Developing a Structural Standard for Smart Contract Electronic Health Records Based on the HL7 Fast Healthcare Interoperability Resources

Christian Pulmano[a] and Proceso Fernandez[b]
Department of Information Systems and Computer Science, Ateneo de Manila University, Quezon City, Philippines

Keywords: Electronic Health Records, HL7 FHIR, Blockchain, Smart Contracts, Solidity.

Abstract: Electronic health records (EHRs) are becoming more essential to patient care, as they provide information that is usually valuable for health and medical decision-making. Each EHR typically comprises validated data provided by medical professionals during patient care and personal health data recorded by the individual patient. Existing implementations of EHRs, however, are often siloed and managed by disparate organizations. Health information exchanges, which aim to interoperate EHRs by different health organizations, are also primarily centralized and can be vulnerable to attacks such as distributed denial of service (DDoS) and data breaches. Blockchain is a potential solution to solving some of the current issues in health information exchange implementations. Smart contracts deployed on blockchain networks can help enforce appropriate health and medical record-keeping and sharing standards. This study aims to create a blockchain-based implementation that allows decentralized EHRs. Smart contracts were developed based on the Health Level 7 Fast Healthcare Interoperability Resource (HL7 FHIR). Experiments were designed and simulated on a local Ethereum blockchain network where 10,000 (generated) patients' records were processed for evaluation. The results show that a structural standard can be applied to EHR smart contracts, but while blockchain solutions for EHR systems may be expected to be less vulnerable than centralized systems, the gas costs in Ethereum are potentially prohibitive and should be carefully considered.

1 INTRODUCTION

Electronic Health Records (EHRs) are important in healthcare. An EHR is a collection of an individual's health records gathered from various sources, including data encoded by medical professionals during consultations (Mahajan et al., 2023) and health and wellness data gathered from personal health tracking. They contain information vital to health and medical decision-making (Arbabi et al., 2022). Therefore, aside from storing EHRs, it is also important to consider the standards and format for which they are shared with relevant decision-makers.

Unfortunately, EHRs are often managed by disparate service providers and stored across multiple facilities in unstandardized formats (Cerchione et al., 2023). Data sharing can be challenging without proper standards, making data inaccessible when needed (Chelladurai and Pandian, 2022). Structural standards define a structure by which data, including the data fields and formats, can be shared across disparate health organizations. One standard that can be considered is the Health Level 7 Fast Healthcare Interoperability Resources (HL7 FHIR), a set of resources that provides health data standards for storing and exchanging health and health-related data (HL7.org, 2023a).

Even if such standards for data storage and exchange are followed, however, there are still other important issues related to EHRs that have to be addressed. For example, typical implementations of EHRs use centralized technologies and are vulnerable to attacks such as distributed denial of service (DDoS), malware, and hacking. If not properly implemented, centralized technologies can also be prone to data tampering. Blockchain, an emerging technology whose strengths include decentralization, immutability, and security, presents a promising solution to some of the current issues in centralized EHR implementations.

A blockchain is a distributed ledger composed of a growing list of connected blocks arranged in chronological order and securely linked together using cryptographic hashes, where a block contains a set of

---

[a] https://orcid.org/0000-0001-7870-8197
[b] https://orcid.org/0000-0001-5370-4544

Pulmano, C. and Fernandez, P.
Developing a Structural Standard for Smart Contract Electronic Health Records Based on the HL7 Fast Healthcare Interoperability Resources.
DOI: 10.5220/0012624100003699

In *Proceedings of the 10th International Conference on Information and Communication Technologies for Ageing Well and e-Health (ICT4AWE 2024)*, pages 169-175
ISBN: 978-989-758-700-9; ISSN: 2184-4984

confirmed transactions (Arbabi et al., 2022; Villarreal et al., 2023). In the healthcare context, transactions can refer to activities like creating a medical record, diagnosing a patient, prescribing a medicine, etc. The ability of blockchain systems to maintain an immutable data ledger can benefit EHRs in terms of data integrity and reliability. Furthermore, smart contracts in a blockchain can enable individuals to manage their EHRs without going through their service providers.

