Designing Decentralized Ledger Technology for Electronic Health Records

Vikram Dhillon

Editor's note: A proposal to implement distributed ledger technology for electronic health records is outlined here. The rationale for integration of distributed ledgers in the healthcare domain is introduced, followed by a discussion of the features enabled by the use of a blockchain. An open source implementation of a distributed ledger is then presented. The article concludes with an examination of opportunities and challenges ahead in deploying blockchains for digital health.

The Electronic Health Record (EHR) is a digital version of a patient's medical history maintained by a healthcare provider over the course of their visits (Figure 1). The record includes patient information on demographics, diagnosis, vital signs, past medical history, progress over time, lab tests and more.[1-4] Healthcare providers such as hospitals and clinics share a portion of this information with other members of the healthcare system such as insurance companies and pharmacies. Even though the primary information systems are integrated across industries, the supporting information systems are often fragmented.[5]

The practice management software implemented by hospital groups and clinics is rarely compatible with other electronic health tools, and poses significant challenges to sending and receiving electronic information. This lack of interoperability in healthcare data is a serious concern because transferred records can be overwritten and present with inaccuracies. Common examples include incorrect lab values appearing in the wrong section of a patient record or missing critical care notes that physicians need to treat patients.[2-4] Patients also experience the inconvenience of additional testing, when the provider fails to transfer previous medical history from the EHR.

Hospitals have tried to solve the interoperability issues and enhance compatibility by using new interfaces and health information exchanges, but the adoption of these new mechanisms remains low.[5, 6] These solutions in turn have their own set of problems such as high costs of deployment, making them unattainable for general use across industry. Another approach to guiding software development has been creation of new standards for data transfer, but the standards are mostly a set of general guidelines. To standardize the development of health information technology (IT), most vendors require more specific standards, which are not available.[5, 6]

A new system to manage health data is required where all interested parties can monitor, access, and analyze consistent and updated information on electronic health records. This system must eliminate information silos by creating a unifying backend that can be shared by the healthcare ecosystem. One emerging technology called a distributed ledger solves this problem by providing a decentralized backend that multiple parties can view and edit consistently without the need to trust any of the other parties.[7, 8]

A distributed ledger is essentially a database without a central authority that can be shared across a network of multiple institutions. All participants in this network have their own identical copy of the ledger stored locally. Any changes to this ledger must be verified by all participants, and these changes are reflected in all copies within a few minutes. The security and accuracy of the ledger is maintained through cryptographic signatures and keys, which control the level of access to the shared ledger.[7, 8] Other parties interested in this network can receive roles and permissions-based access according to rules agreed by the network. A distributed ledger can be defined as a generalized version of a blockchain, not inherently tied to a particular cryptocurrency, for instance the blockchain underpinning Bitcoin. Figure 2 demonstrates how partners access the blockchain and how information access flows between them.

Features of a Blockchain

The design components of the blockchain that enable electronic health records in a payer-provider-patient model include links and keys, front-end development, verifications, and the subscription model.

Links and keys

Hosting data directly on the blockchain can make it bloated and increase its size very quickly. To avoid that, only access-links are stored on the blockchain. This also limits who can access health information, with the electronic health record shared only between interested parties by using these access links. These are then shared as embedded and encrypted links within transactions to the blockchain. The links only become activated and accessible to users who have the appropriate private key to match the public key hash. This type of encryption scheme has been deployed at large in the Pretty Good Privacy (PGP) program (a popular program used to encrypt and decrypt email over the Internet, as well as authenticate messages with digital signatures and encrypted stored files), and a similar scheme would be very successful here. The actual data associated with the EHR can be secured in the cloud and only made accessible to the user upon requests that can be placed through the same EHR user-interface.

