Integrating Blockchain Technologies with the Italian EHR Services

Mario Ciampi, Angelo Esposito, Fabrizio Marangio, Giovanni Schmid, Mario Sicuranza
Institute for High Performance Computing and Networking of the National Research Council of Italy
Naples, Italy
e-mail: {mario.ciampi, angelo.esposito, fabrizio.marangio, giovanni.schmid, mario.sicuranza}@icar.cnr.it

Abstract—**The increase of population life expectancy and of patient mobility makes necessary the development of new reliable and trust models for health provision within a patient-centric approach. For this reason, many proposals have been carried out to make the current healthcare systems able to collect and analyze the great amount of patient clinical data produced by the health organizations in an interoperable way and according to shared processes. However, despite the ability of collecting such clinical data, approaches aimed at assuring that these data are produced strictly following quality processes still lack. This work presents a permissioned blockchain architecture designed to assure integrity of data and processes related to Electronic Health Records coherent with the Italian Interoperability Infrastructure. The proposed architecture is compliant with both the Italian Regulation on Electronic Health Record and the European Regulation on privacy. A proof-of-concept prototype implemented on the top of Hyperledger Fabric framework validates the feasibility of the proposed architecture against two relevant use cases, showing that the application of blockchain technology to the healthcare sector could provide important benefits in terms of process and data integrity and quality.**

Keywords—electronic health record; blockchain; patient-centric architecture

I. Introduction

An important issue for the well-being of citizens is to have health systems suitable to the modern society and able to exploit the most recent technology, so that they can be efficient, reliable, scalable, and capable of providing adequate care to a large number of people, both in the medium and long term.

Evolving these systems – aimed firstly at preventing health diseases through the lifestyle monitoring of people and the use of innovative and non-invasive therapies based on precision medicine – is an essential condition for containing public spending and the sustainability of the same national health systems.

In the attempt to achieve this goal, huge efforts are underway in EU countries to digitize health processes for increasing usability and reliability for patients and healthcare personnel, allowing for a reduction in time and costs.

The areas in which improvements can and must be achieved are still many, and the margins of enhancement allowed by emerging technologies like permissioned blockchains for the secure and transparent processing of distributed workflows can be really substantial, such as to revolutionize prevention and treatment approaches [1]. Indeed, current IT systems are rooted on data producers (e.g., hospitals and healthcare companies) with the aim of collecting health information, while infrastructures and protocols designed to guarantee traceability and a "patient-centric" approach are lacking, if not completely absent. This complicates and makes healthcare costlier for citizens, as well as favoring the incidence of accidental errors and frauds, often with serious consequences in terms of public health.

The integrity and traceability of health documents and processes provided through blockchain technologies can also increase confidence of patients and physicians in emerging fields like Telemedicine.

One of the major IT solution to advance healthcare in the last decade is the Electronic Health Record (EHR), which is a "longitudinal collection of health data and documents about an individual's lifetime with the purpose of supporting continuity of care, education and research, and ensuring confidentiality at all times" [2]. EHRs allow improving the management of health processes by increasing efficiency and decreasing costs.

In [1], we proposed a blockchain network to support the decentralized management of EHRs, specifically designed according to the Italian EHR interoperability architectural model. In the context of that work, we developed a proof-of-concept prototype and performed a set of simulations for showing the effectiveness of our design, and the advantages of deploying a blockchain network for implementing access control and auditability at fine grain.

The present work extends the previous contribution by specifying the integration approach between the blockchain network and the EHR infrastructure, both in terms of architecture and implementation. We point out the main interactions among the different layers composing the system, discuss how these interactions can be implemented through the blockchain technologies, and show their workflow for some use cases. We also present and discuss, with a series of experimental sessions, new and more effective implementations of access control at the blockchain layer, which are in line with the new strategies and tools offered by the blockchain development platform used.

The rest of the paper is organized as follows. Section II describes relevant background and related works. Section III presents our contribution, giving the system requirements and its core architecture. Section IV details the prototype developed, whereas Section V concludes the paper.

II. Background and Related Work

In this section, firstly, a general overview on the Italian EHR interoperability framework is presented, paying attention on its main limitations and how they can be overcome. Secondly, the most significant scientific related work is described.

A. The Italian EHR interoperability framework

The Italian National Health Service (SSN) is a system of organizations, facilities and services that have the purpose of guaranteeing all citizens, under conditions of equality, universal access to the equitable provision of health services. The Italian Constitution provides for legislative protection of the State and the Regions for the protection of health. The State determines the *essential levels of assistance* that must be guaranteed throughout the national territory, while the Regions plan and manage health care in their area in full autonomy [3].

In the last decade, several public health service organizations have undertaken many initiatives in order to improve the quality of health services by applying information and communication technologies. The most significant efforts performed regard the design and implementation of Health Information Systems (HISs) [4], and in particular EHR systems. These ones permit to collect the digital health information related to a patient produced by the healthcare facilities and services on the national territory [5].

In order to overcome the problem of interoperability among the different regional EHR systems, since 2012 specific Italian norms have been issued, leading national Institutions (Agency for Digital Italy, Ministry of Health, Ministry of Economy and Finance, with the technical support of the National Research Council of Italy) to define the national EHR interoperability architectural model. This model is based on a set of regional IT platforms that interact each other by means of a national framework, namely National Interoperability Infrastructure (INI), as shown in Figure 1.

Each regional IT system has the aim of indexing into a registry the digital clinical documents related to its patients, whereas such documents are stored into the data repositories typically located at the health facilities [6].

INI is conform to the registry/repository paradigm based on the IHE XDS Integration profile specifications, with the scope of facilitating the sharing of patient EHRs across health enterprises within an *affinity domain*, which is a set of technical policies and codes shared by a group of healthcare facilities that intend to work together [7].

With regards to the data structure, two different approaches are used for managing digital documents and metadata.

Clinical documents are structured conforming to the HL7 CDA Rel. 2.0 standard [8], which consists of two main sections: *header*, containing contextual data (like patient name, author, etc.); *body*, devoted to the representation of the clinical content. Each type of clinical document is structured according to the Italian Implementation Guides, which are national localizations of the HL7 CDA Rel. 2 standard.

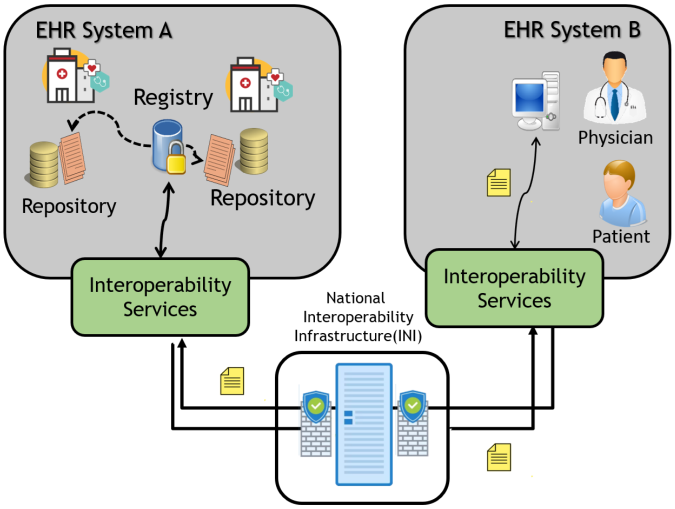
EHR System A
Registry
Repository
Repository
EHR System B
Physician
Patient
Interoperability Services
Interoperability Services
National Interoperability Infrastructure(INI)

Fig. 1. Italian EHR interoperability framework.

Metadata are a set of data that describe the clinical documents produced by the clinical facilities, with the aim of facilitating their indexing and retrieval. Such metadata contain information like patient identifier, author, document reference and so on. They can easily be mapped to the contextual data memorized in the header of the HL7 CDA Rel. 2.0 documents. The structure and the types of the metadata comply with the IHE XDS profile.

Moreover, the processes that formally describe all the activities that each involved actor has to perform have been modeled. Such processes describe the steps to index, search for, and retrieving patient health metadata and documents, wherever they are memorized on the national territory. All the regional IT systems have a regional node, which is the interface among the internal services and resources and the other regional nodes. Each interaction among such regional nodes, based on consolidated international health informatics standards, are mediated by INI. Along with the architectural model and the business processes, the functional and privacy requirements, as well as the technical specification for assuring interoperability are defined and applied [9].

These shared technical specifications permit the regional EHR systems to collect patient health documents, make authorized health professionals able to consult such documents if the patient has provided her/his consent, and interact with other regional EHR systems to exchange clinical information.

B. Issues and new directions

Despite the efforts made so far to develop a national federated architecture for the interoperability of EHR systems in Italy, significant actions are still to be taken in order to ensure an effective and correct implementation of the health interoperability processes.

Specifically, every interoperability process is in part executed through one only regional EHR system (inside a Region) and the other part is executed by other IT systems that interact with the regional system through INI. INI permits to control and partially track the requests coming from regional EHR systems, whereas the interactions occurred within a regional system are logged by this last one. For these reasons, it is complex, for the Regional EHR system, to verify that all the activities of a specific process are appropriately executed, unless to analyze all the event logs generated by the distributed systems involved. Moreover, even considering a regional context, the operations performed are often tracked by different autonomous subsystems, not allowing this way the possibility to certify that the tasks executed are compliant to the desired workflows.

A complementary architecture, opportunely designed to store in a reliable and effective way the operations executed and interoperable with the national architectural framework, would allow ensuring patients, health professionals, and government organizations that the health data are produced according to the specified and shared procedures.

C. Related work

In the last years, many academic and industry works concerning blockchain technologies and their applications have been performed in various sectors besides fintech.

