ELSEVIER

International Journal of Medical Informatics 62 (2001) 27–40

International Journal of Medical Informatics

www.elsevier.com/locate/ijmedinf

The Personal Internetworked Notary and Guardian

Alberto Riva [a], Kenneth D. Mandl [a,b], Do Hoon Oh [a], Daniel J. Nigrin [a], Atul Butte [a], Peter Szolovits [c], Isaac S. Kohane [a,*]

[a] *Children's Hospital Informatics Program, Division of Endocrinology, Children's Hospital, 300 Longwood Avenue, Boston, MA 02115, USA*

[b] *Division of Emergency Medicine, Children's Hospital, Boston, MA 02115, USA*

[c] *Clinical Decision Making Group, Laboratory for Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139, USA*

Received 10 May 2000; received in revised form 20 November 2000; accepted 4 December 2000

Abstract

In this paper, we propose a secure, distributed and scaleable infrastructure for a lifelong personal medical record system. We leverage on existing and widely available technologies, like the Web and public-key cryptography, to define an architecture that allows patients to exercise full control over their medical data. This is done without compromising patients' privacy and the ability of other interested parties (e.g. physicians, health-care institutions, public-health researchers) to access the data when appropriately authorized. The system organizes the information as a tree of encrypted plain-text XML files, in order to ensure platform independence and durability, and uses a role-based authorization scheme to assign access privileges. In addition to the basic architecture, we describe tools to populate the patient's record with data from hospital databases and the first testbed applications we are deploying.

Keywords: Cryptography; Electronic medical records; Life-long medical records; Security; World-Wide Web; XML

1. Introduction

In an era of increased deployment of electronic medical records, patients still have remarkably little access to their own records. As many who have tried to obtain copies of their own record for use in another health-care institution (or sometimes even in a different department within the same health care institution) have experienced, patient control of their record is minimal at best. These diffculties are compounded by the constantly changing affiliations between patients and providers as health-care plans change the providers with whom they contract with for various patient populations. There have been efforts to allow institutions to share their records electronically over networks such as the Internet [1], but because of the competi-

* Corresponding author.
E-mail address: isaac_kohane@harvard.edu (I.S. Kohane).

1386-5056/01/$ - see front matter
PII: S1386-5056(00)00136-2

tive nature of health-care delivery, there has been very little incentive for health-care institutions to support such broad sharing of patient records [2]. Furthermore, the portion of the patient's health history that is captured by institutional medical-record systems represents only a small fraction of the available information about the patient's health history [3]. Perhaps just as problematically, there is little provision in these institutional efforts for a system of reporting that could conceivably meet national research or public-health goals.

With increasing access to the web for a variety of everyday commercial and informational needs, and the commoditization of storage, computation and communication, the possibility of a record controlled by the patient appears to be feasible. Several systems have been proposed and developed, some of which will be described below, but few, if any, meet the combined goals of adequate protection of confidentiality of patient data, portability and security of the data, integration with institutional health-information systems, patient control of the data, and use of selected portions of the patient record for research and public health. Implicit in these goals is the integration of the system with existing institutional information systems (e.g. hospitals and public health authorities) using standard protocols, and storing the records using Internet and Web infrastructure that is likely to endure for a long time, rather than proprietary software.

The Portable Internetworked Notary and Guardian project (PING) is a system currently under development by the Children's Hospital Informatics Program, whose main goal is to develop a secure, distributed and patient-controlled repository of sensitive information on the Web. In this paper, we analyze in detail the motivations for this project, describe the resulting technological and architectural requirements, and present the design and implementation of the first prototype of the PING system. Finally, we describe some applications of PING that are currently being implemented or designed.

2. Motivations

Despite being the legal owner of his or her data, the patient has very little control over how these data are stored, accessed, and used by the different parties involved. The medical record is usually fragmented and scattered among different institutions, according to the medical history of the patient. The information contained in the fragments is usually stored using different, incompatible formats that are often tied to applications whose longevity is not guaranteed. Although multiple solutions for data sharing across and within institutions have been proposed, these solutions show some common limitations. First, they require considerable agreement between various institutions or departments on which data may be exchanged and what common data models may be. Historically, even with agreement on common messaging protocols, such data-sharing efforts have been arduous and problematic [4]. At the same time, patients are increasingly mobile, both in the sense of actually moving from location to location, but also in changing employers and/or health plan, which results in their care changing from provider and institution relatively frequently. Therefore, if patient records are to be truly portable and are to be easily transferable from one set of providers to another, dependence on a particular set of institutional agreements seems both a fragile and untimely solution.