Front-end development

Once a blockchain is deployed to manage EHRs, it becomes the unified and common backbone for digital health. Expanding on the consequences of this development on the future of specialized backend systems, the implication of using this backbone is that each hospital or care provider no longer needs a specific version of databases or software to access patient data. All data are stored in a decentralized manner, and no single entity needs to store a specific portion of the data. There is no longer a need for specific backend protocols or tools, all interested parties just need an underlying protocol for accessing the blockchain to share the same backend. The access-protocol creates distinct user roles (administrators, maintainers, patients) and permission groups with varying levels of privileges. The other advantage is that the development cycle for health data software is simplified significantly, because now the only focus is on the front-end software that accesses the blockchain. Even though different parties have differing levels of access, all user-roles and permissions can be built into the front-end. This also streamlines the costs of future maintenance and development.

Verification

A key hallmark of blockchain-based technologies is the need for verification and consensus from the entire network, and in this case, especially from the interested parties. This rigorous verification ensures that no edits can be made without accountability and visibility to other parties. The verification itself is automated and very rapid; all the parties can approve any changes in a matter of minutes.

This can become even more relevant for health records, because any updates or additional information becomes instantly available to other members on the blockchain. Trusted verification of electronic health can open new avenues of using intelligent agents for anonymous data reporting to track population statistics and trends. Obtaining anonymous data on prescribed medication can help physicians and public health organizations such as the Center for Disease Control and Prevention better understand the over-prescription of antibiotics or pain medication in certain communities.[9]

Subscription model

Maintaining active links to health data on the blockchain can be tied into a subscription model for the patients. Within the subscription, premium features can be added to make it more lucrative, such as instant release of records or free notarization of records on the blockchain such as immunization status. This model allows the blockchain to operate with some level of sustainability for hosting the links, and assist in verification of patient identity. The links can be deactivated after a period of time depending on the subscription ordered by the patient or purchased by the insurance provider.

Current verification schemes embedded in blockchain grant provable integrity of the linked data without any knowledge of the contents. This blockchain can become a platform for running bots and intelligent agents that utilize verified data to provide services at a fraction of the cost compared to traditional services. The auxiliary services such as notarization can provide support for integration of new verticals on top of the front-end EHR. In this manner, new revenue streams can be incorporated within the blockchain.

Hyperledger Project

It is useful to examine an open source project that is creating a distributed ledger for cross-industry use-cases. The Hyperledger Project is a new initiative under the umbrella of Linux Foundation to create an open source enterprise-grade blockchain. It is a collaborative effort towards the development of a distributed ledger that can be used for a variety of use cases to carry out transactions and support advanced features such as smart contracts.[10]

The peculiar open model of development adopted by the Hyperledger Project empowers developers from diverse backgrounds to work on specific areas akin to their level of proficiency. These developers contribute code in small segments that make up a commit (i.e., making a set of tentative changes permanent). This code is then reviewed and tested by other developers and the maintainers in the community. The code is approved when it obtains the “Looks Good To Me” (LGTM) message from the maintainers. Now, the commits can be merged into the main repository in the form of pull requests.

Currently, the Hyperledger Project is guided by a Technical Steering Committee, which includes developers from large corporations that are interested in exploring the commercial viability of the blockchain within the context of products they offer. Over the past few months, two particular implementations of distributed ledger technology have been incubated into the project: Fabric, developed by the joint collaboration between IBM and Digital Assets, and Sawtooth developed by Intel.[10] The project is in very early stages; however, as the ecosystem stabilizes, the core products developed by IBM, Digital Assets and Intel will begin to see external contributions.

This open development model offers incentives to both companies and individual developers to contribute code moving the project forward. Individual developers might want to implement features that satisfy a personal need or fix a bug that only applies to them; however, the solution would reduce the number of bugs in the project and benefit everyone involved. Companies might be only interested in maintaining a small portion of the entire project, but they benefit from advances and contributions made to other parts of the project. The corporate maintainers operate in an open model and use this open source code in their own products. They immensely benefit from bug fixes or feature implementations contributed by the community, because it improves the quality of their product as well. The open-source model proves to be the most effective when operational or security bugs surface. In an open source project, these bugs are often resolved in less than a day with solutions called patches. This speed of development can hardly be attained within a single corporation.