Smart contracts are computer programs that can be deployed to blockchain networks and are considered digital counterparts of traditional paper-based contracts (Arbabi et al., 2022). A key property of a smart contract is that it is self-executing and thus can automate the actions required in an agreement or contract. Vitalik Buterin first published a white paper for Ethereum in 2014, proposing a blockchain with a built-in Turing-complete programming language for developing smart contracts (Buterin et al., 2014). Solidity (Solidity Team, 2023a) is one of the programming languages that can be used to write smart contracts in an Ethereum blockchain. This allows developers to build applications that support different functions, including healthcare, and deploy contracts to the Ethereum blockchain.

Smart contract functions on Ethereum blockchains are executed on the (blockchain) miners' Ethereum Virtual Machine (EVM), thus requiring Transaction initiators to pay the miners what is referred to as *gas fees* as compensation for the used computing resources. These miners are responsible for validating the transactions, creating new blocks and cryptographically linking a newly created block to the blockchain. Miners essentially secure the blockchain network and promote decentralization, and they are incentivized to do so by earning ETH, which is the token for value exchange in the Ethereum network.

The gas fees required for deploying or invoking transactions in smart contracts can also depend on the code design (Li et al., 2023). Thus, software developers have to carefully design the smart contract implementation if they wish to minimize the gas fees per transaction.

In this paper, we partially implemented, as proof of concept, a proposed structural standard for blockchain-based EHRs using smart contracts based on the Health Level 7 Fast Healthcare Interoperability Resource (HL7 FHIR). The smart contract is envisioned to become the technical standard for implementing EHRs in blockchain. Our implementation, written in Solidity, enabled us to estimate average gas fees. Other considerations, such as privacy and authentication, were excluded in this iteration of the study,

The remaining sections of this paper are as follows: Section 2 provides a brief overview of related works, Section 3 outlines the methodology for the study, Section 4 presents the results and discussions, and Section 5 concludes the study.

2 RELATED WORKS

This section provides an overview of related works about blockchain for EHRs and the evaluation of gas costs based on smart contract coding patterns.

2.1 Blockchain for Electronic Health Records

Many efforts have already been made to implement blockchain for EHRs. Blockchain is the underlying technology for Bitcoin, and since its conception, it has been applied to different domains, including health (Frizzo-Barker et al., 2020). One study reviewed a good number of published works from 2016 to 2020 to determine the extent of applying blockchain technology for managing EHRs (Al Mamun et al., 2022). The study chose several publications based on privacy, security, storage scalability, accessibility, and cost analysis. It found that much research was still in the conceptual stage or at most early prototype, and only a small chunk was already in the implementation stage. However, multiple studies were able to consider implementing health standards such as FHIR and HL7 in their design (Zhang et al., 2018; Roehrs et al., 2017; Donawa et al., 2019). More recent publications are also available where smart contracts and artificial intelligence with blockchain are proposed (Mahajan et al., 2023; Haddad et al., 2022), and functionalities for access-based controls are implemented (De Oliveira et al., 2022).

2.2 Evaluating Smart Contract Gas Costs

Multiple studies (Di Sorbo et al., 2022; Li et al., 2023; Zhao et al., 2023) focused on identifying expensive gas patterns in Solidity smart contract source codes. In one study, the researchers identified 19 common gas-expensive patterns or code smells and proposed code recommendations to reduce gas costs. They also developed GasMet, which defines a suite of metrics that can help evaluate the relationships of code patterns to their gas consumption (Di Sorbo et al., 2022). Another study identified gas-expensive patterns and categorized them into storage-related patterns, judgment-related patterns, and loop-related patterns. They proposed GaSaver, a tool that automatically detects gas-expensive patterns in a Solidity smart contract (Zhao et al., 2023).

While previous works have already suggested the feasibility of using blockchain for EHRs, with some of these also providing strategies for integrating health data standards like HL7 FHIR, often these studies combined on-chain and off-chain approaches. On-chain means that data is stored in a node (machine) that is part of the blockchain network, where data is regularly verified since data is replicated across multiple nodes in the blockchain network. On the other hand, off-chain means that data is stored in a machine outside the blockchain network and is not included in block confirmation and verification (Miyachi and Mackey, 2021). Off-chain implementations are meant to increase the scalability of blockchain solutions, but they are more vulnerable to attacks, unlike on-chain implementations (Al Mamun et al., 2022). Our study investigates the feasibility of a structural standard for (full) on-chain EHR smart contract implementations on the Ethereum blockchain and provides insights on gas cost implications for such implementations.