Healthcare, alongside with the supply-chain industry, has probably one of the highest prospects on opportunities from these technologies. A search for the term “blockchain” on PubMed returned 21 results in 2017, 74 results in 2018, 132 in 2019, and 134 results in the first half of 2020. According to a recent study [10] by Frost & Sullivan, blockchain in healthcare is slowly starting to migrate from pilot proof of concept to commercial deployments, mainly across select enterprise-level B2B-focused use cases (e.g., credentialing, claim adjudication, supply chain, and so on). The expected revenue will be 500.7 million US$ in 2022, and pharma companies will probably be the early adopters of blockchain systems compared to other healthcare stakeholders. Various companies have already implemented or are working on putting a blockchain system to the test for a healthcare use case, (e.g., [11] [12] [13] [14] [15]), and as for July 2020 there are at least ten major healthcare blockchain consortia [10] [16]. Below, for the sake of brevity, we will limit our discussion to three major projects, which have resulted in working implementations. Indeed, they exploit different and significant approaches to the management of EHRs that have influenced our work.

MedRec [17] is a project initiated in 2016 by MIT Media Lab and Beth Israel Deaconess Medical Center, with the aim to overcome four important issues in the healthcare context: fragmented data, slow access to medical data, systems interoperability, and patient agency. It provides a decentralized approach in which the permissions, data storage location, and audit logs are maintained in the blockchain, while all healthcare information remains in the already pre-existing EHR systems. The project has developed two blockchain platforms both built on Ethereum’s framework, but with major differences. Version 1.0 [18] was a small-scale, private network with specific APIs, whilst the current version 2.0 is developed using Go-ethereum (Geth) and Solidity, but with changes to the amount of information stored on the blockchain for improving the scaling and privacy properties of transactions. Other major differences concern the consensus and governance protocols. MedRec 1.0 uses the Ethereum’s proof-of-work protocol with appropriate parameters, where the mining process would be performed by medical researchers, who in turn would gain access to aggregated and anonymized data useful to further medical research. However, this approach poses concerns about the security and governance of patient data. In the current version, therefore, the EHR providers maintain the blockchain, resulting in a small and closed set of nodes that can reach consensus without the cost of mining. Providers use a proof-of-authority to append new blocks, and also to determine who is in their group.

Patientory [19] is both the name of a digital health company established in 2015 and a no-profit association for developing and governing the PTOYNet blockchain. PTOYNet is a fork of Quorum, which in turn is an enterprise-focused version of Ethereum, mainly implemented by developers of JPMorgan Chase. Quorum executes smart contracts within the Ethereum Virtual Machine, but uses alternatives to the mining-based consensus protocol of Ethereum; moreover, it has built-in the feature of transaction confidentiality thanks to end-to-end encryption. PTOYNet has been adapted from Quorum in order to store healthcare records and manage their transactions through the PTOY token, providing an ecosystem for healthcare organizations to collaborate and innovate in a completely decentralized fashion. In exchange for PTOY, patients and healthcare organizations are able to use the network to rent health information storage space and execute health-specific smart contract payments and transactions. Patientory Inc. gains its revenue from the Software as a Service (SaaS) annual contract, as well as from population health management services from the aggregation of data on the platform: machine learning services for supporting physicians to perform medical diagnoses, patient-provider UIcare coordination, and patient engagement. In 2018, the company launched on the market a mobile distributed application (DApp), which leverages the services offered by the PTOYNet platform. At the time of writing, the approximate return on investment (ROI) in PTOY if purchased at the time of launch is -98.75% [20].

Medicalchain [13] is an infrastructure to securely store and share EHRs: any interactions with EHRs are recorded as transactions on the network, but the EHRs are encrypted and stored in data stores within appropriate regulatory jurisdictions. Its first implementation was released in February 2018 and is built on a double blockchain. The first blockchain is a permission-based Hyperledger Fabric architecture, which allows varying access levels to the EHRs: users can control who can view their records, how much they see and for what length of time. The second blockchain is Ethereum, which is used to run all the applications and services for the Medicalchain platform through the ERC20-compliant cryptocurrency token MedToken (MTN). MTNs have been offered through an initial coin offering (ICO) crowd selling process started on February 1st, 2018. At the time of writing, Medicalchain has a current supply of 500,000,000 MTN with 308,865,295.76 MTN in circulation, with an approximate ROI of -98.88% [21].

The previous examples should point out the difficulties of realizing a blockchain EHR management system, both in terms of technical deployment and governance. These difficulties are exacerbated by the EU regulations in different ways. For example, the storage of EHRs in the ledger is not only inappropriate since blockchain systems do not have the requisites of massive databases, but they make very difficult to enforce the right to data modification or erasure under particular circumstances, as stated by the Articles 16 and 17 of the General Data Protection Regulation (GDPR) [22]. More generally, blockchains underline the challenges of adhering to the requirements of data minimization and purpose limitation in the current form of the data economy.

III. The Proposed Architecture

In this work, we discuss an innovative blockchain architecture that, as its core functionality, enforces the integrity of the clinical data and processes managed through the Italian EHR infrastructure by implementing the auditing of the actions performed on such data and their resulting status in a decentralized, interoperable, tamper-proof and timeline ledger. The blockchain system acts as middleware and network infrastructure, which is interposed between the application (regional EHR services) and network layers. Depending on the coupling level with the application layer, and the additional requirements with respect to the INI technical specifications [9], the blockchain system could optionally provide for the specification of access control policies, also complementary to policies defined through INI, and the enforcement of their relative authorization rules by means of ACLs (Access Control Lists). This can allow specific healthcare ecosystems belonging to INI (hospital chains, health districts, etc.) to enrich and customize the data management and access to care functions for their patients.

The proposed architecture is compliant with both the recently introduced GDPR and the national EHR interoperability architectural model described in Section II.A. Indeed, the design of the architecture was driven by the functional and non-functional requirements listed in Tables I and II. These requirements stem from the framework of fundamental rights of the GDPR and the organizational constraints for the national EHR interoperability architectural model. They can be subdivided in mandatory (M) and recommended (R) requirements, and further grouped into basic (B) requirements and those deriving from needs related to patients (P) and those arising from the needs of health organizations (O).

Patients' needs are related to their privacy and the rights to data access (Article 15 GDPR) and data portability (Article 20 GDPR), which provide patients with control over what others do with their personal data and what they can do with that personal data themselves.

TABLE I
Mandatory Requirements (MBR: Mandatory Basic Requirement; MPR: Mandatory Patient Requirement; MOR: Mandatory Organization Requirement)

MBR1	Identification and authorization for all the actors
MBR2	Document indexing functions: the reference IT system for a patient has the responsibility of memorizing index metadata related to all his/her documents, even if they are produced and archived by health facilities managed by other IT systems
MBR3	Search and retrieval functions for documents related to a specific patient
MBR4	Backup and restore functions
MBR5	Audit operations are required: it is necessary to track all the operations carried out by all the actors
MBR6	Data and process integrity has to be assured
MPR1	Patients must be able to hide their data from healthcare practitioners
MPR2	Patients need to have the ability to know how and when their data are accessed and for what purpose. This will be possible through the *disclosure* property, as indicated in the EU directives
MPR3	Patients must be able to search for and retrieve their health data in the system
MOR1	The holder of the data treatment is the healthcare organization that produced data
MOR2	Healthcare organizations must provide protection to the data they hold

The basic requirements (MBR1-MBR4 and RBR1) are assured by the implementation of the Italian national interoperability technical specifications for EHR systems [23] [24] for both primary and secondary uses. Instead, MBR5 and MBR6 requirements stem from the blockchain adoption. Indeed, the blockchain functionalities allow to have corroborate and auditable evidence that all workflows at the application layer are correctly executed, provided that these workflows were coded as appropriate (sets of) transactions.

The requirements MPR1, MPR2, and MPR3 are satisfied thanks to the introduction of the blockchain infrastructure and the use of access policies applied to the patient's health documents. The MOR1 requirement is satisfied by the principle defined for the EHR: in fact, the organization of the author of the document is also the owner of it. The use of ACLs allows the patient to have total access control to his/her health documents. Functional requirements MOR2 and ROR1 are assured because data access control is managed through ACL in the blockchain. The RPR1 requirement is guaranteed because all information requesting/accessing health documents is stored in the blockchain. Moreover, in the blockchain, the ACL provides a quick and easy way to modify the access policies on health documents by the Healthcare Organizations, therefore, the requirements ROR2, ROR3 and ROR4 are satisfied.

A. Architecture overview

Each regional EHR system provides both a set of IT services for i) the management of regional health documents and ii) the interoperability with other regional EHR systems. These last services can be used by the actors of other regional systems, through the INI interoperability infrastructure. They are:

TABLE II
RECOMMENDED REQUIREMENTS (RBR: RECOMMENDED BASIC REQUIREMENT; RPR: RECOMMENDED PATIENT REQUIREMENT; ROR: RECOMMENDED ORGANIZATION REQUIREMENT))

RBR1	Anonymization / pseudo-anonymization data functions
RPR1	Patients should be able to provide access to healthcare practitioners that are not entitled to access their data
RPR2	Functions for allowing a patient to send data produced by certified devices to organizations for storage and management
RPR3	Patients must be able to hide their data from specific healthcare practitioners
ROR1	Every healthcare organization can manage security policies in line with regulations, but with a certain level of autonomy
ROR2	Every healthcare organization should be able to design its own security policy and to enforce it. The definition of the access policies must be implemented in total freedom and through a highly flexible mechanism
ROR3	Healthcare organizations should be able to change quickly and easily the access policies of a given document
ROR4	The access control procedure should not add a significant administrative overhead

- *Search for document:* a healthcare professional searches for health documents (by satisfying search criteria) for a patient coming from another Region.
- *Create or update Document:* a healthcare professional creates or updates a healthcare document related to a patient coming from another Region.
- *Delete Document:* a healthcare professional deletes a specific document previously created for a patient coming from another Region.
- *Transfer Patient:* the management of the metadata index related to a patient is transferred to another regional system.
- *Retrieve Document:* a healthcare professional retrieves a specific document.

