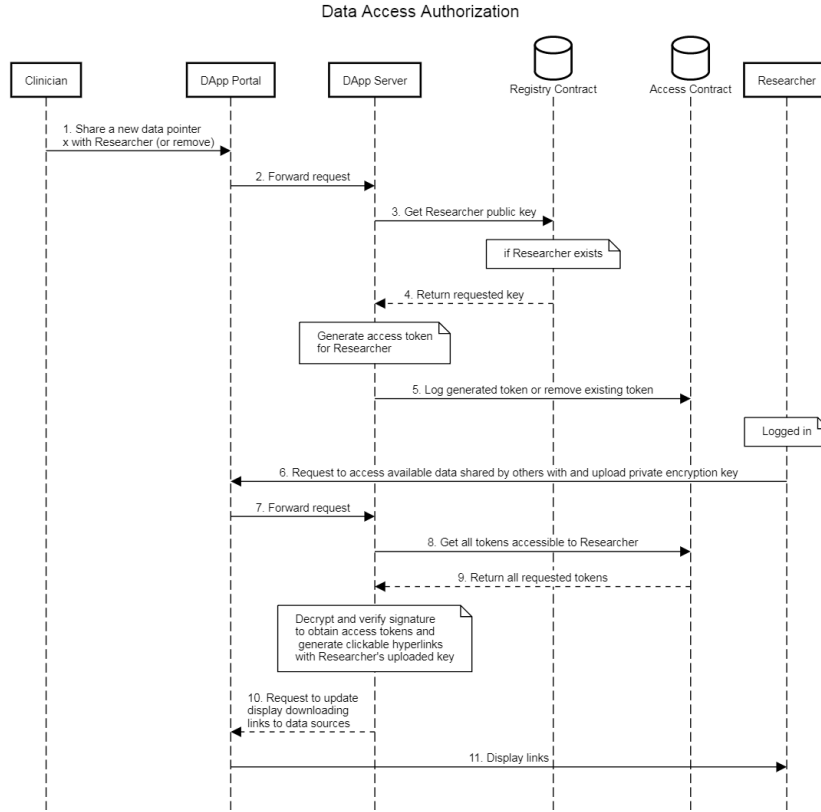


Figure 6: Workflow of Access Authorization in the FHIRChain DApp.

the *portal* and *blockchain*. The *blockchain* component is an independent *model* component containing two smart contracts for ubiquitous storing and persisting event logs of data access.

6.2. Benefits of Our FHIRChain DApp Case Study

Our FHIRChain Dapp case study achieved the following benefits:

- Increased modularity.** To increase modularity, we applied the “separation of concerns” principle [55] to decompose our DApp into independent components. FHIRChain employs a peer-to-peer API exchange protocol that references data pointers stored in a smart contract on the blockchain. In this design, exchanged information becomes lightweight, which increases scalability since system performance remains the same

