Secure and Decentralized Sharing of Medical Imaging Data via Blockchain Consensus

Vishal Patel, MD, PhD

Department of Radiological Sciences ⋄ University of California, Los Angeles

1 Introduction

Imaging studies are one of the leading drivers of modern medical decision-making [1], and thus, their accessibility to healthcare providers and patients is of critical importance. However, current techniques for transferring medical imaging data are inconvenient and occasionally wholly inadequate. In 2014, 15% of patients who visited a healthcare provider reported having to bring a radiological study or other test result to their appointment personally, and 5% needed to have a test or procedure repeated due to the unavailability of prior results [2]. Given the costs of medical image acquisition and the risks associated with delayed access to imaging results, facilitating the secure electronic sharing of radiological studies provides a natural target for improving healthcare efficiency and patient outcomes.

In this paper, we outline a framework that utilizes blockchain technology to enable patients to delegate electronic access to their medical imaging data in a secure manner. We discuss the appropriateness of blockchain technology for this indication, and we describe the relative merits and drawbacks of this approach relative to several alternatives.

1.1 Medical Image Sharing

Despite the widespread availability of digital imaging and high-speed network connectivity, the persistent paradigm for medical image sharing requires that a physical copy (e.g., a CD or DVD) be couriered between providers. There is clear inefficiency and waste inherent in transcribing a digital asset onto optical media which commonly is read only once during image import at the receiving site. Moreover, this workflow imposes an undue responsibility upon the patient to ensure that the data is not lost, damaged, or intercepted in transit.

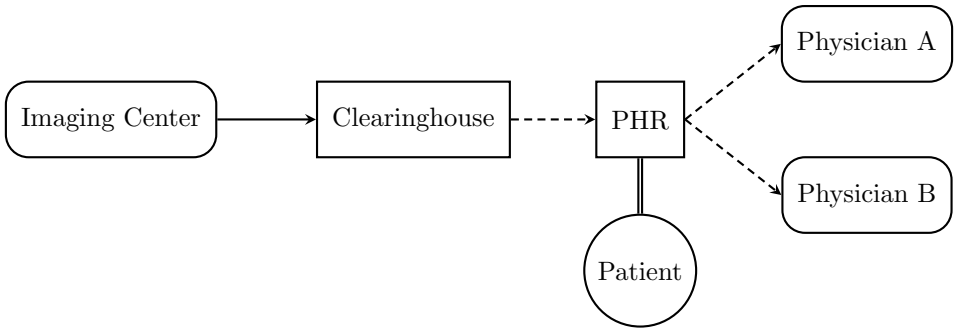

Figure 1: The RSNA Image Share Network. Arrows indicate the flow of imaging data through the network. Dashed transfers require patient interaction and authorization via the PHR. Circles represent the conceptual owners of imaging data, rounded rectangles symbolize data producers or consumers, and block rectangles denote third parties.

In order to address the shortcomings of physical media transfer, the Radiological Society of North America (RSNA) developed the Image Share Network (ISN), which represents the current state of the art for electronic transmission of medical images [3]. Briefly, the sites participating in the ISN centralize image distribution through a third-party Clearinghouse. Acquired images are uploaded to the Clearinghouse, where they are stored, indexed by the cryptographic hash of a secret token, for 30 days. Within this window, a patient can authorize a personal health record (PHR) vendor to download his or her data by divulging the token needed to reproduce the hash. The images stored in the PHR represent the virtual analogue of the optical disc, and the patient can subsequently permit his or her healthcare providers to access this data. The ISN architecture is schematized in Figure 1.

Though the ISN eliminates the need for physical media transfer, its design raises concerns regarding the involvement of intermediaries and the centralization of the sharing infrastructure. For example, the ISN workflow includes the Clearinghouse and PHR vendors as new entities with access to protected health information (PHI); these represent additional points at which an internal or external malicious actor may compromise the network and gain access to sensitive data. Moreover, the provider fulfilling the critical Clearinghouse role is able to exert significant control over the network by limiting which imaging centers and PHR vendors have access to the ISN ecosystem. At the time of this writing, less than a dozen radiology centers are enrolling patients in the network, exactly four PHR vendors (one of which is also the Clearinghouse operator) are authorized to retrieve ISN images, and less than 30% of patients whose images are sent to the Clearinghouse ever download their data into one of these approved PHRs [3, 4]. In examining these issues, the concept of a decentralized architecture presents itself as an option that may decrease the barriers to entry and encourage more widespread adoption of the image sharing network.