3 METHODOLOGY

The study's main objectives are mainly two-fold: (1) develop a smart contract that supports a (full) on-chain EHR implementation based on HL7 FHIR and (2) evaluate its feasibility in terms of gas costs. This section outlines the methodology used to achieve the objective.

3.1 Dataset Generation

A synthetic dataset containing 10,000 patients' personal information records was generated using the Python Faker (Ethicalads, 2023) package. This package has tools that help generate names, birth dates, addresses, and contact data. The number of records was decided based on the amount that could be processed given time limitations but could still provide good record variations for estimating averages later on. A Python script was written to generate data based on the data dictionary of the FHIR Patient resource. The UML diagram in Figure 1 shows the data involved under the resource. It should be noted that the generated dataset contained only general personal data and excluded (generated) historical medical transaction records for the time being since smart contracts for other FHIR resources also need to be designed based on their corresponding data fields.

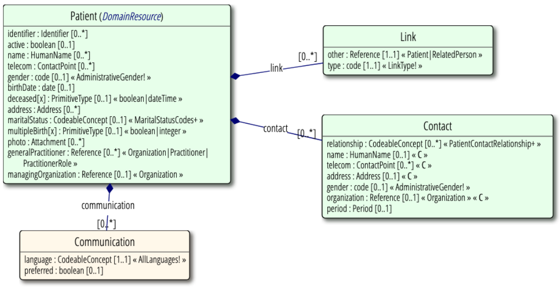

Figure 1: UML diagram of the FHIR Patient Resource Content (HL7.org, 2023b).

3.2 Experiment Designs, Simulation, and Evaluation

The experiments aimed to deploy, for evaluation, some EHR systems that store patient personal records based on the Patient (HL7.org, 2023b) resource of HL7-FHIR. Two main models were designed and implemented for the simulations. The first model implements a hybrid off-chain and on-chain approach, while the second model implements a full on-chain approach. For both models, smart contracts were developed using the Solidity (Solidity Team, 2023a) programming language and deployed to a blockchain network that was setup on a local server using Ganache (Truffle, 2023). A summary of the available functions for each smart contract is shown in Table 1. The models are also illustrated in Figures 2 and 3. The experiments are designed to help assess the gas costs involved in digitalizing patient personal information records, which is one of the first steps in creating patient EHRs, using smart contracts in the Ethereum blockchain.

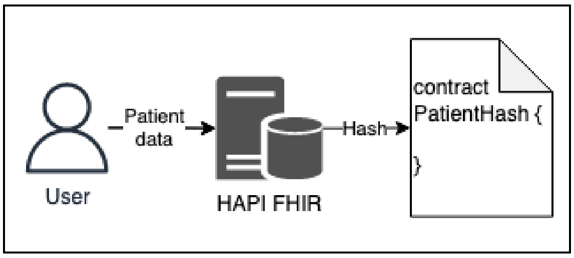

```
contract
PatientHash {

}
```

Figure 2: Illustration of model 1 components and workflow.

Following common implementations, a local HAPI FHIR server (HAPI FHIR, 2023) was set up for the first model to store the FHIR-based patient records. HAPI-FHIR is an open-source project based

Table 1: Smart contract functions.

	Model	Contract Functions
1	HAPI FHIR w/ PatientHash smart contract	constructor
2a	Patient smart contract	constructor, addIdentifier, setActive, addName, addTelecom, setGender, setBirthDate, setDeceasedInfo, addAddress, setMaritalStatus, setMultipleBirthtInfo, addContact, addCommunication, setGeneralPractitioner, setManagingOrganization, addLink
2b	PatientV2 smart contract	constructor, addContact, setGeneralPractitioner, setManagingOrganization, addLink

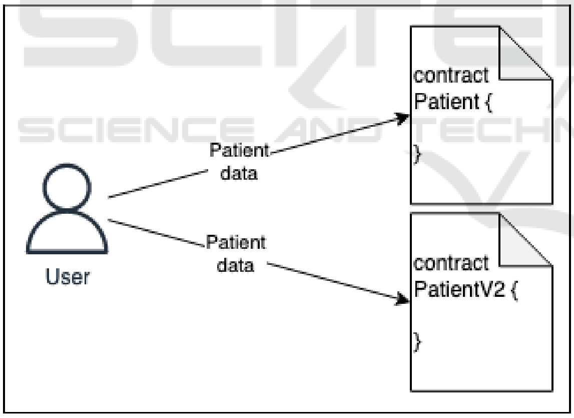

Figure 3: Illustration of model 2 components and workflow.