The revenue generated by companies involved in open source projects is from offering deployments of the code in a ready to use environment. Usually, it is a laborious project, getting the initial development environment setup to run code; but these deployments make it available with one click. For a development team with limited resources, it is much more effective to start in a unified base environment and dedicate their time implementing new features. These deployments often make it incredibly easy to get started and get running. A perfect example is the Fabric blockchain. Eventually it will be offered through Bluemix® (IBM's cloud console) and deployed with a single click on Bluemix® using the browser. This ease of use translates to the revenue stream.

The key insight is that an open model allows for the incentives of individual developers and companies align in a positive manner. The irony is that individuals who contribute significant code to the project end up getting hired by one of the companies involved in the project to maintain the code full time. These incentives benefit the project as a whole and ensure successful long-term product development.[10, 11]

Opportunities and Challenges

The unification of data shared by health services is difficult for a developed society with a sophisticated healthcare system; and the reason is simple: The existing software is entrenched deeply within the infrastructure. Accordingly, it is important to discuss the challenges ahead and new opportunities made possible by the use of a blockchain.

The biggest challenge to architectural overhauls in healthcare IT is the legacy systems that have become a standard part of operations for providers. Currently, some of the largest financial institutions are running private trials involving the blockchain and discovering some implementation challenges. These issues must be resolved long before a viable product enters the healthcare market.

Even though a few blockchain-based solutions are commercially available there is a chasm of uncertainty about their performance and operation with patient data. The issue of regulatory barriers remains unclear, although there is some indication that the blockchain might be compatible with HIPPA (Health Insurance Portability and Accountability Act) regulations. Legislative agencies need a complete and coherent understanding of distributed ledgers in the context of electronic health records; and Congress is already making strides in that direction.

The use of blockchain also enables a reduction in overhead costs, particularly for development and maintenance of legacy health record systems. An obvious benefit of using the blockchain is increased security for data transport. The blockchain works based on consensus, preserving the integrity of the entire system. Several features function as checking mechanisms, including timestamps and multiple-node verification of edits made to the health record.

A subtler feature arises from the decentralized nature of the blockchain: denial of service attacks would be computationally very difficult. There is no single point of failure for attackers to target, but instead a decentralized network involving multiple machines that powers the exchange of data. The blockchain also allows for high-grade end-to-end encryption schemes so that only the holder of the private keys can access the full records. Successful deployment of blockchain for health records allows data transition between related parties in an efficient, consensus-based, seamless manner.

Closing Remarks

Fostering the integration of blockchain in the healthcare space will cause significant transformations in the way electronic health data are processed and made available to patients. The cross-talk between the interested parties will allow for more advanced networks to be built on the blockchain—for instance a health-rewards program to positively reinforce lifestyle changes. Such a program can enroll patients with metabolic diseases that can be managed with a healthy lifestyle. The auxiliary services will expand across several industries to integrate new intelligent agents to automate and streamline healthcare processes. A commitment to investment in blockchain-based healthcare will bring about new business models and help improve existing systems.

Tags: antibiotic, blockchain, Bluemix, CDC, Center for Disease Control and Prevention, cloud, commit, compatibility, consensus, cryptographic, decentralized ledger, Dhillon, Digital Assets, digital health, digital signatures, distributed ledger, EHR, electronic health record, fabric blockchain, health care, healthcare, health information technology, HIPPA, hyperledger project, IBM, insurance, interoperability, Looks Good To Me, pain medication, payer-provider-patient model, practice management, Pretty Good Privacy. public key hash, Sawtooth, verification

Dr. Dhillon is a research fellow at the University of Central Florida, at the Institute of Simulation and Training, where he is studying implementation of emerging technologies into digital health with a focus on Ethereum and smart contracts. Institute for Simulation and Training. The address is Partnership II & III Buildings: 3100 Technology Pkwy, Orlando, FL 32826. Correspondence should be addressed to: dhillonv10@knights.ucf.edu.