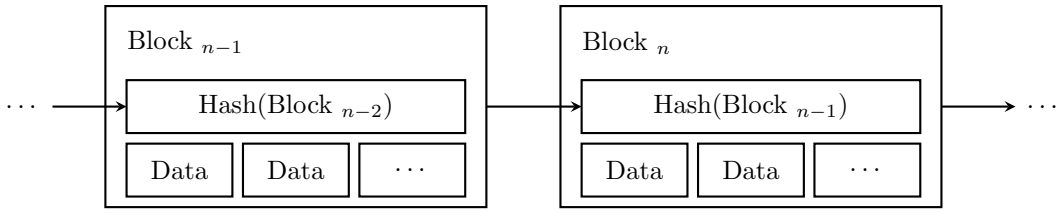

Figure 2: A simple blockchain. In addition to the data records, each block also contains a cryptographic hash of the previous block, thus ensuring block ordering and data integrity.

1.2 Blockchain Technologies

The recently introduced blockchain concept has been the principal innovation driving several prominent decentralized endeavors, and it is thus appropriate to consider its applicability to the task of medical image sharing. A blockchain is, quite simply, a data structure consisting of an ordered sequence of batched entries, termed blocks. The ordering of these blocks is established by storing a cryptographic hash of the immediate prior record within each block (Figure 2). The use of an irreversible hash function as the chaining mechanism also serves to verify the integrity of the previous block, and it is this characteristic that gives rise to the key emergent property of the blockchain as a data store: immutability. Any attempt to tamper with the data in an established block is easily detected since it changes the hash of the altered block and consequently, the hashes of all subsequent blocks in the chain.

A blockchain is maintained by a set of nodes, entities without a preexisting trust relationship that are connected through a peer-to-peer network. For the blockchain to be useful, there must be some mechanism by which the nodes can mutually agree upon the next valid block in the chain. The two most widely deployed schemes for establishing such a distributed consensus are summarized below:

- Proof-of-work is the archetypal process for block validation in which nodes compete to generate the next block by expending computational effort to solve a challenging mathematical problem. The first node to arrive at a solution broadcasts the result to the network; the solution is verified by the remaining nodes, and work begins on the next block [5].

- Proof-of-stake algorithms forego the computational challenge, but offer only a randomly selected subset of nodes the opportunity to produce each block. The probability of selection is weighted by each entity's level of existing investment in the system, typically quantified as the value or duration of asset holdings relevant to that particular blockchain [6, 7].

Blockchains thus enable many separate parties to converge upon a single, immutable record without requiring an authoritative intermediary. As we show below, these properties provide a sufficient core upon which to deploy a decentralized image sharing network.

2 Design

Though the primary application of blockchain technology to date has been to establish ledgers of transactions involving virtual tokens (i.e., cryptocurrencies), our design decisions are driven by a different set of objectives. Specifically, we intend to use a blockchain to store: 1) a list of imaging studies and the patient to whom each study belongs, 2) the set of entities authorized by the patient to access each study, and 3) the endpoint from which each study may be retrieved. In this section, we develop such a system, keeping our description intentionally general to maintain focus on the overall architecture.

The anticipated users of this data structure include the imaging centers at which the studies are acquired and stored, the patients themselves, as well as any healthcare providers, PHR vendors, cloud services, or other designees granted access to an imaging study. All of these entities are represented on the blockchain by the public portion of an asymmetric key pair. Each imaging study is referenced by its globally unique DICOM Study Instance UID, hashed before publication to prevent leakage of PHI [8].

2.1 Block Structure

Each block is comprised of two primary categories of data elements. Block header elements provide important metadata necessary to establish the sequencing and integrity of the blockchain. These include items such as the hash of the prior block (Figure 2), as well as a unique block identifier, timestamp, and total size of the block. An exhaustive list of header fields is omitted as many of these elements are conserved between blockchain implementations, and detailed explanations are available elsewhere [5, 7].

