Peer To Peer Technologies for Health Information Exchange

Michael Dufel
Ipseity Solutions
August, 2016

Introduction

Blockchain technology has been the topic of much interest due to the success of Bitcoin. Bitcoin illustrates the concept that a trusted computer system can be built from the coordination of untrusted participants without any centralized trust system. The success of Bitcoin has translated to an immense amount of research into the applicability of Blockchain based systems for uses beyond digital currencies. Perhaps the most prominent of these efforts is the rise of “smart contracts” in which third parties can verify the terms and conditions of an arrangement. The smart contract efforts by Ethereum promise to extend the functionality of Blockchain to non-currency use cases.

This paper will provide an overview of the general properties of blockchain technologies, the challenges of healthcare information exchanges, and the role of the blockchain in improving data sharing in health information exchanges. An in-depth discussion of how blockchains work can be found from many sources and will not be a focus of this paper.

Blockchain

At its core, a blockchain is a record of ordered discrete events in time. Some refer to a blockchain as a distributed database, but this is a misleading view. The blockchain indeed must store some data, but it is not suited to be a general purpose store of arbitrary data. The driving design requirements behind a blockchain is that these ordered events have to be agreed upon and also trusted by an untrusted community of participants. To date, these primary design requirements have resulted in blockchains being unsuited to storing large data sets and the long term scalability is an open question.

By itself, a blockchain provides little in the way of useful properties. It is simply an ordered chain of transactions that can be easily verified cryptographically as being authentic. Nothing inherent in a blockchain details how the ledger entries are added, by whom, and what those entries mean. Additionally, there is nothing inherently peer-to-peer or fault tolerant about this structure. While the most popular blockchain applications are indeed peer to peer, it is not a necessary property of the blockchain itself. Bitcoin may have invented the blockchain data structure, but the Bitcoin protocol is much more than just a data structure.

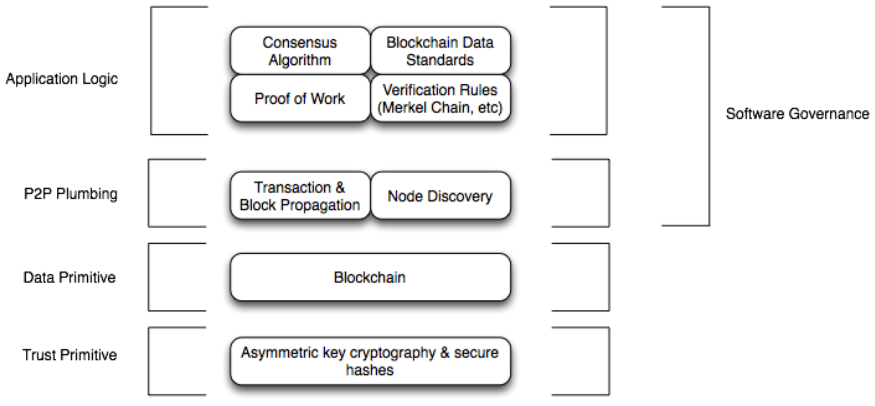

Just as with traditional centralized applications, there is a conceptual

architecture. The foundations of Bitcoin are cryptographic primitives and the blockchain structure. Layered on top of those foundations are simple P2P infrastructures for propagating ledger entries and discovering other participants in the network. Many of the desirable properties of blockchain systems are actually properties of other P2P algorithms and are not directly related to the blockchain itself. For example, the resiliency of the network to dynamic changes in topology is a property found in many peer to peer algorithms and is also a critical feature of the internet routing system.

What Makes Bitcoin Successful

Incentives and Trust

Bitcoin is a success because its algorithms are aligned with the incentives of the participants, and the software algorithms in Bitcoin are designed to reward those that provide value to the network. In Bitcoin terms, an army of independent third parties are needed to check the validity of the ledger entries. Without large numbers of independent third parties (miners), it is possible for one participant to maliciously modify the ledger. The bitcoin algorithm ensures that there are strong profit incentives for the miners, and this ensures that there are many miners that will compete, enhancing the security of the ledger.