The access to services has to be allowed only to authorized actors with respect to national, regional and local access policies (e.g., policies derived from the patient's will).

The actors get access to the system thanks to one of the two authentication methods prescribed in Italy, which are SPID or CNS. SPID [25] is the unique system of access with digital identity to the online services of the Italian Public Administration, in accordance to the Electronic Identification and Trust Services Regulation (eIDAS). CNS [26] is a device (i.e., a Smart Card or USB stick) that contains a "digital certificate" of personal authentication.

Both individuals and companies are identified by an Italian identifier, named *fiscal code* (CF).

The main actors of the national interoperability system are:

- *Patient:* any citizen accredited on the EHR system, who needs health care.
- *Region:* territorial entity with its statute, powers, and functions according to the principles established by the Italian Constitution.
- *Health Organization:* any public/private health company authorized by the Ministry of Health.
- *Admin Officer:* an administrative official in charge of patient registration and accounting for a health company.
- *Organization Physician:* a physician working in a health organization registered in the network, who is in charge of carrying out diagnostic examinations or medical reports for patients, thus creating their health data.

The national access policies are based on the following attributes:

- *role:* the requestor role;
- *locality:* the location of the requestor when she/he performs the request;
- *purpose of use:* the reason for the request;
- *resource type:* the kind of document requested;
- *organization identifier:* the identifier of the health professional's organization;
- *subject identifier:* the identifier of the requestor (health professional's fiscal code);
- *resource identifier:* the patient's fiscal code;
- *consent:* the consent provided from the patient to the health organization to the care treatment;
- *action identifier:* the type of request operation.

The standard used for the exchange of authentication and authorization data is OASIS SAML 2.0. The assertions are transported in the header of SOAP messages, which are exchanged between the regional platforms and INI, in the context of interoperability services. The attributes present in the assertion are compared with the ACL to allow access or not to EHR data and services.

Figure 2 shows the modules relating to the regional EHR services, those related to the Interoperability Services, and the *mapping module* allowing the interaction of the regional EHR services with a permissioned blockchain, where the function of the nodes composing the blockchain network are defined and implemented in relation to the computing facilities of the organizations providing the services. This module is in charge of assuring the correct mapping of participants, data and interactions. In particular, for each regional EHR service request, this module: i) specifies the participants, assets and transactions that are involved in the blockchain system to fulfill the request, ii) encodes a transaction proposal, and iii) submits such proposal to a blockchain peer. It is worthwhile to stress in this respect that the participants in the blockchain network are identified by the mapping module in relation to the entities managed at the higher layers of the overall architecture. Actual participants are indeed enrolled, identified and authenticated at the application layer, which is also responsible for defining user access permissions to data and resources according to attribute-based access control policies. At the blockchain layer, further access control lists are instantiated and enforced for such participants.

B. Blockchain network overview

Our system is a kind of permissioned blockchain where, according to recent design principles [27], network nodes have different functions and can be subdivided in *validating*, *endorsing* and *ordering peers*. This approach decouples agreements about interoperability processes from the consensus

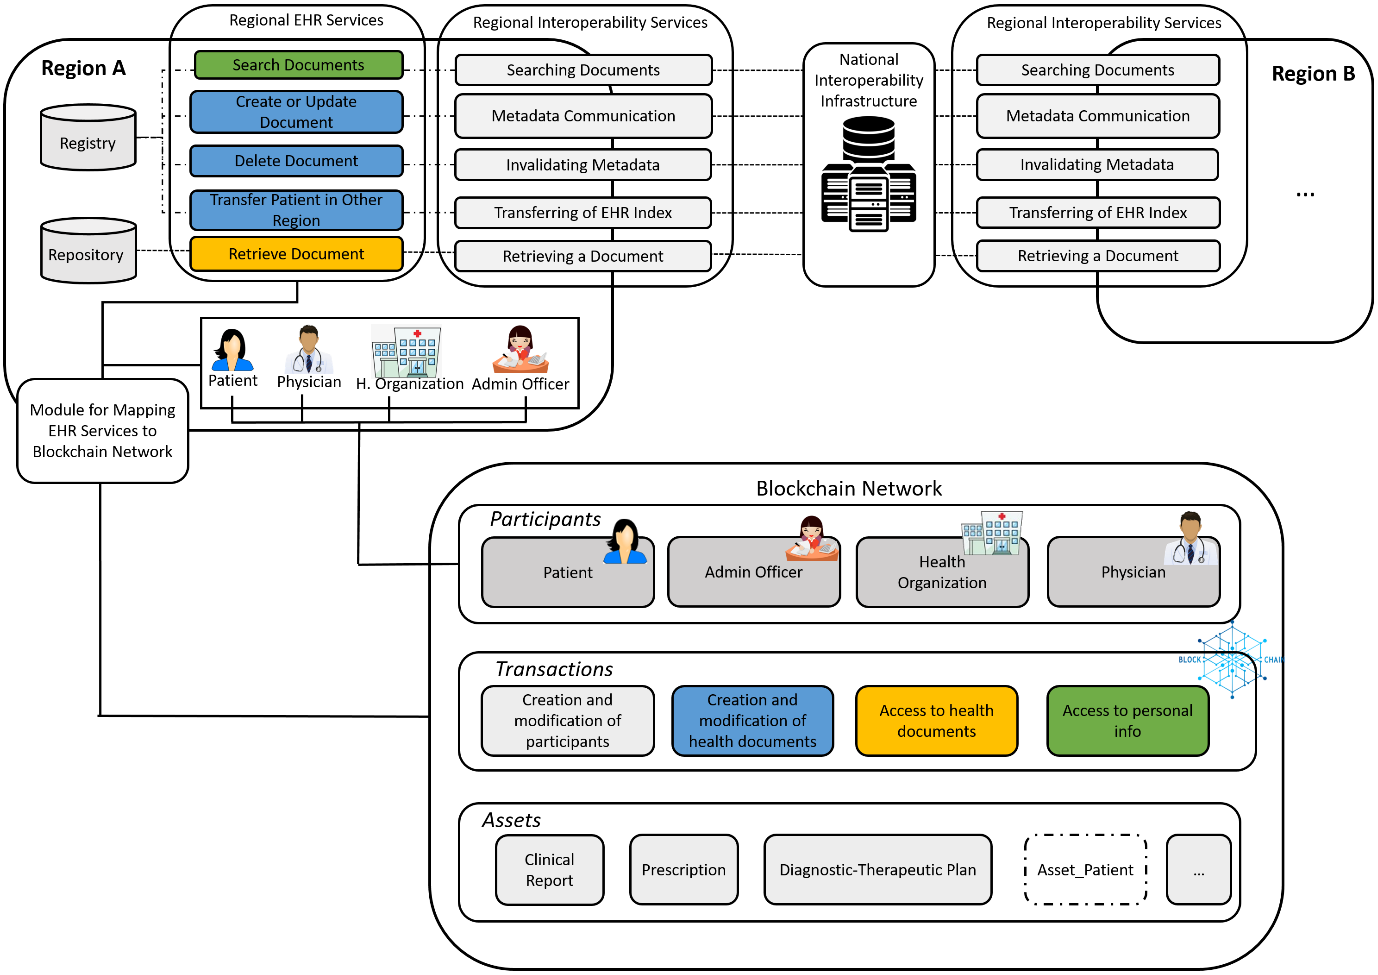

Fig. 2. The proposed blockchain architecture integrated with the EHR services.

concerning the transactions and their ordering, which have to be recorded in the blockchain.

Validating nodes have their own copy of the ledger: they are healthcare-related companies, institutions and control agencies that check for transaction I/O versus the current status of the ledger.

Endorsing peers are validating nodes that, on the basis of a consensus policy provided at the application layer, have got the additional task of checking transaction correctness both syntactically and by running them. Endorsers can be defined on a per-transaction basis, and this role is typically assumed by the entities involved in a given transaction, or the organizations they belong to. For example, in case of a pharmaceutical prescription, the endorsers could be the healthcare company to which the physician who made the prescription belongs or a regional institution representing the SSN. In the Italian scenario, the endorser role could be acted by SistemaTS (the Italian IT framework where all the digital prescriptions are memorized and retrieved by the pharmacies for the dispensations of the medications).

Ordering peers are nodes that – through a suitable consensus protocol, implemented in a dedicated module – have to assemble transactions in blocks and select the next block of the chain for the relevant blockchain. Ordering nodes do not need to store any blockchain, nor they are aware of transaction contents: they just assemble the endorsed transactions received in blocks and communicate the next block to the validating nodes for the relevant blockchain via a gossiping protocol. Ordering nodes can be supplied by the same organizations that provide validating peers, or by different organizations, depending on the governance and trust models defined for the consortium of organizations involved in the blockchain network.

Finally, the users of the blockchain network in our context are patients, physicians and other personnel of the healthcare sector. They require services at the application layer that are encoded as suitable transactions to be submitted to the blockchain by the mapping module (Figure 2).

Transactions define the logic for the management of participants and health documents through the blockchain. According to the national EHR interoperability architectural model described in Section II.A, the actual participants profiles and patients' health documents are stored and accessed through the regional EHR systems. The blockchain network introduced with the proposed architecture keeps track how such profiles and health data are produced or consumed. Specifically, patient profiles have got corresponding blockchain assets just in case the blockchain is used to complement or substitute the EHR systems for the specification and enforcement of access control policies. In these cases, the patient-related asset encodes the patient identifier, plus a list of identities or roles having some privileges on patient's documents. These ACLs are defined according to the patient-centric requirements for access data management, as optionally provided at the application layer, and in function of the roles and data types supported in the technical specification for EHR interoperability [28]. It is worth noting that if the access control policies are the high level ones defined in [29], they can be implemented in the blockchain through smart contracts, as they are defined system-wide rather than at the level of individuals.