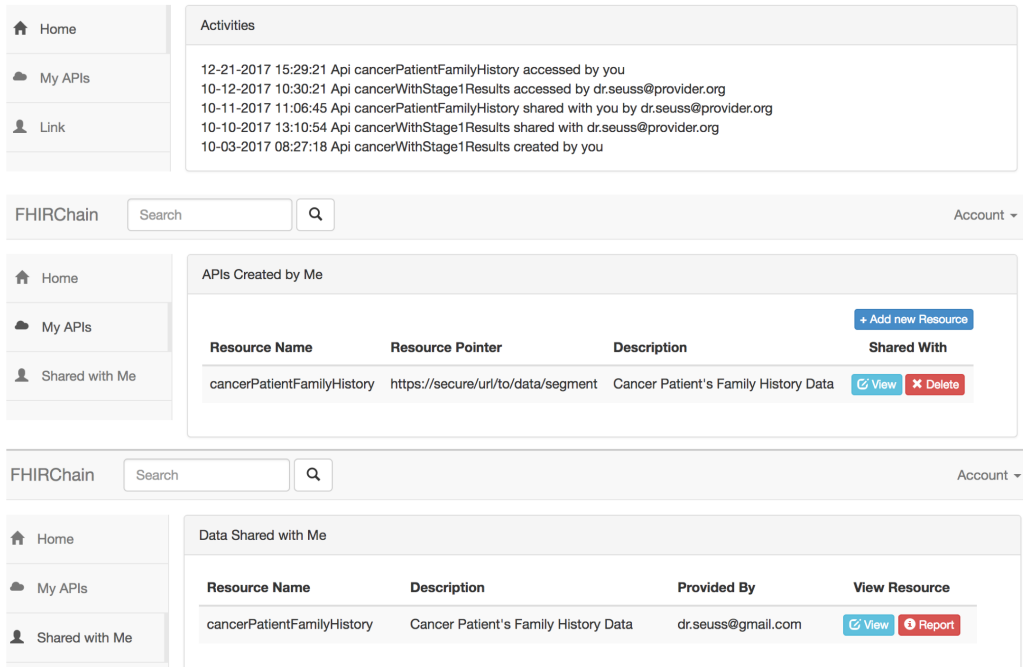


Figure 7: Screenshot of Our FHIRChain-based DApp User Interface.

regardless of the original size of the data. Likewise, data is not transmitted electronically across institutional boundaries, thereby reducing the risk of data being compromised.

- Scalable data integrity.** To ensure scalable data integrity, our design maintains a hash of the original data to exchange in addition to the reference pointer of the data. Suppose that the original data being exchanged is of size N and that the size of its reference pointer is ϵ . The total amount of data stored on-chain in terms of space complexity is then $O(\text{hash}(N) + \epsilon)$. Since the hashed output of a variable-length input can be a fixed value, it consumes a constant amount of space. The size of a data reference pointer would be scalably smaller than the actual data size. This design therefore enhances scalability by using constant-sized representations of the data, rather than using the actual data.
- Fine-grained access control.** To enable fine-grained access control, permissions to access a data source can be given or revoked at will by

providers across various institutions regardless of their trust relationships. By implementing the FHIR standards, more granular access can be granted to selected pieces of data rather than an entire document, which also increases data readability. Moreover, all events related to data sharing and data access are logged in a transparent history for auditability.

- **Enhanced trust.** The DApp applies public key cryptography, which enhances trust to participants in the following ways:
 - **Identifiability and authentication.** Given the computation power today, it is infeasible to impersonate a user without knowing their private key, and the only way a user can be authenticated to use our service is to provide the correct private key paired with their public key registered on the blockchain. On the other hand, it is trivial to create a new public/private key pair in case of a user’s private key being lost or stolen. This “digital identity” approach has been successfully adopted in Estonia’s government and healthcare infrastructure [56].
 - **Permission authorization.** With public key encryption securing their data reference pointers, users can trust that none other than the intended data recipient can view what they have shared. FHIRChain never shares the reference pointer with any user. Instead, RP is used to display the data content when it is decrypted with an authorized user’s private key. In addition, users can approve or revoke data access at any time, and the request takes effect immediately.

6.3. Limitations of Our FHIRChain DApp Case Study

Since our FHIRChain DApp was designed based on several assumptions it incurs the following limitations:

- **Does not address semantic interoperability.** FHIRChain cannot address data exchange challenges related to semantic interoperability that are not yet fully captured by the FHIR standards. To provide semantics to clinical data, therefore, manual inspection and mapping of predefined ontologies from medical and health data experts are required, which remain the focus of our future research in this space.

- **May not be compatible with legacy systems not supporting FHIR.** Many legacy systems may use other messaging standards, such as the more prevalent HL7 v2 standards [57], and do not support FHIR protocols. The goal of this paper, however, is to present the underlying representations and theories of our blockchain-based system. Although we advocate FHIR in the paper because it has been used quite frequently and it supports fine-grained data exchange, the principles behind the system described here can also be used with other standards like HL7 v2 [57].
- **Cannot control clinical malpractice.** The intended users of FHIR-Chain are clinicians interested in collaboratively providing clinical decision support for remote patients. Our current design trusts that the data being exchanged using our DApp is not abused, misused, or unethically redistributed by users. Our future work will explore options to minimize these risks, such as tracking data credibility using cryptographic hashing or zero knowledge proofs [58] (ability to demonstrate the truth of a statement without revealing additional information beyond what its trying to prove [59]) along with each reference pointer. Naturally, clinical malpractice may still occur (as in any other health IT system) since we cannot fully control these human behaviors.
- **DApp deployment costs.** Unlike existing public blockchain, such as Ethereum, our DApp is developed using a private testnet that imposes no interaction costs (*e.g.*, transaction fees). Our DApp would thus not be free of charge if deployed on a public blockchain. The convenience provided by a public blockchain, however, may justify the cost of usage versus the costs of licensing, running, and maintaining a private clinical data exchange infrastructure.