In addition to primal economic incentives, there is a question of why a decentralized currency system is even desirable to begin with. Bitcoin exploits the lack of inherent trust in traditional money systems that are controlled by central institutions. Governments control the central banks, and the central banks control the supply of money and indirectly, the value of that money. Another common function of governments is the control and monitoring of the electronic movement of money. The desire to avoid being monitored by government leads inherently to a reliance on a cash economy and the associated risks and problems dealing with physical money. There are also concerns regarding the suitability of cash to act as a long term store of value. Governments can control the supply of money and can effectively reduce the value of any currency by printing more of it if need to pay it’s bills. The long term stability of a currency is therefore a function of trust in the financial stability of the government that issues the currency. This has led people in countries with unstable and/or untrusted governments to convert from their local currency into “hard” currencies like the US Dollar or hard assets such as gold and silver. These are often physical items that need to be protected from theft and in some cases the mere possession of these items can be illegal. Converting a local currency into hard currencies in electronic form is difficult because it may require opening a bank account in a foreign country, which is either illegal or requires expensive travel. Bitcoin eliminates all these problems by providing a cap on the number of bitcoins (money supply) that can be created, low to zero transaction fees, independence from government controlled entities(Central Banks), and a system of storage that does not require money to protect.

Cost

Bitcoin is also innovative in that it is possible to operate a complex network over consumer grade hardware. By expecting that any other participant can and will fail, the software overcomes the need for hardware redundancy. The challenges to high availability and data redundancy is a large driver in the infrastructure related costs in centralized systems. Bitcoin is massively data redundant in that every full node stores a full copy of the blockchain. Due to the massive amount of data replication, Bitcoin is both highly available and the data is highly resilient. For example, it is not possible for a portion of the ledger to go missing due to a large number of nodes leaving the network. This means that low cost hardware can be used on the network. A large driver of computing cost lies in the complexity of maintaining a highly available and resilient system of storage. It should be noted that due to the economic incentives that Bitcoin offers to miners, there is a mining arms race that is leading to the use of large amounts of computing resources to mine blocks. This arms race for mining can be eliminated in a consortium based blockchain. Alternatively, the Bitcoin “proof of work” algorithm could be substituted with a “proof of stake” algorithm that does not require expensive hardware for mining.

Blockchain Scalability Limitations

Bitcoin solves important currency problems, but it does have some downsides and limitations. Bitcoin has a fundamental limit of handling no more than 7 transactions per second. This is in contrast to the Visa network that can process many thousands of transactions per second. Beyond its limitations in handling large numbers of transactions, it does not handle large amounts of data in an efficient manner. The larger the blockchain, the costs to maintain the hardware and the network to move the data around will grow. If the blockchain is too large, the number of network participants will drop due to the costs needed to operate a node. This could compromise the trust in the blockchain itself. These limitations are also present to a degree in Ethereum based blockchain applications.

Blockchain Investment Drivers

Blockchain systems are finding investment due to the potential to replace third party entities such as transaction clearinghouses or intermediaries providing trust. This can be driven by a lack of trust in a clearinghouse, by the costs involved in operating such a clearinghouse, or both.

Beyond simple currency transactions, there is also a large amount of interest in smart contract blockchains such as Ethereum. Smart Contracts offer a way for a decentralized

third party to verify the terms and conditions of a piece of executable code known as a contract. Due to the very low transactions fees for this service, Ethereum may rise as a fundamental enabler of distributed applications that have value, but not enough value for the creation of a traditional centralized structure.

Other Peer To Peer Technologies

Blockchain based systems such as Bitcoin and Ethereum are innovative because they solve the very hard problem of assembling a trusted system from untrusted parties. However, for networks that don't require that extreme level of trust, there are other peer to peer technologies that are worthy of examination.