Unlike participants, *each* health document is represented as a blockchain asset. This asset contains a set of metadata derived by [23] and the link to the actual document. In the application scenarios envisioned in this work, the aforementioned link is only used for tracking purposes, since the indexing functions are offered at the application layer; however, this element could be used as a real hyperlink to the document if the indexing functions were provided via blockchain. Some of the fields specified in assets encoding health documents are (see Figure 3):

- *authorPerson:* defines the CF identifier of the author, in our case the physician that created the asset;
- *authorRole:* defines the role of the author (like general practitioner);
- *authorInstitution:* defines the CF identifier of the company in which the physician who created the asset works;
- *patientID:* the CF identifier of the participant for whom the document is created;
- *classCode:* defines the class of the document (prescription – PRS, medical report – REF, and so on);
- *confidentialityCode:* defines the level of confidentiality of the asset (unrestricted, low, moderate, normal, restricted, very restricted);
- *mimeType;* identifies the MIME type of the indexed document.

Transactions are articulated in the following four sets, depending on their scope:

- *Creation and modification of participants:* various transactions permit to create and modify the blockchain assets related to individual participants. Participants are univocally identified in the system by their CF, which can be set and modified only by the creator of the participant, following the rules given in Section III.A. The whole process is managed by the mapping module in accordance with the access control rules for the participants defined at the application layer. Typically, as detailed previously, assets are created only in case participants represent patients, and only when the blockchain is used to manage fine-grained, patient-centric access control policies.

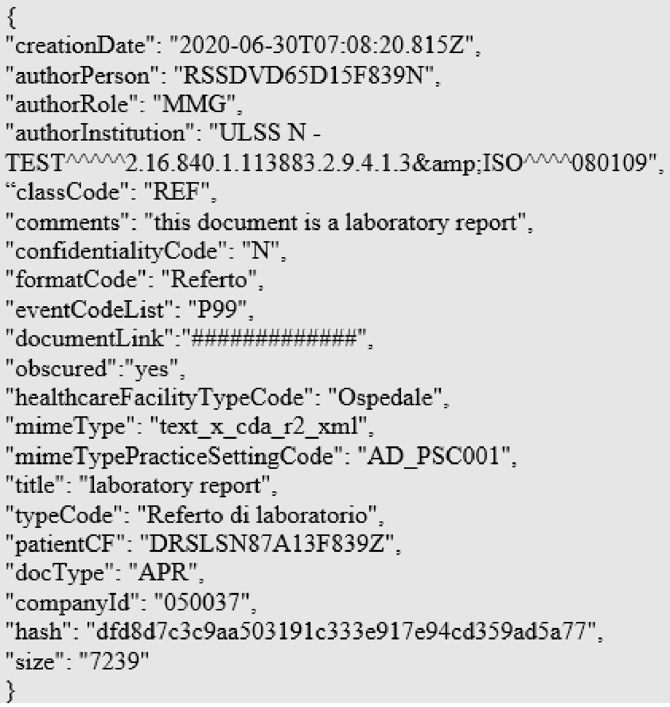

```
{
"creationDate": "2020-06-30T07:08:20.815Z",
"authorPerson": "RSSDVD65D15F839N",
"authorRole": "MMG",
"authorInstitution": "ULSS N -
TEST^^^^^^2.16.840.1.113883.2.9.4.1.3&amp;ISO^^^^^080109",
“classCode": "REF",
"comments": "this document is a laboratory report",
"confidentialityCode": "N",
"formatCode": "Referto",
"eventCodeList": "P99",
"documentLink":"##############",
"obscured":"yes",
"healthcareFacilityTypeCode": "Ospedale",
"mimeType": "text_x_cda_r2_xml",
"mimeTypePracticeSettingCode": "AD_PSC001",
"title": "laboratory report",
"typeCode": "Referto di laboratorio",
"patientCF": "DRSLSN87A13F839Z",
"docType": "APR",
"companyId": "050037",
"hash": "dfd8d7c3c9aa503191c333e917e94cd359ad5a77",
"size": "7239"
}
```

Fig. 3. Example of a blockchain asset encoding a health document.

- *Creation and modification of health documents:* consistently with the fact that only agents previously authorized by the high-level policies in [23] and/or by a patient can create or update their health documents, these rules apply also for the related assets managed in the blockchain. Only the creator of a health document (and its corresponding asset) can subsequently modify it, but in any case this will be tracked in the blockchain through a suitable transaction. If provided as functionality by the access control policy, the patient can give read access for the document to other participants in the network, and this will be tracked in the blockchain through a specific transaction affecting the asset encoding the patient's profile (see next item).
- *Access to health documents:* this kind of transactions allows the access to the health documents of a patient. By default, other than by their creators, health documents can be read by the patients to which they refer to and by the practitioners indicated in [23], in function of the purpose of use of the document. If the blockchain is used to implement patient-centric access control policies, these last are implemented as a specific set of read ACLs provided in the patient's profile. By tracking access requests, this kind of transactions implements the MBR4 requirement of disclosure (see Section III.A)
- *Access to personal info:* patients must give their explicit consent to other participants (e.g., healthcare companies) for reading the information encoded in their asset. This kind of transactions implements the requirement MPR1 and is regulated by another set of specific read ACLs in

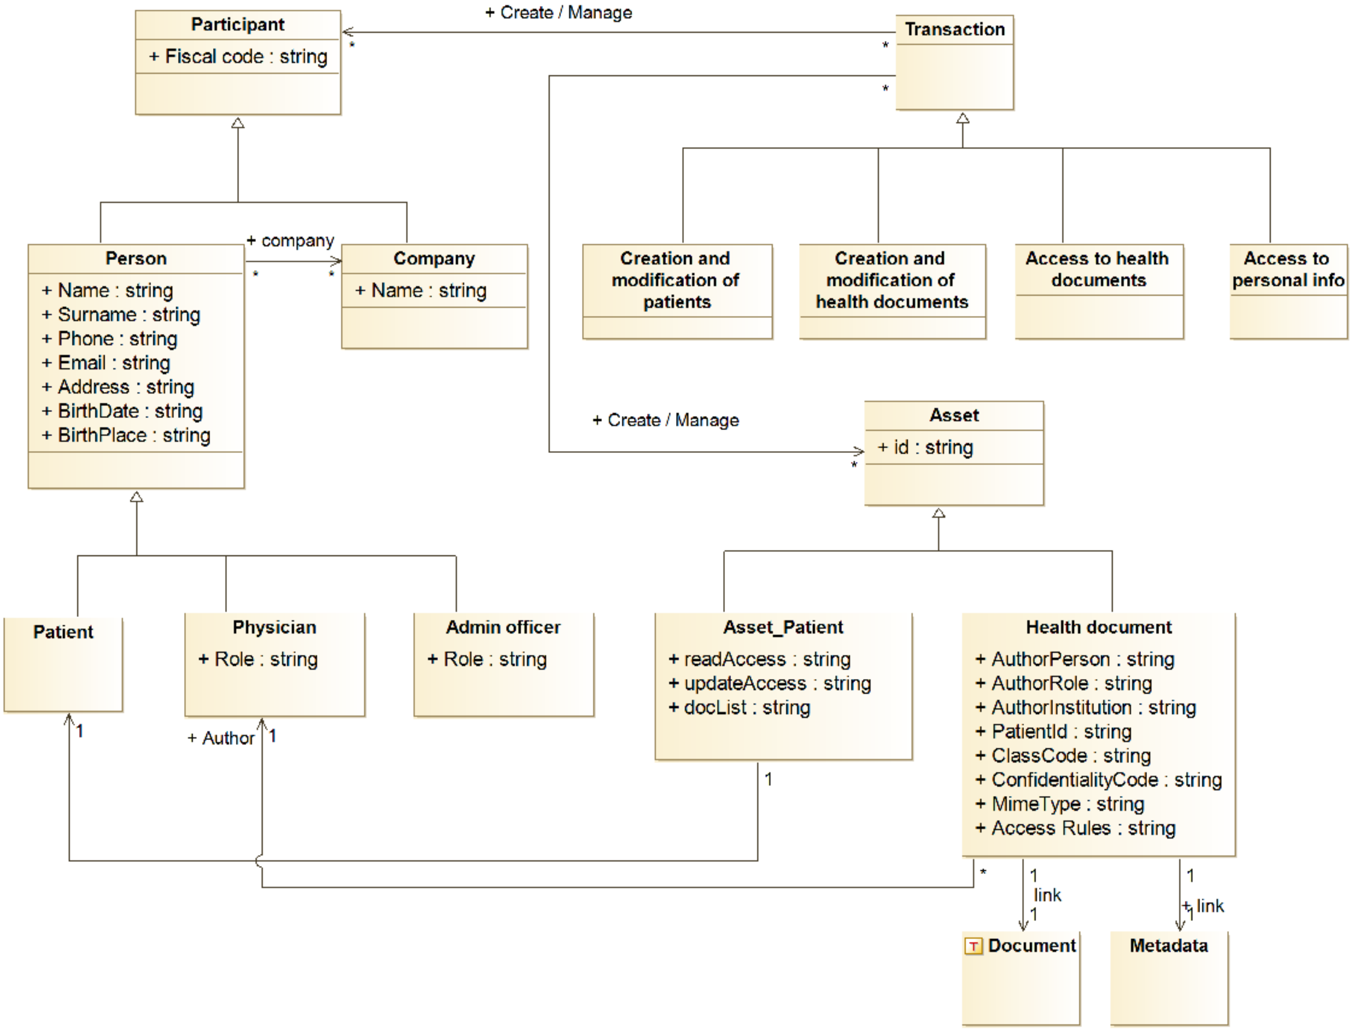

Fig. 4. Class Diagram of participants, data and transactions managed in the blockchain network.

the patient's profile.