on Java that developers can use to provide FHIR-based API resources. In this study, a centralized HAPI-FHIR server was deployed to convert personal information records to the FHIR Patient resource data format. The hash value of the Patient resource record was then computed and submitted as a parameter to the constructor function of the smart contract deployed in a local Ethereum blockchain network. The hash value of the record in the blockchain network can be used to verify if the record has been tampered with since altering the patient record should yield a different hash value. The actual patient records are stored in a centralized database system in this model.

For the second model, a fully on-chain approach was developed. Under this approach, personal data is directly submitted to a new smart contract deployed on the local Ethereum blockchain network. The smart contract accepts data such as name, birthdate, address, contact information, and other needed information in the HL7 FHIR Patient resource. Two versions were developed and evaluated for this model. In the first version, the constructor function does not accept any parameters, and data can only be submitted using the specific functions in the contract developed to handle a particular piece of data input. The functions handle data attributes defined in the Patient resource of HL7 FHIR.

The second iteration of the second model is an optimized version where the constructor function already accepts parameters and stores patient information in the contract. Due to limitations in the stack size of the EVM, not all data attributes were included as parameters of the constructor function; only data about the patients themselves were included. Functions that process data external to the patients, i.e., general practitioner, managing organization, and other links, were retained as separate functions. Based on recommendations from previous studies, the second version also uses bytes data type instead of string for handling text data (Di Sorbo et al., 2022; Li et al., 2023; Zhao et al., 2023).

A Python script was developed to simulate the experiments for the two models where all models used the same dataset. The Web3 package (Ethereum Foundation, 2023) was used to invoke transactions from the Python script to smart contracts in the Ethereum blockchain network. All programs used for the experiments were installed on a machine with a macOS Sonoma 14.0 operating system, 8 GB memory, and an Apple M1 processor. The gas costs for all invoked transactions were recorded during the experiment runs. The gas costs for the different models were then compared and evaluated.

4 RESULTS AND DISCUSSIONS

Table 2 summarizes the total gas cost and average gas cost per record for the smart contracts. While it is expected that the second model would require higher gas costs than the first model, the purpose of the experiments was to evaluate the magnitude of the difference. As shown in Table 3, the total gas cost to process 10,000 records for Model 1 is above 1.8 billion gas units. Model 2a incurred the highest gas cost at around 44.1 billion gas units, while

Table 2: Total gas costs and estimated gas fees for each model as of 2024-01-03T11:00:00Z.

	Model 1 HAPI FHIR w/ PatientHash	Model 2a Patient	Model 2b PatientV2
Records	10,000	10,000	10,000
Gas price (Gwei)	20	20	20
Gas price (USD)	0.000002367	0.000002367	0.000002367
Total gas cost	1,872,025,468	44,099,957,978	26,143,190,020
Gas cost/record	187,203	4,409,996	2,614,319
Gas fee/record (Gwei/record)	3,744,051	88,199,916	52,286,380
Gas fee/record (USD/record)	9	209	124

Table 3: Total gas costs for each smart contract function.

Function	Model 1 HAPI FHIR w/ PatientHash	Model 2a Patient	Model 2b PatientV2
Contract deployment	1,872,025,468	18,878,090,000	24,386,657,948
addIdentifier		2,077,611,622	
setActive		338,402,376	
addName		3,007,115,284	
addTelecom		1,766,766,264	
setGender		387,823,784	
setBirthDate		451,550,000	
setDeceasedInfo		475,833,369	
addAddress		4,056,553,516	
setMaritalStatus		450,645,024	
setMultipleBirthInfo		543,069,056	
addContact		9,098,202,123	
addCommunication		809,973,488	
setGeneralPractitioner		452,386,380	451,496,380
setManagingOrganization		451,947,088	451,277,088
addLink		853,988,604	853,758,604
Total	1,872,025,468	44,099,957,978	26,143,190,020

Table 4: Average gas costs and fees for each smart contract function.