Furthermore, these inter-institutional agreements do not provide for maintaining a patient accessible record. Although health-

care systems increasingly afford patients a look at some subset of their records, no institution is providing access to the entirety of a patient's records, nor are there any mechanisms in place to allow the transfer of that record with the patient from one set of providers to another set. Additionally, it is quite clear that under a variety of contractual arrangements between patients and their providers, certain subsets of their record may be released to third parties for quality control, fraud detection, but also for actuarial and managed care analysis. There are a number of reasons why a patient may not want their data fully accessible for these purposes, not least of which might be the confidentiality of particularly sensitive information such as sexually transmitted diseases, psychiatric problems and social disorders. Because of the nature of the contractual arrangements between the patients and the providers, and the providers and third parties, there may be very little that patients can do under current arrangements to limit some of the current existing access to what they would consider confidential information. This also suggests the need for alternative solutions to provider-based institutional health-information systems.

Another overall consideration is the longevity of patient data in electronic medical record systems. We have already seen, in the brief history of electronic medical information systems, the expense that it takes to maintain legacy health-information systems that contain patient specific data and that with time become increasingly incompatible with state-of-the-art information technologies. Although compatibility with modern information technologies can be maintained, it is only at ever increasing expense and effort. Recent experience with addressing the Year 2000 problem should be fairly convincing of the scope of effort required to maintain large health-information systems current with state-of-the-art technologies. Finally, in a durable personal medical record, all transactions (data entries, deletions, annotations, etc.) should be irrepudiably signed and timestamped. In other words, it should always be possible to know when a certain operation was performed, and by whom.

3. Requirements

The goal of the PING project is to design and implement a secure, distributed, user-controlled data-storage system. In this section, we outline the basic requirements of the PING architecture, and the technological and architectural solutions that we adopted in order to meet them. The basic components of the architecture are the PING *database* (a set of repositories on the Internet where PING records are stored), the PING *server* (a software system that provides access to the PING database), and any number of *agents*, that interact with the PING server to manipulate the information contained in the PING record. The agents can be autonomous, or act on behalf of a human user.

3.1. Data representation and storage

In order to ensure its longevity and durability, the medical record should be stored in a format that combines expressiveness with high flexibility and that can be easily generated and parsed. A solution tied to a proprietary data format is clearly unacceptable, since it would not guarantee the information contained in the medical record to still be readable in the future. For the same reason, we want the hierarchical structure and access control list of the PING record to be independent of any particular file system, database management system or operating system. This is particularly impor-

tant since we should be able to deploy PING on file systems where the patient has only minimal (or no) privileges to change the security properties of files and directories. In other words, PING should only depend on the lowest common denominator form of distributed storage; one that is not dependent on any particular vendor and is most widely distributed. Static web pages fit this requirement for generality and ubiquity.

The PING architecture must ensure interoperability between the various agents involved in the processing of the medical record. For example, health-care institutions should be able to update laboratory studies; the patient should be able to enter daily dietary information; patient monitors ranging from a glucometer to a portable ECG monitor should be able to write the biophysical data they generate to the record. This implies the existence of a standardized data model and set of vocabularies to express these data in a set of common terminologies. PING must also support the most popular messaging transactions for medical data such as the Health Level 7 (HL7) [5] and the X12 standards, by providing a communications and translation infrastructure to bring information from multiple sources into the PING database.

3.2. Security

An extremely important requirement, one that is at the core of the whole architecture of PING, is that only appropriately authenticated and authorized parties should be allowed to access the contents of the PING record. The PING server should therefore grant or deny access privileges to the record to different agents according to their identity: only agents that have successfully proven their identity should be allowed to operate on the contents of the PING database, and only through actions for which they have received an explicit permission. However, we want the owner of the PING record to be able to store it on a location on the Internet of his or her choice (typically, a web server on which the owner has write access, such as their America On-Line account). Since we have no control over the hardware and software characteristics of the chosen site, this implies that the PING server cannot rely on the PING database being secure. As will be explained in detail in the next section, PING uses cryptography both to authenticate the agents (i.e. proving their identity) and to authorize them to perform the requested actions, by making the contents of the database effectively unreadable to all other agents. Moreover, the data should be irrepudiably signed, so that it can be proven that the contents of any PING record are exactly what was received from the original information source, and no unauthorized modification took place.