A class diagram representing the various participants, data and transactions that are managed in the blockchain network is given in Figure 4.

C. Use cases

This section aims to give a more comprehensive overview of the overall architecture resulting from our proposal of integration of a blockchain network with the Italian EHR Interoperability Framework. It describes the interactions among its different layers in the two use cases of document search and document retrieval.

Document search

Figure 5 shows the sequence diagram related to the search for one or more health documents by a physician. The physician, authenticated on the EHR regional system, uses the Search Documents service depicted in Figure 2, which forwards the request to INI through the Searching Documents interoperability service. These interactions are concisely indicated as an "Access" phase performed by the physician in Figure 5, who at this point will have the search request submitted to INI. INI carries out the validation of the request by verifying if the user has the access right to the service and, if these checks are passed, then forwards the request to the Regional Services. In turn, a node implementing a regional service: i) forwards the request to a blockchain peer, which reads the ACLs provided by the asset of the patient to whom the documents belong; ii) compares those ACLs with the metadata encoded in the assets of the required documents; and iii) returns a list of metadata that the physician is authorized to access (the list may be empty).

Document retrieval

Once the physician has received a list of references to the documents satisfying the search query, she/he can carry out the document retrieval request, according to the same logic described above. In this case, as shown in Figure 5, the

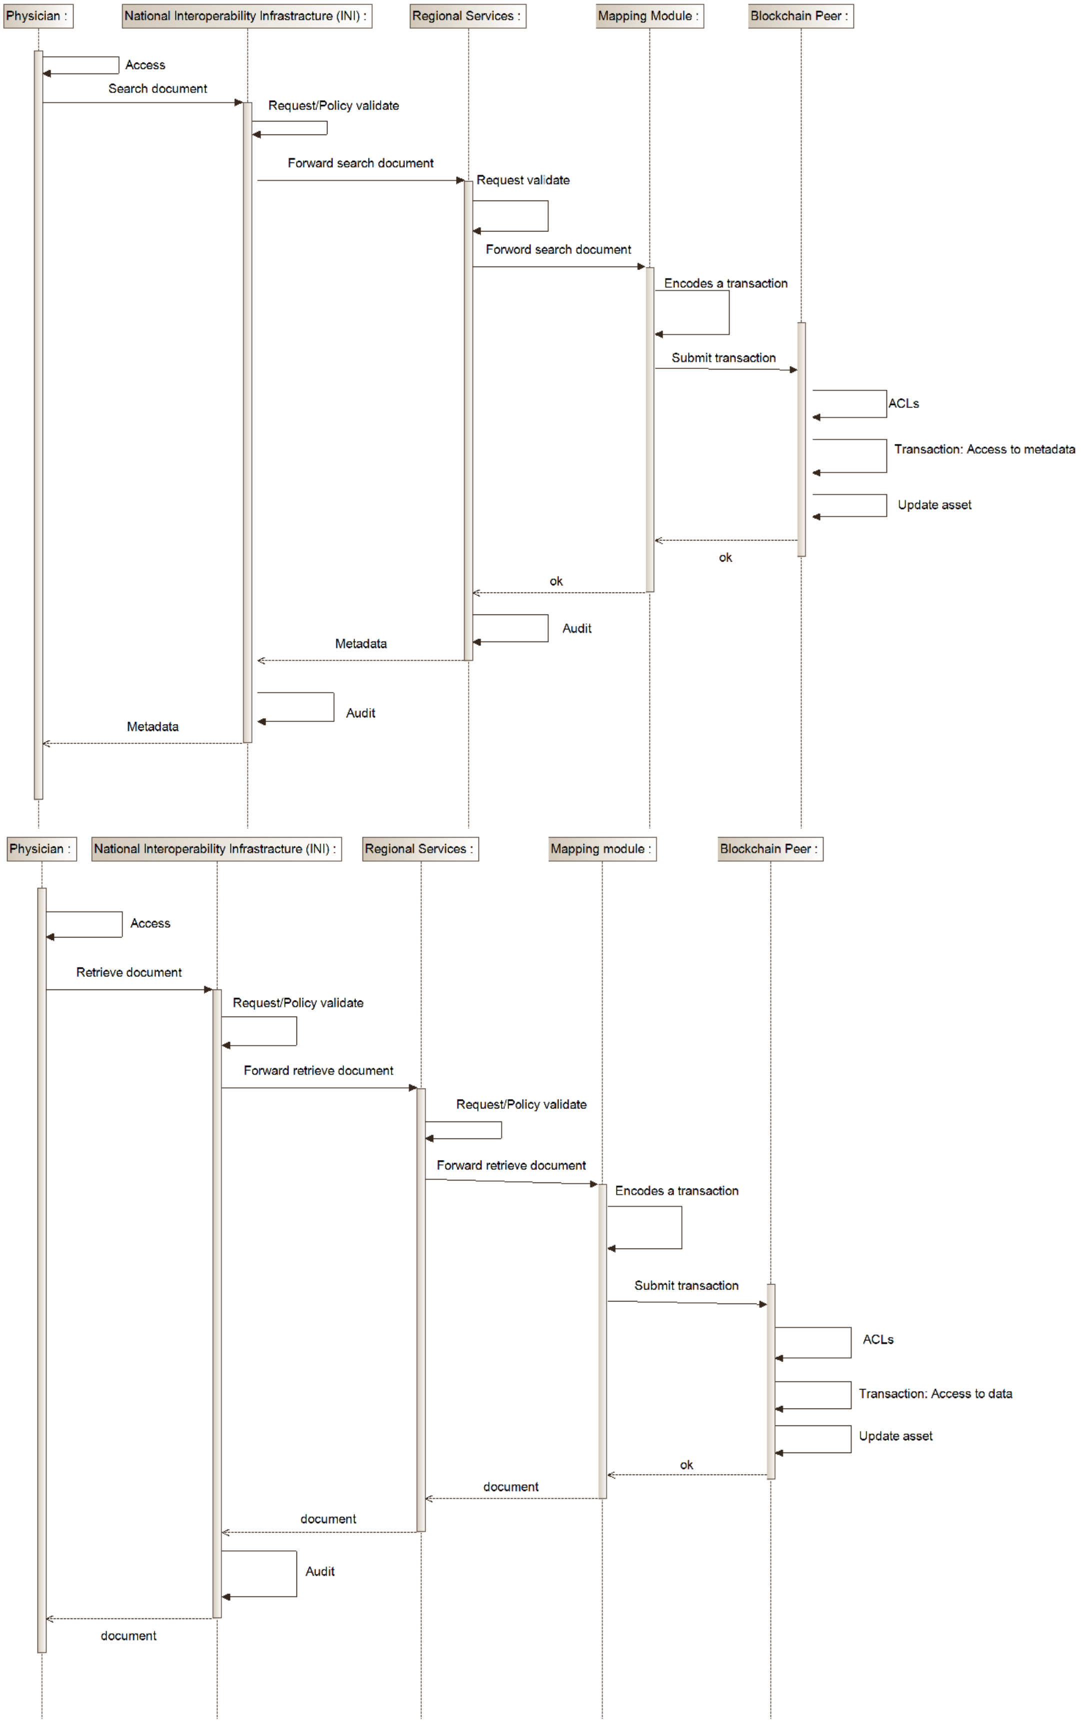

Fig. 5. Sequence Diagrams that illustrates the interactions among the actors of the architecture for document search (top) and retrieval (bottom).

blockchain peer, after the comparison of the ACLs coded in the patient's asset with the metadata contained in the document's asset, returns to the physician the required health document, or an "access denied" error.

IV. Implementation and results

In [1], we implemented a prototype of a permissioned blockchain network, in order to assess the proposed architecture. At the time, we deployed a blockchain network and enrolled some templates of participants using Hyperledger Composer [28]. Then, we used the Hyperledger Fabric v1.2 runtime [29] to perform a set of simulations, in order to show the effectiveness of the enforcement of patient-centric access control policies through the blockchain for some relevant use cases. The ultimate scope of those simulations was to show that through blockchain technologies it is possible to easily and effectively implement all the functional requirements illustrated in Section III.

In this extended contribution, we pursue the same goal but with the following major differences:

- *participants are not actually enrolled in the blockchain network.* This complies with the overall architectural framework described in Section III.A and with the role and functions of the blockchain network as detailed in Section III.B;
- *ACLs are no more defined through Hyperledger Composer.* This stems primarily from the fact that Hyperledger Composer was considered an obsolete project by September 2019. However, implementing the ACLs by working directly at the chaincode and/or asset levels allows for the specification of more flexible and powerful ACLs, as we will show shortly.

The complete set of functional requirements listed in Tables I and II gives rise to a patient-centric data management framework, where individuals have the capability to set permissions for their documents in a punctual way (e.g., set for a specific physician the read access to a specific document). However, according to the technical specifications [23], patients can manage their documents in a much less punctual way. Permissions to read and edit documents are indeed enforced on the basis of system-wide rules stemmed from roles such as *general practitioner*, *hospital physician* or *pharmacist*, which the patient cannot change. What patients can do are some actions like obscuring a given document to all other actors, obscuring all their information and documents, or delegate some other people to manage their health documents (as in the case of patients with severe health conditions). In such respect, the recommended requirements listed in Table II can be seen as complementary to the mandatory requirements already provided by standard implementations of the regional EHR services, and that can be effectively and reliably implemented through the blockchain technologies.