The other main category of block data consists of transaction elements. These are the unique data fields that differentiate this blockchain, and their definitions fully determine the types of information that the blockchain structure is able to store. The image sharing blockchain is characterized by the transactions enumerated below. For clarity, we omit the transactional metadata (unique identifier, timestamp, etc.) and envision a patient P who undergoes a radiographic examination X while at hospital H and wishes to share those images with primary care doctor D. We follow the conventional security protocol notation, and thus K_α and K_α^{-1} represent the public and private keys, respectively, of actor α. We then have the following minimal transaction set:

1. Define Source: This transaction establishes a source of medical imaging data by linking a public key to a uniform resource locator (URL). In our example, hospital H signs a transaction associating its public key with the URL endpoint that it uses to service image transfer requests from authorized entities. This need be done only once. The transaction thus defines the signed tuple $\{K_H, \mathrm{URL}_H\}_{K_H^{-1}}$.

2. Define Study: This transaction establishes a source as the creator and a patient as the owner of a radiological study with a particular unique identifier UID. The

tuple stored in the blockchain is $\{K_H, \{K_P, K_H, \text{Hash}(\text{UID}_X)\}_{K_P^{-1}}\}_{K_H^{-1}}$. Hospital H broadcasts this transaction after obtaining the necessary electronic signature from patient P.

3. Allow Access: This transaction enables the owner of a radiological study to authorize another party to retrieve his or her imaging data from the source endpoint URL. Patient P signs a transaction granting this ability to doctor D in our example. The inclusion of the signed tuple $\{K_P, K_D, K_H, \text{Hash}(\text{UID}_X)\}_{K_P^{-1}}$ within a validated block embeds this consent into the blockchain.

These three types of transactions are sufficient for the image sharing blockchain to satisfy the primary objectives of the ISN. Additional transaction types not explored here may provide further functionality, such as enabling patients to revoke future access to their data or allowing imaging centers to change their URL endpoints.

2.2 Validation and Consensus

Special consideration must be devoted to the process of block creation and validation in the setting of medical image sharing. Proof-of-stake systems for achieving distributed consensus have the advantage of imposing minimal computational and energy burdens on their participants. In this setting, however, we must somehow motivate the block generators to ensure that valid blocks are produced in a timely fashion and that a single chain quickly predominates any forks. Alternatively, we may instead develop a disincentive for block producers to engage in malicious or noncontributory behavior. This latter approach is similar to the concept of bonded validators in cryptocurrency proof-of-stake systems; in that construct, only nodes that have established a security deposit may participate in chain extension, and any misbehaving node is forced to forfeit its investment [9].

We note that the nature of the blockchain provides for straightforward auditing of each node's activities, including the number of blocks generated, failures to produce a block when eligible to do so, and attempts to publish invalid blocks. In addition, a node operator is able to prove its ownership of a node simply by signing a message with the private key corresponding to that node's identifying public key.

Given these preliminaries, we build the image sharing blockchain using a proof-of-stake scheme in which the probability of a node producing the next block is driven by the number of Define Study transactions designating that node's public key as the image source. This choice effectively restricts the block generation process to facilities that acquire and store medical images—a group expected to be responsive to monetary incentives put forth by a large (e.g., governmental) payer. By linking reimbursement rates to the results of periodic reviews of the relevant nodes' activities on the blockchain, we can ensure participation that yields a reliable chain of valid blocks produced at regular intervals.

2.3 Image Transfer

There are no medical images stored on the blockchain; the chain of transactions instead represents only a list of the key owners that are permitted to access each study. The actual image transmission requires that the image recipient send a signed request to the URL endpoint of the imaging source that created the study. We leverage the existing work by the Integrating the Healthcare Enterprise (IHE) initiative, which has specified a standard form for retrieving documents across domains, the ITI-43 transaction [10].

The image source public key and hashed study UID satisfy the required elements of the ITI-43 request, termed repositoryUniqueId and documentUniqueId in IHE parlance. Moreover, we require that the requesting entity authenticates itself by digitally signing the transaction using the private key corresponding to the public key that was previously granted access by the patient [11].

Upon receipt of such a request, the image source verifies that the signature is correct, checks that the repositoryUniqueId specifies its own public key, confirms that the hashed UID corresponds to a study that it previously published, and confirms—via the blockchain—that the patient has granted the requester access to these images. If all criteria are satisfied, the source returns an ITI-43 response containing the imaging study. Both the request and response are transmitted over a TLS-secured link to prevent eavesdropping.

3 Discussion

Deploying a blockchain as a distributed access control list to enable the sharing of medical records represents an unexplored application of a nascent technology, and a thorough discussion of its anticipated benefits and limitations is thus warranted. In this section, we outline the major points of consideration when implementing a blockchain of the form described above.