To overcome these limitations in future work, we will deploy our DApp in a permissioned consortium blockchain platform with trusted parties to ensure consensus through a variation of proof-of-work that incentivizes mining with cryptocurrency rewards. For instance, [35] proposes to use aggregated data as mining rewards in their system, while MultiChain [60] enforces a round-robin mining protocol in their blockchain. With the ability to replace monetary incentives to maintain consensus on the blockchain, the cost to use this blockchain-based service will be lower in the long run, although the initial deployment may still be expensive.

Although permissioned systems may be prone to collusion due to the 51% attack problem [29], the permissioned system used for healthcare would be maintained and managed by relatively large-scale entities/stakeholders within the healthcare industry. Unless majority of them (major hospitals, insurance companies, *etc.*) collude, therefore, the chance of experiencing this type of attack is quite low. Moreover, legal actions would most likely occur immediately upon the attack.

7. Concluding Remarks

This paper described the FHIRChain prototype we designed to provide patients with more collaborative clinical decision support using blockchain technology and the FHIR data standards. Complemented by the adoption of public key cryptography, our FHIRChain design addressed five key requirements provided by the ONC interoperability roadmap, including user identifiability and authentication, secure data exchange, permissioned data access, consistent data formats, and system modularity.

The following are the key lessons we learned from designing and implementing our DApp based on FHIRChain:

- **FHIRChain can provide trustless, decentralized storage for necessary meta information and audit logs.** FHIRChain alleviates proprietary vendor-lock found in conventional health IT systems by leveraging its blockchain component as a decentralized storage of necessary reference information as secure access points into those databases. It enables the sharing of clinical data without established trusts, providing clinicians with secure and scalable collaborative care decision support. In addition, each public key generated for a user is stored in the blockchain via a smart contract used to associate healthcare participants with their digital identities. Similarly, permission authorizations established between those participants are recorded in a smart contract as well, creating a traceable permission database with an audit log of data exchange history (*i.e.*, meta information involved during the data exchange and not the actual data). Storing these data on the blockchain ensures that our app is not subject to a single point of failure or corruption of records so that it is always accessible by healthcare participants.

- **FHIRChain facilitates data exchange without the need to upload/download data thus maintains data ownership.** The FHIR standards provide resource APIs to reference specific pieces of structured data while maintaining original data ownership. By adopting FHIR and combining it with blockchain technologies, FHIRChain creates lightweight reference pointers to siloed databases and exchange these pointers via the blockchain component instead of actual data. For telemedicine clinics or clinics in rural areas in particular, this approach can overcome network limitations by enabling scalable data sharing without requiring data to be uploaded to some other centralized repository, through which data can be shared and downloaded by other parties. In addition, this approach reduces risks of compromised data and ensures that original data ownership is respected. The reference pointers are encrypted with the intended recipients public key, *i.e.*, digital identity to permission data access. When successfully authenticated (*i.e.*, reference pointers are correctly decrypted) the data will be downloaded directly from the source and present properly formatted data to the user.
- **Public key cryptography can be effective for managing digital health identity in data sharing.** FHIRChain creates public keys as digital health identities associated with each collaborating care entity (provider or organization administrator). The benefits to this strategy include: (1) *easy authentication* since a clinician only needs to provide their private key associated with their identity, (2) *integrity* since by signing the exchanged reference pointers FHIRChain can easily verify that it was provided by the signed provider and has not been modified, and (3) *remedy to lost or stolen keys* since a new key can be created easily to replace the old key and associate with the same user. There is a drawback, however, to using digital identities for patients in a general clinical setting. Managing these identities—private keys—is hard because private keys are harder to remember than conventional passwords and require technical training for patients to manage their own keys. Nevertheless, there are approaches for managing private keys for larger populations, such as using key wallets [61, 25] or embedding private keys to physical medical ID cards [62].

In summary, our FHIRChain-based DApp demonstrates the potential of blockchain to foster effective healthcare data sharing while maintaining the

security of original data sources. FHIRChain can be further extended to address other healthcare interoperability issues, such as coordinating other stakeholders (*e.g.*, insurance companies) across the industry and providing patients with easier (and secure) access to their own medical records.