BitTorrent

BitTorrent is a fully Peer To Peer file sharing system. It's notable properties is that files are split up into smaller chunks and distributed into the network. This has the advantageous property of distributing network traffic for popular file downloads, while also providing redundancy and high availability. For large amounts of data, BitTorrent is vastly more scalable and efficient than a Blockchain, while still providing for cheap consumer grade hardware, high availability and redundancy. While BitTorrent is well known for serving pirated music and videos, it has also found use as a content distribution system for commercial companies. BitTorrent can be operated as a private network and in the context of healthcare, would need to be secured to prevent unauthorized downloads.

Distributed Hash Tables

Distributed Hash Tables are a distributed version of one of the most primitive and yet powerful data structures in computer science. Similar to a blockchain, a DHT is a primitive building block for higher purpose P2P applications. For Healthcare, a DHT can be the building block for implementing distributed Master Patient Indexes, search indexing for data, storing files, and implementing a distributed message pub/sub system. DHT's evenly spread information among the various nodes in the network, and can provide data redundancy and high availability.

Healthcare Information Exchange Overview

Business Incentives

To date, there has been a misalignment in the incentives for sharing electronic health records. The greatest benefits of medical records are bestowed on different parties then

the ones that have to pay for them. Electronic Health Record systems are very expensive to purchase, operate, and maintain and those costs are paid for by healthcare providers. These providers operate as a for-profit business and in most cases, the benefits to any individual provider to share data for the benefit of the patient do not exceed the cost of the system. This fact has hindered the use and adoption of EHR systems. To force a business case, Medicare leveraged its large influence in the healthcare industry to introduce a carrot & stick payment incentive approach called Meaningful Use. This program has been highly effective in driving up the adoption of both standards and use of EHR systems, but the fundamental disconnect between those who invest in technology and those who receive the benefits still remains. The practical effect of this tension will be that providers will tend to do only what the government dictates them to do, while creating unnecessary tension between the government and care providers.

Cost

Health Information Exchanges have common infrastructure needs:

- Master Patient Index
- Provider Registry or Provider Index
- Document Search/Storage/Retrieval (Query based exchange)
- Secure Direct Messaging such as DIRECT (Directed exchange)
- Event Routing (Ensuring events such as Admit-Discharge-Transfer are routed to the right parties)
- Lab and Prescription Orders & Results

To date, some these services have been provided by regional HIEs and private organizations such as SureScripts. This mix of public and private exchanges are not all connected to each other, which means that the data is still not in a nationwide logical store and is still sequestered in a lot of regional data silos. These HIEs need to solve the same problems, but are implementing solutions independently. This leads to a large amount of waste spent on providing these services in a redundant and yet disconnected manner. These costs are currently being subsidized by state and federal grants, but in the absence of grants, these costs will end up being paid by the individual participants or by taxpayers at large.

Trust

Centralized systems offer many benefits, but they tend to centralize the control of that information into the hands of a few. They also offer the potential for abuse and attack by hackers. In the United States, there is also a great resistance to the concept of socialized or government run healthcare. This resistance is a primary reason why there is no government operated centralized health information repository.

A New Paradigm

Fundamentally, the main impediments to wide scale sharing and exploitation of Electronic Health Records involve Cost and Trust. If there is adequate trust, there could be a single centralized data store for patient records. The lack of trust translates directly into separation of data into large numbers of controlled silos. We believe that the trust issues can be addressed by a consortium managed peer to peer network with adequate privacy and data security controls that allow participants to have a voting voice in rules and operation of the network as well as individual control of the data they provide. The membership in a consortium means that there is an effective voice for stakeholders to evolve the functions of the network. The use of privacy and security controls, especially strong auditing provided in part by a blockchain smart contract system, will provide trust the that network is not being abused and that any abuse can be identified and corrected. Such a system would combine the benefits of a centralized health record system with the trust properties of a distributed system. We believe that the cost problem can be completely resolved by the combination of a low-cost peer to peer network, and the introduction of a data marketplace that is able to funnel money in an efficient manner from those who can exploit the data to those to provide it.