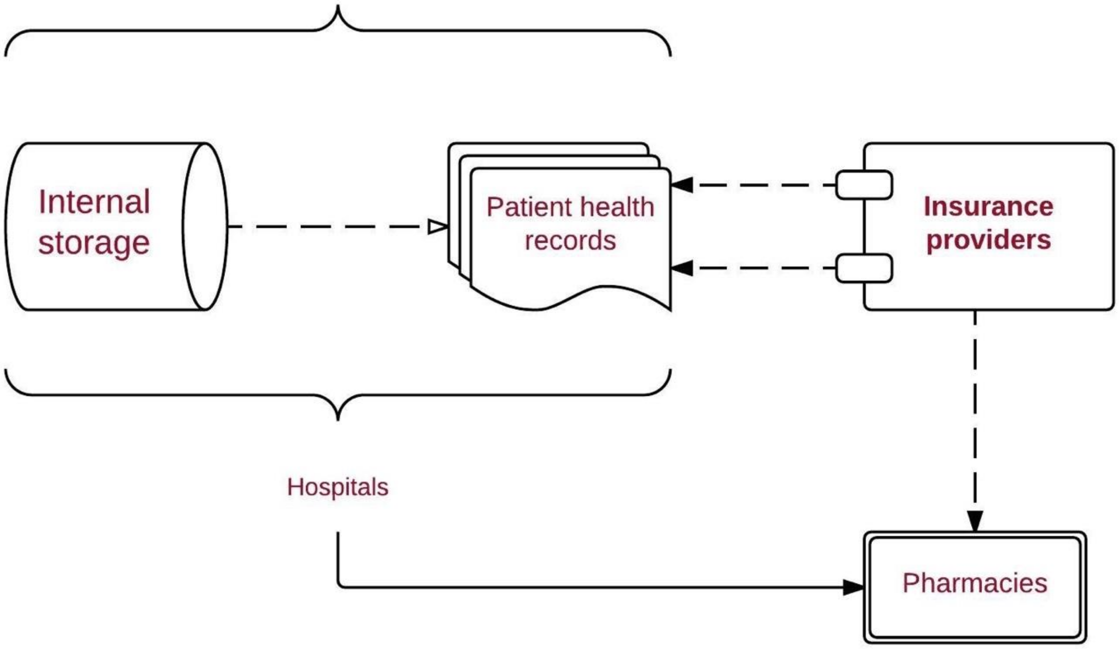

Figure 1. Current view of the digital healthcare landscape

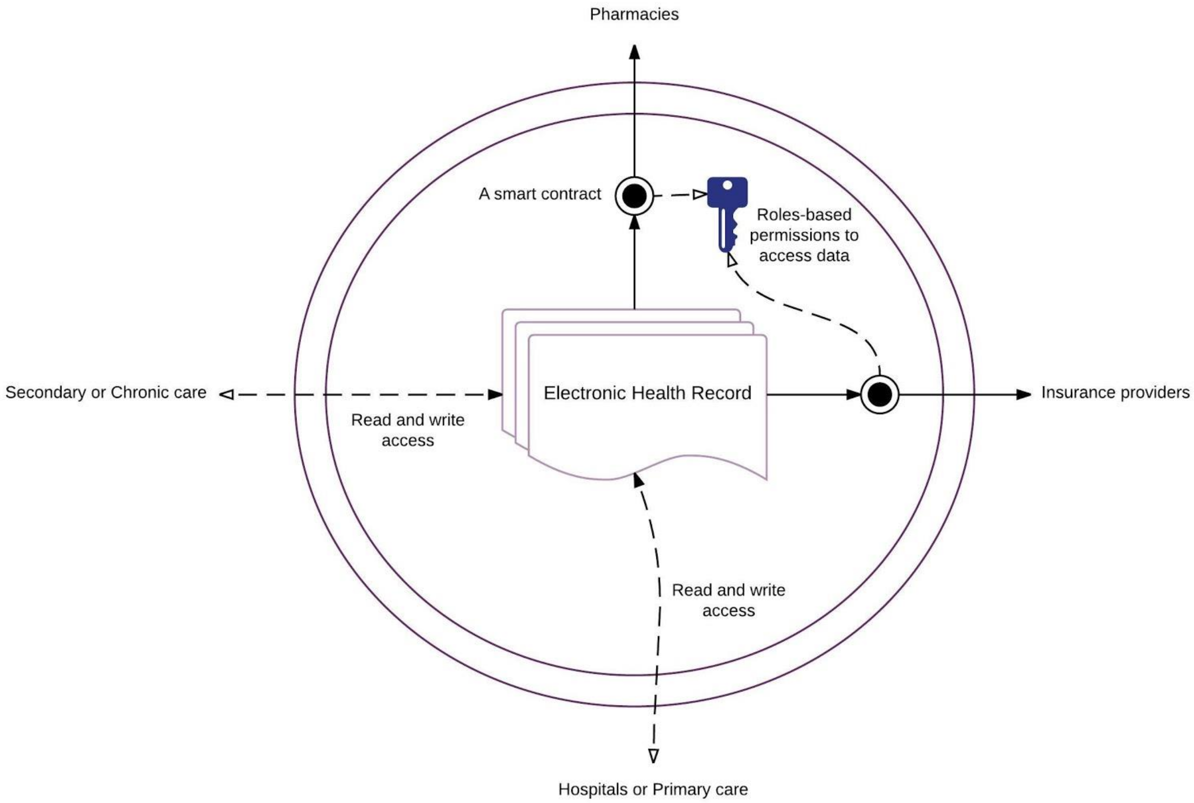

Figure 2. An overview of the Healthchain. In this scenario, hospitals or primary care facilities have unrestricted access to the electronic health records (EHR). The EHR can be edited directly by hospitals, once approved by other members of the distributed ledger. Secondary or chronic care facilities for patients also have unrestricted access. The other two parties in this healthchain are bound by roles and permissions, which are executed on the ledger through a smart contract.

References

1. Jensen, P.B., Jensen, L.J. and Brunak, S., 2012. Mining electronic health records: towards better research applications and clinical care. Nature Reviews Genetics, 13(6), pp.395-405.

2. Häyrinen, K., Saranto, K. and Nykänen, P., 2008. Definition, structure, content, use and impacts of electronic health records: a review of the research literature. International journal of medical informatics, 77(5), pp.291-304.

3. Jha, A.K., DesRoches, C.M., Campbell, E.G., Donelan, K., Rao, S.R., Ferris, T.G., Shields, A., Rosenbaum, S. and Blumenthal, D., 2009. Use of electronic health records in US hospitals. New England Journal of Medicine, 360(16), pp.1628-1638.