3.1 Advantages and Interoperability

As constructed in Section 2 and schematized in Figure 3, the image sharing blockchain achieves the objective of enabling patient-controlled image sharing without the need for a central clearinghouse. Moreover, the PHR is no longer a required additional intermediary, but rather a tool that the patient may optionally utilize to manage health records. In addition to this disintermediation, the blockchain architecture provides several additional key benefits, many of which have been explicitly targeted by the Office of the National Coordinator for Health Information Technology (ONC) as critical to the development of an interoperable health system [12].

For example, we have selected a proof-of-stake system with implicit bonding (Section 2.2), which facilitates auditing of each entity's involvement in the system and provides a simple means for large payers (e.g., the Centers for Medicare & Medicaid Services) to motivate participation. The blockchain design thus greatly simplifies the process of

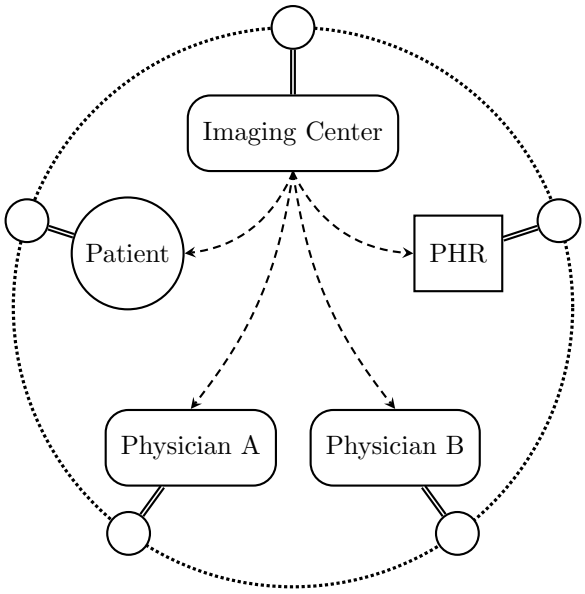

Figure 3: The image sharing blockchain. Each participant operates a node (∘) on the network which establishes the blockchain (⋯⋯). The patient provides access to chosen entities by posting blockchain transactions. Imaging data is transferred directly from the source to these authorized recipients; no central intermediary is required.

establishing an economic environment in which interoperability is a sound business decision, a priority identified in ONC's nationwide interoperability roadmap [12, §A].

Moreover, the peer-to-peer network of nodes and associated blockchain architecture represent design decisions that effectively address another interoperability criterion: the sharing of PHI with patient-authorized recipients over a secure, private, tamper-resistant network infrastructure [12, §C]. As described in Section 1.2, the chaining mechanism inherently results in an immutable data record, such that any attempt to modify established blocks is immediately evident. Beyond this, the transactions that define the image sharing blockchain are based on well-tested and widely-deployed practices utilizing public key cryptography; as long as an entity secures its private key, all information on the blockchain signed by that entity may be trusted.

We further note that the transaction types described in Section 2.1 satisfy additional properties of the interoperable health system. For example, a signed Allow Access transaction provides a straightforward representation of authorization to access electronic health information [12, §E]. In addition, the set of Define Source transactions on the blockchain effectively establishes a directory of resource locations that can be easily referenced to locate and request imaging studies [12, §M].

One of the blockchain's most unique strengths lies in its ability to establish an anonymous record while still enabling participants to authenticate themselves when required. We recall that patients are represented on the blockchain as randomly-generated public keys. As long as entities use unique key pairs to manage each study (see Section 3.2) and the private

keys are kept secure, patients cannot be identified by analyzing the blockchain record. When needed, however, any entity can prove ownership of a public key simply by signing a message with the corresponding private key. As previously suggested, an imaging center may be required to do so when undergoing an audit to determine reimbursement rates. In this manner, the asymmetric cryptography driving the blockchain satisfies the need for verifiable identity and authentication of all participants [12, §D], while at the same time protecting privacy.

Interestingly, this same property solves the problem of cross-domain identity matching, another goal identified on the interoperability roadmap [12, §L]. Existing patient matching techniques relying on standardized demographic data and other heuristics yield suboptimal match rates, particularly across geographic regions or health systems [13]. The public key identity system avoids the problems associated with incomplete, inaccurate, or outdated demographic data entirely; instead, all studies associated with provably patient-owned public keys are included in that patient's medical record.