In the rest of this section, we will show how it is possible to use Hyperledger Fabric chaincode and assets to implement ACLs with different levels of granularity, so to enforce at the blokchain layer access control policies that integrate more or less extensively those provided by the regional EHR services.

Figures 6 and 7 show two possible implementations of a patient asset, in relation to more or less stringent ACLs.

The patient asset in Figure 6 has a complex structure in order to indicate: i) who can read or update the asset (through the *readAccess* and *updateAccess* strings); ii) who can create a patient's document (through the *docCreateAccess* string); and, iii) who can read or update specific patient's documents (through the *docReadAccess* and *docUpdateAccess* strings). The patient asset in Figure 7 has a less complex structure, corresponding to the fact that patients can still decide who can read or update their assets, but they can not set access permissions related to specific health documents. In both cases,

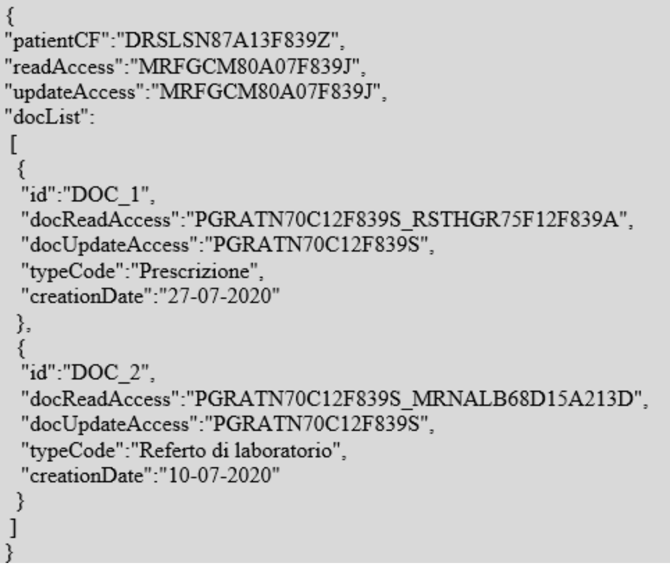

```
{
"patientCF":"DRSLSN87A13F839Z",
"readAccess":"MRFGCM80A07F839J",
"updateAccess":"MRFGCM80A07F839J",
"docList":
 [
  {
   "id":"DOC_1",
   "docReadAccess":"PGRATN70C12F839S_RSTHGR75F12F839A",
   "docUpdateAccess":"PGRATN70C12F839S",
   "typeCode":"Prescrizione",
   "creationDate":"27-07-2020"
  },
  {
   "id":"DOC_2",
   "docReadAccess":"PGRATN70C12F839S_MRNALB68D15A213D",
   "docUpdateAccess":"PGRATN70C12F839S",
   "typeCode":"Referto di laboratorio",
   "creationDate":"10-07-2020"
  }
 ]
}
```

Fig. 6. Asset with document-oriented ACL management.

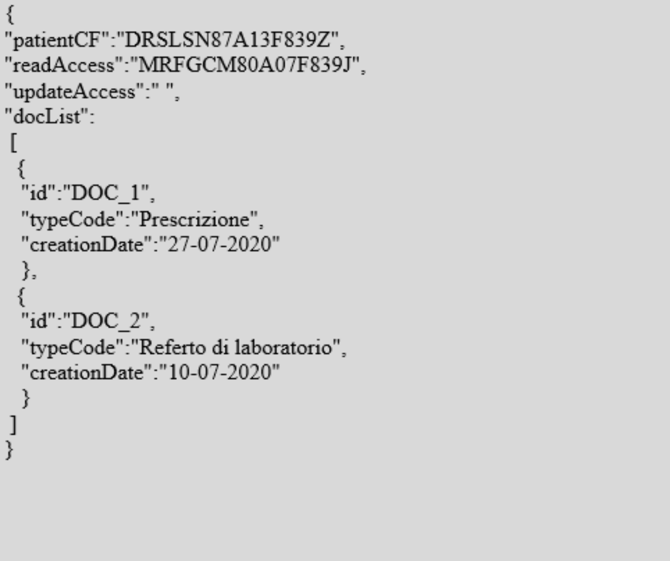

```
{
"patientCF":"DRSLSN87A13F839Z",
"readAccess":"MRFGCM80A07F839J",
"updateAccess":" ",
"docList":
 [
  {
   "id":"DOC_1",
   "typeCode":"Prescrizione",
   "creationDate":"27-07-2020"
   },
  {
   "id":"DOC_2",
   "typeCode":"Referto di laboratorio",
   "creationDate":"10-07-2020"
   }
 ]
}
```

Fig. 7. Asset without document-oriented ACL management.

every patient's health document is listed in the *docList*, which represents the timeline of all patient's documents and is used to set the ACLs related to specific documents.

A. Fine-grained ACLs

Let us first consider an access control scenario like that in our previous work [1], where patients can manage punctual permission to authorize specific participants to read or update their documents.

Figure 8, Figure 9, Figure 10 and Figure 11 show a comparison of the ACLs implemented with Hyperledger Composer and the new implementation via chaincode. Specifically, these figures show the possible implementation of two ACLs, for the management of the read permission to the patient asset and a health document, respectively.

While Hyperledger Composer allows to set permissions at a higher level, working with chaincode requires the knowledge of a chaincode-oriented programming language (in this case Java), but allows the specification of more complex, fine-grained ACLs.

For instance, Hyperledger Composer requires a participant and a resource to set up an ACL. Thus, the participant must be explicitly enrolled in the blockchain network through a suitable *membership service provider*.

In the context of our architecture (see Section III), this is a strong limitation that would preclude from having more assets communicating with each other, and from managing ACLs through assets other than through identities.

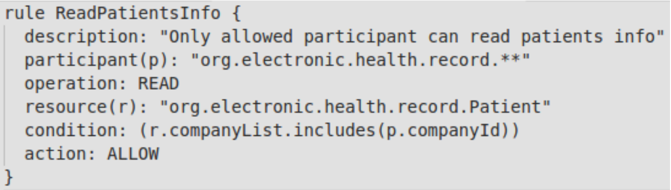

```
rule ReadPatientsInfo {
  description: "Only allowed participant can read patients info"
  participant(p): "org.electronic.health.record.**"
  operation: READ
  resource(r): "org.electronic.health.record.Patient"
  condition: (r.companyList.includes(p.companyId))
  action: ALLOW
}
```

Fig. 8. Hyperledger Composer ACL for reading patient info.

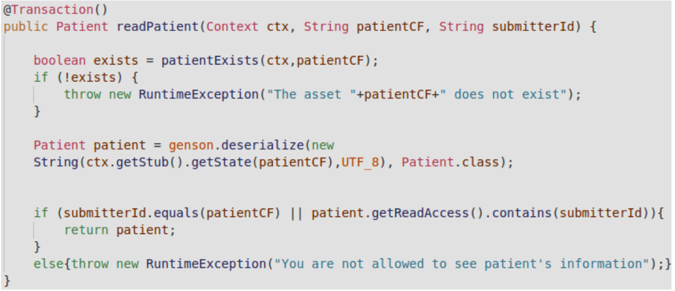

```
@Transaction()
public Patient readPatient(Context ctx, String patientCF, String submitterId) {

    boolean exists = patientExists(ctx,patientCF);
    if (!exists) {
        throw new RuntimeException("The asset "+patientCF+" does not exist");
    }

    Patient patient = genson.deserialize(new
    String(ctx.getStub().getState(patientCF),UTF_8), Patient.class);


    if (submitterId.equals(patientCF) || patient.getReadAccess().contains(submitterId)){
        return patient;
    }
    else{throw new RuntimeException("You are not allowed to see patient's information");}
}
```

Fig. 9. Chaincode ACL for reading a patient asset.

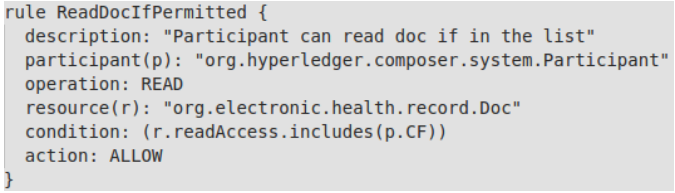

```
rule ReadDocIfPermitted {
  description: "Participant can read doc if in the list"
  participant(p): "org.hyperledger.composer.system.Participant"
  operation: READ
  resource(r): "org.electronic.health.record.Doc"
  condition: (r.readAccess.includes(p.CF))
  action: ALLOW
}
```

Fig. 10. Hyperledger Composer ACL for reading a document.

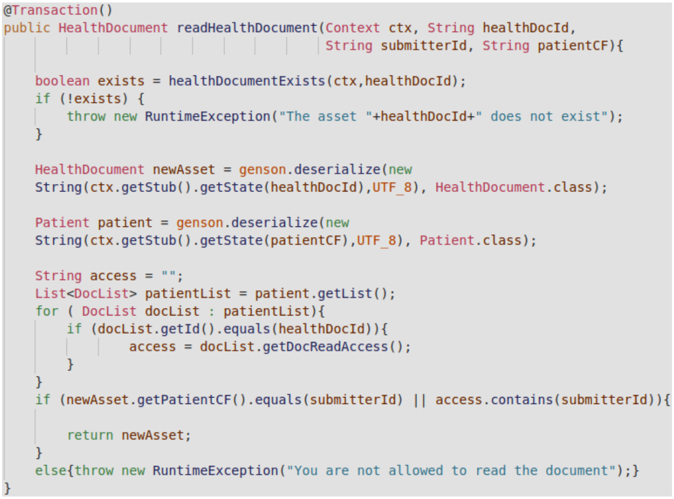

```
@Transaction()
public HealthDocument readHealthDocument(Context ctx, String healthDocId,
                                          String submitterId, String patientCF){

    boolean exists = healthDocumentExists(ctx,healthDocId);
    if (!exists) {
        throw new RuntimeException("The asset "+healthDocId+" does not exist");
    }

    HealthDocument newAsset = genson.deserialize(new
    String(ctx.getStub().getState(healthDocId),UTF_8), HealthDocument.class);

    Patient patient = genson.deserialize(new
    String(ctx.getStub().getState(patientCF),UTF_8), Patient.class);

    String access = "";
    List<DocList> patientList = patient.getList();
    for ( DocList docList : patientList){
        if (docList.getId().equals(healthDocId)){
                access = docList.getDocReadAccess();
        }
    }
    if (newAsset.getPatientCF().equals(submitterId) || access.contains(submitterId)){

        return newAsset;
    }
    else{throw new RuntimeException("You are not allowed to read the document");}
}
```

Fig. 11. Chaincode ACL for reading a document.