4. Beasley, J.W. and Sinsky, C.A., 2014. Electronic health records. Annals of internal medicine, 161(9), p.680.

5. Hripcsak, G. and Albers, D.J., 2013. Next-generation phenotyping of electronic health records. Journal of the American Medical Informatics Association, 20(1), pp.117-121.

6. Blumenthal, D. and Tavenner, M., 2010. The "meaningful use" regulation for electronic health records. New England Journal of Medicine, 363(6), pp.501-504.

7. Mainelli, M. and Smith, M., 2015. Sharing ledgers for sharing economies: an exploration of mutual distributed ledgers (aka blockchain technology). The Journal of Financial Perspectives, 3(3), pp.38-69.

8. Walport, M., 2016. Distributed Ledger Technology: Beyond Blockchain. UK Government Office for Science, Tech. Rep, 19.

9. Dixon, B.E., Gibson, P.J., Frederickson, C.K. and Rosenman, M., 2014. Measuring Population Health Using Electronic Health Records: Exploring Biases and Representativeness in a Community Health Information Exchange. Studies in health technology and informatics, 216, pp.1009-1009.

10. K. Christidis; M. Devetsikiotis, "Blockchains and Smart Contracts for the Internet of Things," in IEEE Access , vol.PP, no.99, pp.1-1. doi: 10.1109/ACCESS.2016.2566339

11. Vukolic, M., 2016. Eventually Returning to Strong Consistency. Data Engineering, p.39.