Function	Model 1 HAPI FHIR w/ PatientHash	Model 2a Patient	Model 2b PatientV2
Contract deployment	187,203	1,887,809	2,438,666
addIdentifier		207,761	
setActive		33,840	
addName		300,712	
addTelecom		176,677	
setGender		38,782	
setBirthDate		45,155	
setDeceasedInfo		47,583	
addAddress		405,655	
setMaritalStatus		45,065	
setMultipleBirthInfo		54,307	
addContact		909,820	
addCommunication		80,997	
setGeneralPractitioner		45,239	45,150
setManagingOrganization		45,195	45,128
addLink		85,399	85,376
Total	187,203	4,409,996	2,614,319

the updated smart contract for the second model (2b) only required around 26.1 billion gas units. The gas units required to process a patient record are approximately 187 thousand, 4.4 million, and 2.6 million for models 1, 2a, and 2b, respectively. As of 2024-01-03T11:00:00Z, the average gas price in the Ethereum main net is 20 Gwei. Assuming the smart contracts will be deployed in the main network, the gas fees per record would be approximately 3.7 million Gwei (9 USD), 88.1 million Gwei (209 USD), and 52.3 million Gwei (124 USD) for Models 1, 2a, and 2b, respectively.

Tables 3 and 4 show the total and average gas costs per record for each function available in the smart contracts. A contract deployed to the blockchain network corresponds to one patient record. Model 1 required the lowest gas cost for deploying the smart contracts for 10,000 records at 1.8 billion gas units since it has minimal features and required inputs. Model 2a required around 18.9 billion gas units for contract deployment, which is lower than the value for Model 2b at around 24.4 billion gas units. Model 2b has a higher value since its constructor already accepts multiple parameters, unlike Model 2a, which doesn't accept any parameter. It can also be observed from the table that the more gas-expensive functions are those requiring a struct (Solidity Team, 2023b) type parameter such as addIdentifier(), addName(), addAddress(), etc.

The results show that the first model is the least gas-expensive strategy. However, there are also other considerations to implementing a specific model. In terms of decentralization, the EHRs in the first model are still stored in a centralized HAPI FHIR server, which is vulnerable to attacks and hacking. The hash values recorded on-chain help with the verification of the records, but additional safeguards must still be implemented to ensure the centralized server's security and data integrity. In the second model, EHRs are stored and validated by a decentralized network of nodes.

The two versions in the second model show that smart contract patterns should also be considered. This is evident in the difference in gas cost values between the first and second versions. After reorganizing how data fields are encoded to the smart contract, the total gas costs also changed significantly, despite processing the same dataset.

It is also important to consider that a full on-chain HL7 FHIR implementation will require including other resources besides the Patient resource, including Observations, Encounters, Procedures, etc. The inclusion of other resources will also require additional gas costs. Implementing the Patient resource alone already incurred high values. Therefore, selecting only specific data points that benefit from the on-chain implementation might be more appropriate.

5 CONCLUSIONS AND RECOMMENDATIONS

The study provides insights into an on-chain implementation of EHRs using smart contracts based on the HL7 FHIR interoperability standard. Two models were designed in the experiments, where a hybrid off-chain-on-chain implementation was compared against a full on-chain implementation. A synthetic patient dataset was generated and processed for evaluation.

The study concludes that gas costs for deploying HL7 FHIR-based smart contacts should be considered. The high gas costs for implementing the Patient resource alone suggest that not all the data can be implemented on-chain. Instead, HL7 FHIR resources and data attributes that will benefit highly from on-chain implementation should be carefully selected. The design pattern at which smart contracts are developed is also relevant, considering that implementing different patterns can also affect the gas requirements of the contract.

Future work can look more closely into the gas mechanics of the Ethereum blockchain to have more insights into how gas costs are computed based on smart contract patterns. Additional experiments can also be conducted to evaluate the implementation in other Layer 1 blockchains, or possibly even Layer 2 blockchain networks over the Ethereum or other Layer 1 blockchains.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the Department of Information Systems and Computer Science and the Ateneo de Manila University for their support of the study.

REFERENCES

Al Mamun, A., Azam, S., and Gritti, C. (2022). Blockchain-based electronic health records management: a comprehensive review and future research direction. *IEEE Access*, 10:5768–5789.