3.3. Distribution

PING records should be fully distributed, that is, a PING document might include references to other PING documents on other PING servers throughout the Internet. This is necessary because it provides for the maintenance of very large PING records. For example, a radiology or photographic archive for an individual would take up a lot more storage than the entire textual history of that patient's care, and therefore, it might make sense to dedicate a specific PING site just to support the imaging requirements for a group of patients. Distribution over the net also enables versioning and mirroring mechanisms that are relatively robust.

3.4. Executable procedures

One possible kind of information to store in PING documents would be executable

procedures such as decision support programs to help the patient in the management of their information. For instance, a possible addition to PING might be a program that could process the immunization history of the patient stored in the PING pages and then provide an alert when the patient was due for another vaccination. However, we have elected to provide a very clear separation between declarative representation of patient data and any executable inferences. In fact, PING offers very limited and well-defined provisions for storing and activating executable procedures. The reasons for this are threefold. First and foremost, medical logic and medical procedures or guidelines are highly context-dependent and change over time: today's appropriate immunization guideline may be inappropriate tomorrow. Second, despite multiple efforts over the past several years to formalize the representation of medical reasoning and management procedures for computers, such as the Arden syntax [6] or the Guideline Interchange Format [7], there is yet to be a successful widely accepted knowledge representation for executable medical protocols or procedures. Third, the universe of possible procedures that might apply to the patient data is much larger than the universe of different types of data that we are likely to store in PING, and therefore, their representation and execution is best separated from the actual representation of patient specific data.

4. Implementation

The current prototype implementation of PING is designed to take advantage to the maximum extent possible of existing standards and technologies, to help ensure interoperability and longevity. In this perspective, our first design choice was to implement the system in the World Wide Web environment. The architecture of the WWW, although extremely successful and in these days almost ubiquitous, is actually very simple: it can be viewed as an infrastructure to deliver documents from a single source to a potentially unlimited number of requesters. The information flow is unidirectional, there is no explicit way of obtaining information on the structure of the retrieved documents, and there is little support for sophisticated access control.

In order to implement the functionalities needed to support the PING framework, we propose an evolution of this architecture, aimed at allowing the web to be used as a means to store and exchange information in a secure and controlled way. Our main goal is to be able to use the WWW as a 'secure digital repository' of structured information, and to allow for a bi-directional flow of such information between the repository and trusted external agents (human users, software, hardware devices). The architecture we are going to describe strives to meet the above-described desiderata and is therefore tailored on the requirements for building a distributed, secure, patient-owned medical record infrastructure (however, it is not tied to any particular application domain and could therefore be used to support a wide range of secure on-line transactions). This evolution of the Web does not necessitate any fundamental change in the Web protocols such as http: the process by which the PING server obtains a file containing a portion of the record from a web server is essentially the same that takes place when downloading a binary file instead of an html page from a Web site. Moreover, we can take advantage of the wide availability of http clients (i.e. Web browsers) by 'wrapping' PING commands inside http requests. Upon receiving a

request expressed using the http protocol, the server extracts the relevant commands from the request, performs the desired operation, and replies generating a suitable html page. With very minimal additions to the PING server, we can therefore turn PING into a Web-enabled application, seamlessly integrated with the network, and employing the same user interface that all Internet users are familiar with.

4.1. Data-storage format

With regard to the problem of the encoding of the PING record, the solution we propose is to use a set of eXtensible Markup Language (XML) [8] documents encoded in ASCII format. In addition to the features outlined above, the main reason for choosing XML is that it has been adopted widely throughout the information industries. In particular, XML is one of the languages for which the HL7 data model has been specified [9]; by adopting the HL7 DTD for the medical portion of the PING record, we will therefore adhere to a widely accepted data model for the interchange of medical information.

As stated above, we have also chosen not to encode the hierarchical nature of patient record documents within a particular file and directory structure of any operating system because we wanted the access permissions to these records to be independent of the set of security features provided by any particular operating and file management system. Instead, we have defined an XML DTD to describe a 'virtual' hierarchical, secure file system in an OS-independent way. This allows us to use XML documents to represent both the directories and the objects contained within them. In general, such documents consist of a *header* part and a *data* part. The header part contains meta-information such as the author of the document contents, the times of creation, last access and last modification, and the access rights that apply to them. The data part contains the actual data, or a link to an external document containing it; in the case of an object representing a directory, the data part contains pointers to other directories and/or objects. The whole directory structure can therefore be stored securely together with the PING record data. An additional benefit of this choice is that we do not expose the directory structure of the PING record. This is important because the mere presence of a subdirectory (e.g. psychiatric admissions, or sexually transmitted diseases) might disclose confidential information.