Simulations have been carried out in the Hyperledger Fabric 2.0 runtime in order to verify that the implementations of the ACLs via chaincode, along the same lines as previously illustrated, permit to enforce the authorization rules performed by INI. These simulations consist in the realization of a scenario in which a patient is able to provide explicit authorization to a specific health organization to access his/her own prescription document produced by a general practitioner. The implementation of this scenario was carried out by exploiting the feature provided by the chaincode concerning the possibility of making the assets relating to the patient and that relating to the health document able to communicate each other.

The results reached give corroborate evidence that all the requirements listed in Tables I and II could be easily and effectively implemented trough blockchain technologies.

B. Obscuration of patient's documents

This section provides a proof-of-concept for the implementation of the "document obscuration" functionality using the proposed architecture based on the blockchain technology. This capability could be necessary for some type of documents containing sensitive data about major health problems or addictions. In some cases, the practitioner at document creation time has to specify that it has to be obscured because of national regulations; however, patients have to be able to obscure or make their documents visible at their choice [23].

We will illustrate a very simple workflow, where a document is first created as visible by a practitioner and then is obscured by the patient (to whom it refers to).

The document creation is realized by the transaction showed in Figure 12.

This transaction requires an Id for the health document and a JSON String with all the fields of the asset as showed in Figure 3. During this phase, the physician can set the field "obscured" either on "Yes" or on "No" (see Figure 3); in this

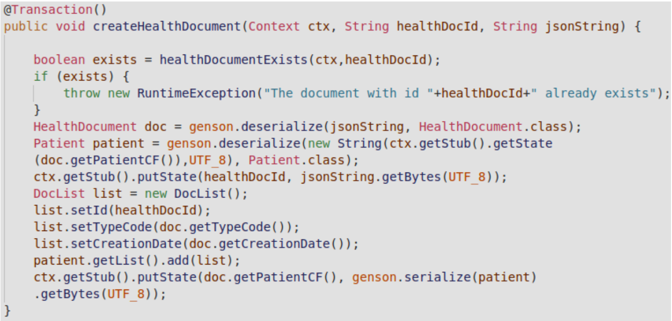

```java
@Transaction()
public void createHealthDocument(Context ctx, String healthDocId, String jsonString) {

    boolean exists = healthDocumentExists(ctx,healthDocId);
    if (exists) {
        throw new RuntimeException("The document with id "+healthDocId+" already exists");
    }
    HealthDocument doc = genson.deserialize(jsonString, HealthDocument.class);
    Patient patient = genson.deserialize(new String(ctx.getStub().getState
    (doc.getPatientCF()),UTF_8), Patient.class);
    ctx.getStub().putState(healthDocId, jsonString.getBytes(UTF_8));
    DocList list = new DocList();
    list.setId(healthDocId);
    list.setTypeCode(doc.getTypeCode());
    list.setCreationDate(doc.getCreationDate());
    patient.getList().add(list);
    ctx.getStub().putState(doc.getPatientCF(), genson.serialize(patient)
    .getBytes(UTF_8));
}
```

Fig. 12. Transaction to create a health document.

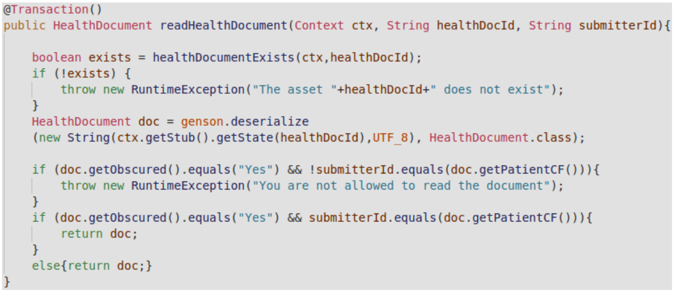

```java
@Transaction()
public HealthDocument readHealthDocument(Context ctx, String healthDocId, String submitterId){

    boolean exists = healthDocumentExists(ctx,healthDocId);
    if (!exists) {
        throw new RuntimeException("The asset "+healthDocId+" does not exist");
    }
    HealthDocument doc = genson.deserialize
    (new String(ctx.getStub().getState(healthDocId),UTF_8), HealthDocument.class);

    if (doc.getObscured().equals("Yes") && !submitterId.equals(doc.getPatientCF())){
        throw new RuntimeException("You are not allowed to read the document");
    }
    if (doc.getObscured().equals("Yes") && submitterId.equals(doc.getPatientCF())){
        return doc;
    }
    else{return doc;}
}
```

Fig. 13. Transaction to read a health document.

example, the choice is “No”, whilst the document ID was set to “TEST_DOC”.

At this point, both the patient to whom the document refers to and the other authorized participants in the network can read the document launching the transaction in Figure 13, which results in an outcome like that showed in Figure 14.

Actually, any other participant in the blockchain, can read the document. This is because the field “obscured” is set on “No” for this document,as shown by Figure 15, and the access control policies concerning participants in this case are enforced at the application layer.

Now, if the patient sets to “Yes” the “obscured” field through the transaction shown in Figure 16 and then launches the read transaction, he/she can read the document and verify that the obscured field has actually been modified (Figure 17).

The last step of this simulation consists in showing that with any other ID different from that of the patient, the obscured document cannot actually be read. Figure 18 shows the result of a read attempt by a random ID, which turns out in the error “No document satisfying your request”, as provided by the readHealthDocument transaction.

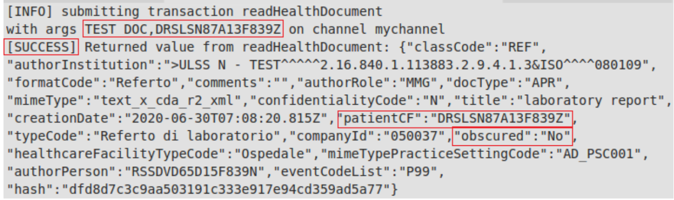

```
[INFO] submitting transaction readHealthDocument
with args TEST_DOC,DRSLSN87A13F839Z on channel mychannel
[SUCCESS] Returned value from readHealthDocument: {"classCode":"REF",
"authorInstitution":">ULSS N - TEST^^^^^2.16.840.1.113883.2.9.4.1.3&ISO^^^^080109",
"formatCode":"Referto","comments":"","authorRole":"MMG","docType":"APR",
"mimeType":"text_x_cda_r2_xml","confidentialityCode":"N","title":"laboratory report",
"creationDate":"2020-06-30T07:08:20.815Z","patientCF":"DRSLSN87A13F839Z",
"typeCode":"Referto di laboratorio","companyId":"050037","obscured":"No",
"healthcareFacilityTypeCode":"Ospedale","mimeTypePracticeSettingCode":"AD_PSC001",
"authorPerson":"RSSDVD65D15F839N","eventCodeList":"P99",
"hash":"dfd8d7c3c9aa503191c333e917e94cd359ad5a77"}
```

Fig. 14. Result of read transaction submitted by the patient.

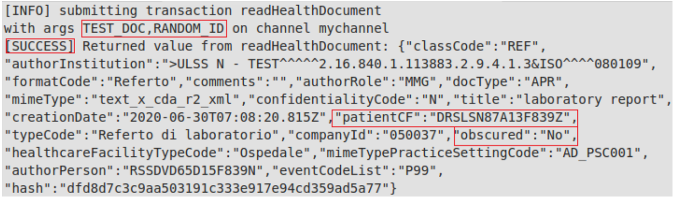

```
[INFO] submitting transaction readHealthDocument
with args TEST_DOC,RANDOM_ID on channel mychannel
[SUCCESS] Returned value from readHealthDocument: {"classCode":"REF",
"authorInstitution":">ULSS N - TEST^^^^^2.16.840.1.113883.2.9.4.1.3&ISO^^^^080109",
"formatCode":"Referto","comments":"","authorRole":"MMG","docType":"APR",
"mimeType":"text_x_cda_r2_xml","confidentialityCode":"N","title":"laboratory report",
"creationDate":"2020-06-30T07:08:20.815Z","patientCF":"DRSLSN87A13F839Z",
"typeCode":"Referto di laboratorio","companyId":"050037","obscured":"No",
"healthcareFacilityTypeCode":"Ospedale","mimeTypePracticeSettingCode":"AD_PSC001",
"authorPerson":"RSSDVD65D15F839N","eventCodeList":"P99",
"hash":"dfd8d7c3c9aa503191c333e917e94cd359ad5a77"}
```

Fig. 15. Result of read transaction performed by a random Id.