In our future work, we plan to refine the simulations for more rigorously evaluating the performance of our FHIRChain prototype. We will do so by deploying and comparing a number of different blockchain configurations in a testbed environment, such as using the blockchain template provided by Amazon Web Services [63]. Moreover, we will research techniques for identity management targeting the patient population.

Acknowledgements

The work presented in this paper was sponsored in part by funding from Varian Medical Systems.

References

- [1] M. Berman, A. Fenaughty, Technology and managed care: patient benefits of telemedicine in a rural health care network, *Health economics* 14 (2005) 559–573.
- [2] C. Castaneda, K. Nalley, C. Mannion, P. Bhattacharyya, P. Blake, A. Pecora, A. Goy, K. S. Suh, Clinical decision support systems for improving diagnostic accuracy and achieving precision medicine, *Journal of clinical bioinformatics* 5 (2015) 4.
- [3] H. Singh, T. D. Giardina, A. N. Meyer, S. N. Forjuoh, M. D. Reis, E. J. Thomas, Types and origins of diagnostic errors in primary care settings, *JAMA internal medicine* 173 (2013) 418–425.
- [4] R. Kaushal, K. G. Shojania, D. W. Bates, Effects of computerized physician order entry and clinical decision support systems on medication safety: a systematic review, *Archives of internal medicine* 163 (2003) 1409–1416.
- [5] G. D. Schiff, O. Hasan, S. Kim, R. Abrams, K. Cosby, B. L. Lambert, A. S. Elstein, S. Hasler, M. L. Kabongo, N. Krosnjar, et al., Diagnostic error in medicine: analysis of 583 physician-reported errors, *Archives of internal medicine* 169 (2009) 1881–1887.

- [6] D. B. Taichman, J. Backus, C. Baethge, H. Bauchner, P. W. De Leeuw, J. M. Drazen, J. Fletcher, F. A. Frizelle, T. Groves, A. Haileamlak, et al., Sharing clinical trial data: A proposal from the international committee of medical journal editors, *Annals of internal medicine* 164 (2016) 505–506.
- [7] E. Warren, Strengthening research through data sharing, *New England Journal of Medicine* 375 (2016) 401–403.
- [8] N. Geifman, J. Bollyky, S. Bhattacharya, A. J. Butte, Opening clinical trial data: are the voluntary data-sharing portals enough?, *BMC medicine* 13 (2015) 280.
- [9] G. E. Gross, The role of the tumor board in a community hospital, *CA: a cancer journal for clinicians* 37 (1987) 88–92.
- [10] J. Rieke, H. Bartelink, Telemedicine and its impact on cancer management, *European Journal of Cancer* 36 (2000) 826–833.
- [11] C. L. Marshall, N. J. Petersen, A. D. Naik, N. V. Velde, A. Artinyan, D. Albo, D. H. Berger, D. A. Anaya, Implementation of a regional virtual tumor board: a prospective study evaluating feasibility and provider acceptance, *Telemedicine and e-Health* 20 (2014) 705–711.
- [12] L. Levit, A. P. Smith, E. J. Benz Jr, B. Ferrell, Ensuring quality cancer care through the oncology workforce, *Journal of Oncology Practice* 6 (2010) 7–11.
- [13] M. Terry, Medical identity theft and telemedicine security, *Telemedicine and e-Health* 15 (2009) 1–5.
- [14] A. S. Downey, S. Olson, et al., Sharing clinical research data: workshop summary, National Academies Press, 2013.
- [15] G. Hripcsak, M. Bloomrosen, P. FlatelyBrennan, C. G. Chute, J. Cimino, D. E. Detmer, M. Edmunds, P. J. Embi, M. M. Goldstein, W. E. Hammond, et al., Health data use, stewardship, and governance: ongoing gaps and challenges: a report from amia’s 2012 health policy meeting, *Journal of the American Medical Informatics Association* 21 (2014) 204–211.
- [16] G. Hartvigsen, M. A. Johansen, P. Hasvold, J. G. Bellika, E. Arsand, E. Arild, D. Gammon, S. Pettersen, S. Pedersen, et al., Challenges in telemedicine and ehealth: lessons learned from 20 years with telemedicine in tromso, *Studies in health technology and informatics* 129 (2007) 82.