Throughout the construction of the blockchain, we have endeavored to make best use of existing standards, file formats, and infrastructure. We continue to leverage the ubiquitously supported DICOM format for medical imaging data [12, §I], and we rely on the existing standards developed by IHE to perform the actual image transfer (Section 2.3) [12, §K]. Furthermore, we continue to use imaging centers, many of which have already invested substantially in data storage and archiving technology, as the repositories of imaging studies to make the most efficient use of existing infrastructure.

Considering all of these elements, the image sharing blockchain concept may ultimately lead to a greater ability for patients to electronically access their health information and share it at their discretion [12, §N]. Furthermore, we note that the bonded proof-of-stake blockchain architecture and the underlying cryptographic elements are application-independent and easily generalized to cover domains beyond medical imaging.

3.2 Limitations

Privacy concerns are, justifiably, major considerations when using blockchain technologies to share health information. Historically, the dominant principle for protecting health-related data has been to keep the records themselves generally inaccessible except to those directly involved in a patient's care. The blockchain privacy model, however, is more similar to one often used when conducting medical research: the data records themselves are widely accessible, but the patients to whom they refer are either secret or anonymized. In this setting, we must take steps to minimize the risk that an analysis of the blockchain transactions, perhaps combined with outside information, will permit a public key to be linked definitively to a particular individual. We have alluded previously to the requirement that patients and the entities with which they share images should generate a new key pair for each study. By using a new public key (effectively a new virtual identity) to manage each study, a patient prevents unauthorized actors from linking multiple studies to a single identity, a process that may result in the leakage of PHI; a similar technique

is used to provide anonymity on cryptocurrency networks [5]. It is equally important for the recipients of imaging data to mask their identities in a similar manner since sharing images with a public key known to be associated with an oncologist, for example, leaks information about a patient's diagnosis. We consider this privacy model to be a relative drawback of the blockchain technique as it is error-prone and requires multiple parties to act in a very deliberate manner in order to preserve patient privacy.

The security model also differs significantly between the ISN and the image sharing blockchain. On one hand, the blockchain removes a central point of failure—the Clearinghouse, a breach of which exposes all imaging studies flowing through the ISN. Nevertheless, because the ISN is a highly restricted network, it may simply reject all network traffic from unauthorized addresses, thus forcing any attack to relay via one of the small number of imaging centers or PHR vendors authorized to communicate with the Clearinghouse. No such filtering is possible on the image sharing blockchain, an open network with no central authority limiting participation. The attack surface is thus much larger; each imaging center must adequately secure its URL endpoints, and each node operator must take care to ensure the secrecy of its private keys. We note, though, that due to the decentralized architecture, the exposure of any single private key is unlikely to affect as large a number of individuals as an ISN Clearinghouse breach. The ultimate reliance on asymmetric cryptography also means that the loss of a private key results in an inability to manage the corresponding resource, requiring some off-blockchain recovery process to re-establish ownership. As with privacy, decentralization results in a more complex security model, likely overall more prone to breaches than a centralized scheme.

Finally, we must consider the issue of whether a blockchain-based sharing network is even permissible given current regulations. Several sections of the HIPAA Privacy Rule are relevant in assessing the feasibility of the proposed approach. For example, it remains to be determined if the use of random public keys and hashed study identifiers provides sufficient de-identification to exempt covered entities from the standard disclosure restrictions (45 C.F.R. §§ 164.502(d)(2), 164.514(a) and (b)). If no such exemption is possible, then we must also investigate whether patients' digital signatures on blockchain transactions serve as adequate documentation to authorize the release of PHI (45 C.F.R. § 164.508(c)). Additional federal and local regulations also likely apply, and given the severe penalties associated with the unlawful disclosure of PHI, healthcare institutions are expected to be hesitant in their adoption of this unconventional sharing architecture.

3.3 Alternatives

Though a comprehensive review of all related efforts is beyond the scope of this document, we briefly review here the main differences between the image sharing blockchain and several prominent similarly-oriented endeavors. Throughout this article, we have drawn comparisons to the RSNA ISN, and we wish to note for completeness that the ISN also targets other potential use cases that we have not addressed in this work [3].

In addition, we note that several have promoted the concept of "private" blockchains

in which participation is limited to a set of trusted organizations. This architecture again requires registration via an authoritative intermediary, with all the associated drawbacks of centralization. Having made such a concession, the practical benefit of deploying a blockchain at all becomes somewhat questionable; as we have shown, the distributed architecture increases the complexity of the privacy and security models and produces an inefficient representation of access permissions compared to a centralized scheme.