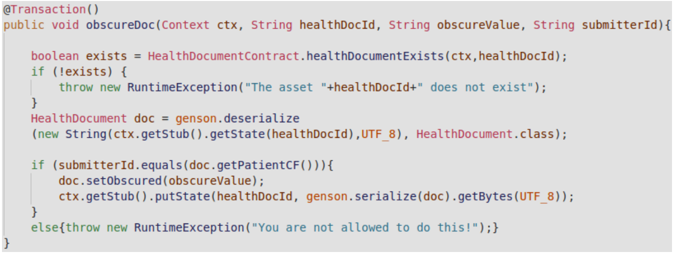

```java
@Transaction()
public void obscureDoc(Context ctx, String healthDocId, String obscureValue, String submitterId){

    boolean exists = HealthDocumentContract.healthDocumentExists(ctx,healthDocId);
    if (!exists) {
        throw new RuntimeException("The asset "+healthDocId+" does not exist");
    }
    HealthDocument doc = genson.deserialize
    (new String(ctx.getStub().getState(healthDocId),UTF_8), HealthDocument.class);

    if (submitterId.equals(doc.getPatientCF())){
        doc.setObscured(obscureValue);
        ctx.getStub().putState(healthDocId, genson.serialize(doc).getBytes(UTF_8));
    }
    else{throw new RuntimeException("You are not allowed to do this!");}
}
```

Fig. 16. Transaction to obscure a document.

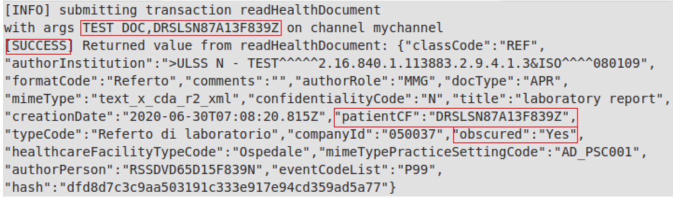

```
[INFO] submitting transaction readHealthDocument
with args TEST_DOC,DRSLSN87A13F839Z on channel mychannel
[SUCCESS] Returned value from readHealthDocument: {"classCode":"REF",
"authorInstitution":">ULSS N - TEST^^^^^2.16.840.1.113883.2.9.4.1.3&ISO^^^^080109",
"formatCode":"Referto","comments":"","authorRole":"MMG","docType":"APR",
"mimeType":"text_x_cda_r2_xml","confidentialityCode":"N","title":"laboratory report",
"creationDate":"2020-06-30T07:08:20.815Z","patientCF":"DRSLSN87A13F839Z",
"typeCode":"Referto di laboratorio","companyId":"050037","obscured":"Yes",
"healthcareFacilityTypeCode":"Ospedale","mimeTypePracticeSettingCode":"AD_PSC001",
"authorPerson":"RSSDVD65D15F839N","eventCodeList":"P99",
"hash":"dfd8d7c3c9aa503191c333e917e94cd359ad5a77"}
```

Fig. 17. Read transaction launched by the patient after obscuration.

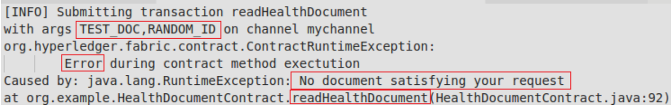

```
[INFO] Submitting transaction readHealthDocument
with args TEST_DOC,RANDOM_ID on channel mychannel
org.hyperledger.fabric.contract.ContractRuntimeException:
        Error during contract method exectution
Caused by: java.lang.RuntimeException: No document satisfying your request
at org.example.HealthDocumentContract.readHealthDocument(HealthDocumentContract.java:92)
```

Fig. 18. Error in read transaction.

All the above tests were performed using the Hyperledger Fabric extension for VSCode that supports versions of the Fabric framework from 1.4 onwards [30].

V. Conclusion and Future Work

This paper has presented a blockchain architecture for the decentralized management of clinical documents collected in EHRs, compliant with the GDPR. The proposed architecture is designed for facing the integrity and traceability issues concerning the current national EHR framework for the interoperability of the regional systems in Italy. The architecture lies on a network that represents the new core components that, integrated with the federated EHR IT system, permits to easily and effectively implement the health processes in a verifiable and correct manner. The proposed network is coupled with a suitable access control and security framework to protect patient’s health data. This framework respects a set of functional and non-functional requirements identified on the basis of the Italian norms and the GDPR principles, without prejudicing neither the usability of the system nor its scalability and management. A proof-of-concept prototype of the architecture has been developed to prove its feasibility in two real scenarios. For this reason, a set of transactions opportunely identified are mapped with the application services. Even if the proposed work is customized for the Italian context, the methodology adopted permits to simply decline it to other contexts.

Future work is planned for implementing a testbed in order to evaluate the effectiveness of the EHR management system resulting by integrating it with the blockchain network illustrated in this work.

REFERENCES

[1] M. Ciampi, A. Esposito, F. Marangio, G. Schmid, and M. Sicuranza, "A blockchain architecture for the italian ehr system," in *The Fourth International Conference on Informatics and Assistive Technologies for Health-Care, Medical Support and Wellbeing HEALTHINFO 2019*. IARIA, November 2019, pp. 11–17.

[2] I. Iakovidis, "Towards personal health record: current situation, obstacles and trends in implementation of electronic healthcare record in europe," *International journal of medical informatics*, vol. 52, no. 1-3, pp. 105–115, 1998.

[3] "Italian Ministry of Health," http://www.salute.gov.it/portale/lea/dettaglioContenutiLea.jsp?lingua=italiano&id=5073&area=Lea&menu=vuoto/, [retrieved on 2020.11.17].

[4] A. D. Black, J. Car, C. Pagliari, C. Anandan, K. Cresswell, T. Bokun, B. McKinstry, R. Procter, A. Majeed, and A. Sheikh, "The impact of ehealth on the quality and safety of health care: a systematic overview," *PLoS med*, vol. 8, no. 1, p. e1000387, 2011.

[5] F. Aminpour, F. Sadoughi, and M. Ahamdi, "Utilization of open source electronic health record around the world: a systematic review," *Journal of research in medical sciences: the official journal of Isfahan University of Medical Sciences*, vol. 19, no. 1, p. 57, 2014.

[6] M. Ciampi, M. Sicuranza, A. Esposito, R. Guarasci, and G. De Pietro, "A technological framework for ehr interoperability: Experiences from italy," in *International Conference on Information and Communication Technologies for Ageing Well and e-Health*. Springer, 2016, pp. 80–99.

[7] "Integrating the Healthcare Enterprise," http://www.ihe.net/, [retrieved on 2020.11.17].

[8] "HL7 Version 3 Clinical Document Architecture (CDA) Release 2," https://www.hl7.org/implement/standards/product_brief.cfm?product_id=7/, [retrieved on 2020.11.17].

[9] "Agency for Digital Italy of the Presidency of the Council of Ministers," https://www.fascicolosanitario.gov.it/, [retrieved on 2020.11.17].

[10] "Frost and Sullivan, Global Blockchain Technology Market in the Healthcare Industry 2018-2022," https://www.businesswire.com/news/home/20191016005587/en/, [retrieved on 2020.11.17].

[11] "Patientory," https://patientory.com/technology/, [retrieved on 2020.11.17].

[12] "GemOS," https://enterprise.gem.co/, [retrieved on 2020.11.17].

[13] "Medicalchain," https://medicalchain.com/, [retrieved on 2020.11.17].

[14] "EncrypGen," https://encrypgen.com/the-dna-economy/, [retrieved on 2020.11.17].

[15] "SimplyVitalHealth," https://github.com/SimplyVitalHealth/drs/, [retrieved on 2020.11.17].

[16] "Hashed Health," https://hashedhealth.com/newsletter-sept-2019/, [retrieved on 2020.11.17].

[17] "MedRec," https://medrec.media.mit.edu/technical/, [retrieved on 2020.11.17].

[18] A. Azaria, A. Ekblaw, T. Vieira, and A. Lippman, "Medrec: Using blockchain for medical data access and permission management," in *2016 2nd International Conference on Open and Big Data (OBD)*. IEEE, 2016, pp. 25–30.

[19] C. McFarlane, M. Beer, J. Brown, and N. Prendergast, "Patientory: A healthcare peer-to-peer emr storage network v1." *Entrust Inc.: Addison, TX, USA*, 2017.

[20] "CoinMarcketCap Patientory," https://coinmarketcap.com/it/currencies/patientory/, [retrieved on 2020.11.17].

[21] "CoinMarketCap Medicalchain," https://coinmarketcap.com/it/currencies/medical-chain/, [retrieved on 2020.11.17].

[22] "GDPR," https://eur-lex.europa.eu/eli/reg/2016/679/oj/, [retrieved on 2020.11.17].

[23] "Italian EHR Interoperability: Italy Affinity Domain," https://www.agid.gov.it/sites/default/files/repository_files/regole_tecniche/affinitydomainitalia_versione_2.1.pdf, [retrieved on 2020.11.17].

[24] "Italian EHR Interoperability: Framework and Basic Services Dataset," https://www.agid.gov.it/sites/default/files/repository_files/regole_tecniche/frameworkdataset_versione_2.1.pdf, [retrieved on 2020.11.17].

[25] "SPID," https://www.spid.gov.it/, [retrieved on 2020.11.17].

[26] "CNS," https://sistemats1.sanita.finanze.it/portale/modalita-di-accesso-con-ts_cns, [retrieved on 2020.11.17].

[27] M. Vukolić, "Rethinking permissioned blockchains," in *Proceedings of the ACM Workshop on Blockchain, Cryptocurrencies and Contracts*, 2017, pp. 3–7.

[28] "Hyperledger Composer," https://hyperledger.github.io/composer/latest/, [retrieved on 2020.11.17].

[29] "Hyperledger Fabric," https://hyperledger-fabric.readthedocs.io/, [retrieved on 2020.11.17].

[30] "IBM Blockchain Platform," https://marketplace.visualstudio.com/items?itemName=IBMBlockchain.ibm-blockchain-platform, [retrieved on 2020.11.17].