- [17] M. Maheu, P. Whitten, A. Allen, *E-Health, Telehealth, and Telemedicine: a guide to startup and success*, John Wiley & Sons, 2002.
- [18] R. LaRose, S. Stover, J. L. Gregg, J. Straubhaar, The impact of rural broadband development: Lessons from a natural field experiment, *Government Information Quarterly* 28 (2011) 91–100.
- [19] A. B. Bondi, Characteristics of scalability and their impact on performance, in: *Proceedings of the 2nd international workshop on Software and performance*, ACM, pp. 195–203.
- [20] D. Bender, K. Sartipi, H17 fhir: An agile and restful approach to healthcare information exchange, in: *Computer-Based Medical Systems (CBMS), 2013 IEEE 26th International Symposium on*, IEEE, pp. 326–331.
- [21] R. L. Richesson, J. Krischer, Data standards in clinical research: gaps, overlaps, challenges and future directions, *Journal of the American Medical Informatics Association* 14 (2007) 687–696.
- [22] R. Das, Does blockchain have a place in healthcare, *Forbes*. <https://www.forbes.com/sites/reenitadas/2017/05/08/does-blockchain-have-a-place-in-healthcare> (2017).
- [23] M. Mettler, Blockchain technology in healthcare: The revolution starts here, in: *e-Health Networking, Applications and Services (Healthcom), 2016 IEEE 18th International Conference on*, IEEE, pp. 1–3.
- [24] A. Azaria, A. Ekblaw, T. Vieira, A. Lippman, Medrec: Using blockchain for medical data access and permission management, in: *Open and Big Data (OBD), International Conference on*, IEEE, pp. 25–30.
- [25] S. Nakamoto, *Bitcoin: A peer-to-peer electronic cash system*, 2008.
- [26] P. Zhang, M. A. Walker, J. White, D. C. Schmidt, G. Lenz, Metrics for assessing blockchain-based healthcare decentralized apps, in: *2017 IEEE 19th International Conference on e-Health Networking, Applications and Services (Healthcom)*, pp. 1–4.
- [27] P. Zhang, J. White, D. C. Schmidt, G. Lenz, Applying software patterns to address interoperability in blockchain-based healthcare apps, the 24th Pattern Languages of Programming conference (2017).
- [28] K. DeSalvo, E. Galvez, Connecting health and care for the nation: a shared nationwide interoperability roadmap version 1.0, *Health IT Buzz* (2015).

- [29] V. Buterin, et al., Ethereum white paper, 2013.
- [30] D. Johnston, S. O. Yilmaz, J. Kandah, N. Bentenitis, F. Hashemi, R. Gross, S. Wilkinson, S. Mason, The general theory of decentralized applications, dapps, GitHub, June 9 (2014).
- [31] R. Krawiec, D. Housman, M. White, M. Filipova, F. Quarre, D. Barr, A. Nesbitt, K. Fedosova, J. Killmeyer, A. Israel, et al., Blockchain: Opportunities for health care, in: Proc. NIST Workshop Blockchain Healthcare, pp. 1–16.
- [32] J. B. Peter B. Nichol, Co-creation of trust for healthcare: The cryptocitizen. framework for interoperability with blockchain (2016).
- [33] I. G. B. S. P. S. Team, Blockchain: The chain of trust and its potential to transform healthcare our point of view (2016).
- [34] T.-T. Kuo, H.-E. Kim, L. Ohno-Machado, Blockchain distributed ledger technologies for biomedical and health care applications, Journal of the American Medical Informatics Association 24 (2017) 1211–1220.
- [35] A. Ekblaw, A. Azaria, J. D. Halamka, A. Lippman, A case study for blockchain in healthcare:medrec prototype for electronic health records and medical research data, in: Proceedings of IEEE Open & Big Data Conference.
- [36] K. Peterson, R. Deeduvanu, P. Kanjamala, K. Boles, A blockchain-based approach to health information exchange networks, 2016.
- [37] A. Dubovitskaya, Z. Xu, S. Ryu, M. Schumacher, F. Wang, Secure and trustable electronic medical records sharing using blockchain, arXiv preprint arXiv:1709.06528 (2017).
- [38] A. Gropper, Powering the physician-patient relationship with hie of one blockchain health it, 2016.
- [39] Direct project, Available at <https://www.healthit.gov/policy-researchers-implementers/direct-project>, ????
- [40] G. Greenspan, Blockchains vs centralized databases, Available at <https://www.multichain.com/blog/2016/03/blockchains-vs-centralized-databases>, Accessed 2017-12-31.
- [41] C. for Disease Control, Prevention, et al., Hipaa privacy rule and public health. guidance from cdc and the us department of health and human services, MMWR: Morbidity and mortality weekly report 52 (2003) 1–17.