Fostering Incentives For Better Outcomes

The aligning of incentives among the distributed and independent stakeholders is crucial to obtaining better outcomes. Insurance companies, researchers, public health agencies, and accountable care organizations, all directly benefit from access to data, and aggregated data in particular. With the lack of a centralized storage mechanism, getting access to aggregated data is a difficult and expensive task as data must be retrieved from a variety of sources and correlated. With a logically centralized source of data, the cost of obtaining high quality data will be dramatically reduced.

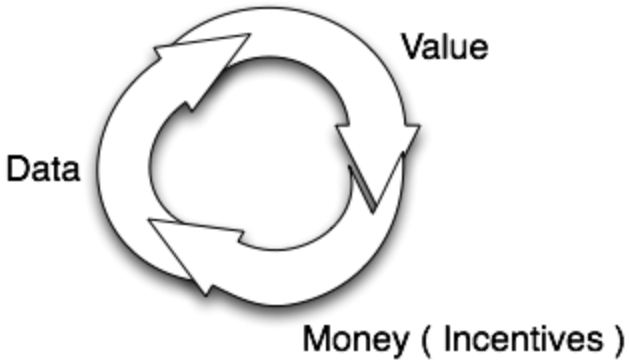

Even though the costs to operate a peer to peer network are low, they are not insignificant to the providers. Even if hardware and software costs are minimized, there are still costs associated with data entry. By providing a common data marketplace mechanism, those that can benefit from healthcare data (in accordance with state and federal laws) can both compensate the providers for collecting and publishing the data, and provide funding for the sustaining and governance of the network.

It should be noted, that this data marketplace would be necessary to fund the management of the network, reimburse the providers for the costs associated with providing the data, and provide a market based signaling mechanism to encourage the publishing of useful and exploitable data. The marketplace would not be intended enrich the participants as this would likely violate HIPAA regulations.

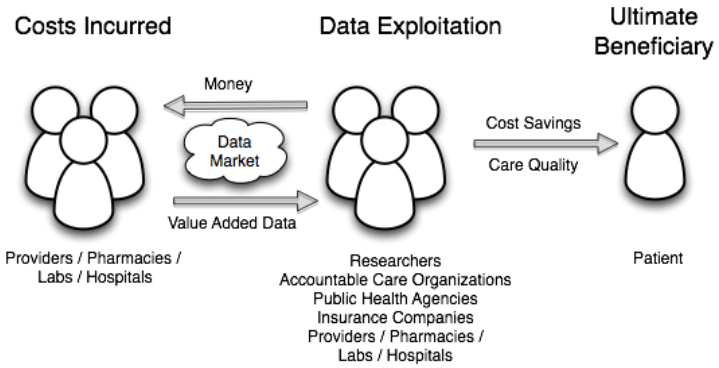

Reducing Costs

We propose the creation of a peer to peer solution that leverages the strengths of different technologies in a synergistic manner. Many of the common services in traditional HIEs can be implemented in a peer to peer manner.

	Implementing Technologies		
Services	**Blockchain**	**DHT**	**BitTorrent**
Master Patient Index	NO	YES	NO
Provider Registry	NO	YES	NO
Data Search	NO	YES	NO
Data Storage & Retrieval	NO	NO	YES
Direct Messaging	NO	YES	NO
Event Bus	NO	YES	NO
Data Marketplace	YES	NO	NO
Consent & Compliance	YES	NO	NO

As shown in the table above, blockchain provides only one component of an overall Peer To Peer solution. A blockchain based smart contract system would yet play a critical role as a provider could create a smart contract associated with his portion of the patient records. The retrieval of information would be governed by the terms of the contract. These controls can allow for interactions among different parties like between patients and providers for consent purposes. Furthermore, the ability to control information retrieval from the network via a smart contract can increase the level of trust

in the network and result in greater information sharing overall. No longer does data have to reside in individual silos for the purposes of control and security.