Arbabi, M. S., Lal, C., Veeraragavan, N. R., Marijan, D., Nygård, J. F., and Vitenberg, R. (2022). A survey on blockchain for healthcare: Challenges, benefits, and future directions. *IEEE Communications Surveys & Tutorials*.

Buterin, V. et al. (2014). A next-generation smart contract and decentralized application platform. *white paper*, 3(37):2–1.

Cerchione, R., Centobelli, P., Riccio, E., Abbate, S., and Oropallo, E. (2023). Blockchain's coming to hospital to digitalize healthcare services: Designing a distributed electronic health record ecosystem. *Technovation*, 120:102480.

Chelladurai, U. and Pandian, S. (2022). A novel blockchain based electronic health record automation system for healthcare. *Journal of Ambient Intelligence and Humanized Computing*, pages 1–11.

De Oliveira, M. T., Reis, L. H. A., Verginadis, Y., Mattos, D. M. F., and Olabarriaga, S. D. (2022). Smartaccess: Attribute-based access control system for medical records based on smart contracts. *IEEE Access*, 10:117836–117854.

Di Sorbo, A., Laudanna, S., Vacca, A., Visaggio, C. A., and Canfora, G. (2022). Profiling gas consumption in solidity smart contracts. *Journal of Systems and Software*, 186:111193.

Donawa, A., Orukari, I., and Baker, C. E. (2019). Scaling blockchains to support electronic health records for hospital systems. In *2019 IEEE 10th Annual Ubiquitous Computing, Electronics & Mobile Communication Conference (UEMCON)*, pages 0550–0556. IEEE.

Ethereum Foundation (2023). gm — web3.py 6.11.4 documentation.

Ethicalads (2023). Welcome to faker's documentation! — faker 18.13.0 documentation.

Frizzo-Barker, J., Chow-White, P. A., Adams, P. R., Mentanko, J., Ha, D., and Green, S. (2020). Blockchain as a disruptive technology for business: A systematic review. *International Journal of Information Management*, 51:102029.

Haddad, A., Habaebi, M. H., Islam, M. R., Hasbullah, N. F., and Zabidi, S. A. (2022). Systematic review on ai-blockchain based e-healthcare records management systems. *IEEE Access*, 10:94583–94615.

HAPI FHIR (2023). hapifhir/hapi-fhir: Hapi fhir - java api for hl7 fhir clients and servers.

HL7.org (2023a). Fhir v5.0.0.

HL7.org (2023b). Patient - fhir v5.0.0.

Li, J., Zhao, Z., Su, Z., and Meng, W. (2023). Gas-expensive patterns detection to optimize smart contracts. *Applied Soft Computing*, 145:110542.

Mahajan, H. B., Rashid, A. S., Junnarkar, A. A., Uke, N., Deshpande, S. D., Futane, P. R., Alkhayyat, A., and Alhayani, B. (2023). Integration of healthcare 4.0 and blockchain into secure cloud-based electronic health records systems. *Applied Nanoscience*, 13(3):2329–2342.

Miyachi, K. and Mackey, T. K. (2021). hocbs: A privacy-preserving blockchain framework for healthcare data leveraging an on-chain and off-chain system design. *Information Processing & Management*, 58(3):102535.

Roehrs, A., Da Costa, C. A., and da Rosa Righi, R. (2017). Omniphr: A distributed architecture model to integrate personal health records. *Journal of biomedical informatics*, 71:70–81.

Solidity Team (2023a). Home — solidity programming language.

Solidity Team (2023b). Types — solidity 0.8.24 documentation.

Truffle (2023). Ganache - truffle suite.

Villarreal, E. R. D., García-Alonso, J., Moguel, E., and Alegría, J. A. H. (2023). Blockchain for healthcare management systems: A survey on interoperability and security. *IEEE Access*, 11:5629–5652.

Zhang, P., White, J., Schmidt, D. C., Lenz, G., and Rosenbloom, S. T. (2018). Fhirchain: applying blockchain to securely and scalably share clinical data. *Computational and structural biotechnology journal*, 16:267–278.

Zhao, Z., Li, J., Su, Z., and Wang, Y. (2023). Gasaver: A static analysis tool for saving gas. *IEEE Transactions on Sustainable Computing*, 8(2):257–267.