- [42] A. J. Menezes, P. C. Van Oorschot, S. A. Vanstone, Handbook of applied cryptography, CRC press, 1996.
- [43] A. Al Omar, M. S. Rahman, A. Basu, S. Kiyomoto, Medibchain: A blockchain based privacy preserving platform for healthcare data, in: International Conference on Security, Privacy and Anonymity in Computation, Communication and Storage, Springer, pp. 534–543.
- [44] X. Yue, H. Wang, D. Jin, M. Li, W. Jiang, Healthcare data gateways: found healthcare intelligence on blockchain with novel privacy risk control, Journal of medical systems 40 (2016) 218.
- [45] R. Shea, Simple contracts are better contracts: What we can learn from the meltdown of the dao, Available at <https://medium.com/@ryanshea/simple-contracts-are-better-contracts-what-we-can-learn-from-the-dao-6293214bad3a>, Accessed 2017-12-31.
- [46] D. Hardt, The oauth 2.0 authorization framework (2012).
- [47] Introduction to the isa, Available at <https://www.healthit.gov/isa/>, Accessed 2017-12-31.
- [48] M. Kim Futrell, Structured data (2013).
- [49] K. Kawamoto, T. Hongsermeier, A. Wright, J. Lewis, D. S. Bell, B. Middleton, Key principles for a national clinical decision support knowledge sharing framework: synthesis of insights from leading subject matter experts, Journal of the American Medical Informatics Association 20 (2013) 199–207.
- [50] A. Leff, J. T. Rayfield, Web-application development using the model/view/controller design pattern, in: Enterprise Distributed Object Computing Conference, 2001. EDOC'01. Proceedings. Fifth IEEE International, IEEE, pp. 118–127.
- [51] T. Ylonen, C. Lonvick, The secure shell (ssh) protocol architecture (2006).
- [52] H. Krawczyk, The order of encryption and authentication for protecting communications (or: How secure is ssl?), in: Advances in CryptologyCRYPTO 2001, Springer, pp. 310–331.
- [53] D. Crockford, The application/json media type for javascript object notation (json) (2006).
- [54] Hapi-fhir, Available at <http://fhirtest.uhn.ca/>, Accessed 2017-12-31.

- [55] H. Ossher, P. Tarr, Using multidimensional separation of concerns to (re) shape evolving software, *Communications of the ACM* 44 (2001) 43–50.
- [56] R. M. Alvarez, T. E. Hall, A. H. Trechsel, Internet voting in comparative perspective: the case of estonia, *PS: Political Science & Politics* 42 (2009) 497–505.
- [57] R. H. Dolin, L. Alschuler, S. Boyer, C. Beebe, F. M. Behlen, P. V. Biron, A. Shabo, H17 clinical document architecture, release 2, *Journal of the American Medical Informatics Association* 13 (2006) 30–39.
- [58] C. Rackoff, D. R. Simon, Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack, in: *Annual International Cryptology Conference*, Springer, pp. 433–444.
- [59] G. Greenspan, Understanding zero knowledge blockchains, Available at <https://www.multichain.com/blog/2016/11/understanding-zero-knowledge-blockchains/>, Accessed 2017-12-31.
- [60] G. Greenspan, Multichain private blockchain white paper, Available at <https://www.multichain.com/download/MultiChain-White-Paper.pdf>, 2015.
- [61] S. Even, O. Goldreich, Y. Yacobi, Electronic wallet, in: *Advances in Cryptology*, Springer, pp. 383–386.
- [62] G. Anthes, Estonia: a model for e-government, *Communications of the ACM* 58 (2015) 18–20.
- [63] Aws blockchain templates, <https://aws.amazon.com/blockchain/templates/>, ????