Finally, we acknowledge the efforts of the HEART Working Group in developing the the User-Managed Access (UMA) profile [14]. This represents a unique approach to distributing access control that divides the role of the image source into that of a resource server, which stores and provides the images, and an authorization server, which verifies that requesting parties have access to the images. UMA provides many of the same benefits as the blockchain architecture, and it also suffers from similar limitations. However, the decoupling of authorization from the image source raises the question of what incentive drives the operators of the authorization servers. If there is not sufficient benefit to be gained from providing such a service, the UMA scheme has the potential to degenerate into a more centralized network.

4 Conclusions

We have reviewed the basic principles of blockchain technologies and provided an overview of a blockchain implementation that may serve as a tool to enable the patient-controlled, cross-domain sharing of medical images without the need for a central authority. There are a number of unique advantages to such an approach and we have specifically highlighted the ways in which the blockchain satisfies many of the requirements of an interoperable health system. However, there are also several important limitations to this technology, and it will be essential to remain cognizant of these and to consider the relative merits of available alternatives prior to embarking upon any large-scale implementation of a blockchain as a basis for sharing health information.

References

[1] P. V. Pandharipande, A. T. Reisner, W. D. Binder, A. Zaheer, M. L. Gunn, K. F. Linnau, C. M. Miller, L. L. Avery, M. S. Herring, A. C. Tramontano, E. C. Dowling, H. H. Abujudeh, J. D. Eisenberg, E. F. Halpern, K. Donelan, and G. S. Gazelle. CT in the emergency department: A real-time study of changes in physician decision making. *Radiology*, 278(3):812–821, 2016.

[2] V. Patel, W. Barker, and E. Siminerio. ONC Data Brief, No. 30: Trends in consumer access and use of electronic health information. Office of the National Coordinator for Health Information Technology, 2015.

[3] S. G. Langer, W. Tellis, C. Carr, M. Daly, B. J. Erickson, D. Mendelson, S. Moore, J. Perry, K. Shastri, M. Warnock, and W. Zhu. The RSNA Image Sharing Network. *Journal of Digital Imaging*, 28(1):53–61, 2015.

[4] D. S. Mendelson. Sharing medical images with patients & providers with RSNA Image Share Validation. The Sequoia Project. Healthcare Information and Management Systems Society, 2016.

[5] S. Nakamoto. Bitcoin: A peer-to-peer electronic cash system. `http://bitcoin.org/bitcoin.pdf`, 2008. (Accessed: 7 Aug 2016).

[6] S. King and S. Nadal. PPCoin: Peer-to-peer crypto-currency with proof-of-stake. `https://peercoin.net/assets/paper/peercoin-paper.pdf`, 2012. (Accessed: 7 Aug 2016).

[7] Nxt Community. Whitepaper:Nxt. `https://nxtwiki.org/wiki/Whitepaper:Nxt`, 2016. (Accessed: 7 Aug 2016).

[8] A. W. Kamauu, S. L. Duvall, and D. E. Avrin. Using Java to generate globally unique identifiers for DICOM objects. *Journal of Digital Imaging*, 22(1):11–14, 2009.

[9] V. Zamfir. Introducing Casper "the friendly ghost". `https://blog.ethereum.org/2015/08/01/introducing-casper-friendly-ghost/`, 2015. (Accessed: 7 Aug 2016).

[10] IHE International. IHE IT infrastructure technical framework. `http://www.ihe.net/Technical_Frameworks/`, 2016. (Accessed: 7 Aug 2016).

[11] M. Bartel, J. Boyer, B. Fox, B. LaMacchia, and E. Simon. XML signature syntax and processing (second edition). `https://www.w3.org/TR/xmldsig-core/`, 2008. (Accessed: 7 Aug 2016).

[12] The Office of the National Coordinator for Health Information Technology. Connecting health and care for the nation: A shared nationwide interoperability roadmap, 2015.

[13] G. Morris, G. Farnum, S. Afzal, C. Robinson, J. Greene, and C. Coughlin. Patient identification and matching final report. Office of the National Coordinator for Health Information Technology, 2014.

[14] J. Richer. Health relationship trust profile for user managed access 1.0. `http://openid.bitbucket.org/HEART/openid-heart-uma.html`, 2016. (Accessed: 7 Aug 2016).