Security Requirement	Solution
Integrity	Records hashed and signed by record creator
Confidentiality	Data at rest encrypted using symmetric key cryptography such as AES. Data in transit secured using TLS.
Auditing & Accounting	Blockchain based smart contract
Availability	P2P Networks are Highly Available. Data is both replicated and distributed across many peers.
Authentication	Consortium authenticates covered entities and other HIPAA allowed entities. Providers / Covered Entities authenticate patients at point of care.
Authorization	Blockchain based smart contract

Standards and Interoperability

Data Interoperability

Data will always be heterogenous and available in different formats and standards. These formats and standards are continuously evolving and the network needs to be logically separated from these standards in order for it to evolve gracefully. The Peer to Peer network would not impose requirements on the format of the data stored and would function like a “schemaless” database or NoSQL data store. This ensures that data of all types can be stored and the details of interoperability can be sorted out by the participants. A data marketplace driven by a blockchain or smart contract system could be used to incentivize the inclusion of data that adds value to others. Thus, the network would incentivize good behavior and interoperability. Even a large data store with different data formats is still valuable. Most data does follow some standard and can be made visual through a web browser as the lowest common denominator of sharing.

API Standards

A Peer To Peer network would be primarily a “back end” service infrastructure. As such, it is expected that existing IHE and HL7 API standards such as FHIR, and XDS.b can act as a standard “front end” for interoperability with current EHR systems.

Privacy & Consent

The same legal principles that apply to currently existing private and public HIEs would apply to a consortium based peer to peer HIE. It would be the responsibility of the network participants to follow the laws of their jurisdiction with respect to what can be shared on the network and the collection of patient consent. It should be noted, that with a peer to peer HIE, a patient's consent can now be stored and shared on a nationwide basis, which reduces the burden on the patient for managing consents for every provider they visit. Additionally, with a national HIE, there will be the ability for the patient to see who has accessed their records and will be in a good position to spot fraud and abuse. With respect to a blockchain based smart contract system, it is also entirely possible for patients to create their own "smart contract" that governs access to their records.

Summary

We believe that a smart contract system such as Ethereum can act as a key component in a peer to peer health information exchange by providing the mechanism for a trusted peer to peer data marketplace to incentivize data sharing. A smart contract system can increase trust in the network by providing privacy and security functions by acting as a policy decision point controlled by the data holders and the patients.

References

[1] S. Nakimoto, Bitcoin: A Peer-to-Peer Electronic Cash System, https://bitcoin.org/bitcoin.pdf

[2] Vitalik Buterin, A Next-Generation Smart Contract and Decentralized Application Platform
https://github.com/ethereum/wiki/wiki/White-Paper

[3] Niam Yaraghi, A Sustainable Business Model for Health Information Exchange Platforms: The Solution to Interoperability in Health Care IT
http://www.brookings.edu/research/papers/2015/01/30-sustainable-business-model-health-information-exchange-yaraghi

[4] Improving Health Through Interoperability and Information Sharing
http://bipartisanpolicy.org/wp-content/uploads/2015/11/BPC-Improving-Health-Interoperability.pdf

[5]Andrew Quentson, Ethereum Announces Unlimited Scalability Roadmap,
https://www.cryptocoinsnews.com/ethereum-announces-unlimited-scalability-roadmap/

[6] Bram Cohen, Incentives Build Robustness in BitTorrent,
http://www.bittorrent.org/bittorrentecon.pdf

[7] Frank Dabek, A Distributed Hash Table,
https://pdos.csail.mit.edu/papers/fdabek-phd-thesis.pdf

[8] Jennifer Bresnick, How Health Information Exchange Models Impact Data Analytics, http://healthitanalytics.com/news/how-health-information-exchange-models-impact-data-analytics

[9] Key Challenges to Enabling Health Information Exchange and How States Can Help, https://www.healthit.gov/sites/default/files/state_hie_evaluation_stakeholder_discussions.pdf