Although the PING architecture allows full functionality using only a flat file system, particular implementations could instantiate the data model and syntax within a proprietary database or file system. The motivations articulated above argue against such a system, which will ultimately limit personal control by the patient over the fate of their record.

A recent New York Times article described an expert panel [10] convened to discuss how best to design a millennial capsule. There was overwhelming consensus that all digital media were currently far from adequate to store any data for any period of time extending into decades. Rather, the panel suggested that analog media be used for long-term storage such as acid-free paper. For this reason, as well as the current lack of guaranteed reliability of any network-based communication media, PING provides the user with two formats of archival storage:

1. Printout in human-readable ASCII of the entire PING record with structured tags. This would allow a clean printout to be readily scanned and parsed into the native XML format of PING. Note that this

printout is an archival form and not the special application printouts that are generated on-demand for patients (e.g. growth charts or immunization histories).
2. Tab-delimited file readable by most word-processing and spreadsheet applications. Although, in many ways, this storage form is a step back from the exibility of the native PING storage, it meets a desirable standard of care in giving patients full control of their record and providing a set of media that are known to last longer than all existing digital media.

4.2. *Authentication*, *authorization*, *security*

Security issues play a central role in the architecture of PING. Storage of sensitive medical and personal data on a publicly accessible web server can only be allowed if strong provisions are in place to ensure that only authorized agents can access and manipulate them. We plan to make heavy use of cryptographic technology, in both software and hardware forms, for this purpose. In particular, we will address separately the three issues of authentication (establishing the identity of an agent with a high degree of certainty), authorization (assigning the correct privileges to the agents), and encryption (enforcing the security policies by preventing unauthorized agents to access the information in PING). In order to comply with the above-outlined requirements of generality, exibility, and platform independence, we did not rely on the security features of the machine hosting the PING database, but we implemented our own security architecture. In doing so, we tried to adhere to widely used cryptographic protocols and algorithms, taking advantage of existing technology and open-source software where possible.

The PING system uses role-based authentication to determine the access rights that apply to an individual agent. Access rights, here called privileges, are described in terms of five atomic operations: Create (add an object to the PING database), Read (access the data part of an object), Modify (change the data part of an object), Delete (remove an object from the database), and Annotate. The last operation refers to the ability to add annotations to any record obtained from any data source, independently of its format and semantics, while maintaining the integrity of the original data: that is, annotations are in themselves first-class objects within the PING framework.

A role represents a set of agents possessing certain privileges over an object in the PING database. A role can represent a single individual (e.g. my friend Ann), an unspecified individual, whose identity can change over time (e.g. my psychiatrist), or a group of individuals (the physicians of an emergency department); moreover, PING defines three special-purpose roles: Owner (the owner of the data in the PING database), Author (the agent who created an object) and Other (a 'catch-all' role to describe privileges that apply to agents with no other role). Note that an agent can belong to more than one role; in particular, all agents belong to the Other role. Note also that the Author of a PING document is not necessarily its Owner: in some cases, the Owner will grant other agents permission to create documents in his or her PING database (e.g. a lab submitting exam results). The PING server maintains a database of agents, listing their identity, their authentication credentials, and all the roles that have been assigned to them. It also records the identity of the owner of each PING database. The information contained in the header part of every PING object tells the system, among other things, who its Author is and what privileges apply to the specified roles. Fig. 1 shows an example of a PING object.

The header part of this object specifies that it was created by the agent identified by the name 'agent1', and that its owner can read it, delete it and annotate it, while its author can read it and modify it; all other agents have read-only access to this object. In the actual implementation, the header contains other elements, not shown here for simplicity. The object data are contained in an external XML file, pointed to by the URL attribute in the Data tag.

The goal of the authentication phase is to determine the identity of the agent requesting access to the PING record with the highest possible degree of certainty. Several different factors influence the choice of a personal identification scheme. By far the most widespread solution in use today is password-based authentication; although very easy to implement, this scheme is generally recognized as offering very weak security. An alternative, software-based solution is to store a secure identifier, such as a private cryptographic key, on the client machine; in this case, however, the drawbacks are that users are forced to access the system from a single machine, and that allowing multiple users to use the same machine poses high security risks. At the other end of the spectrum, we find authentication based on biometric hardware devices, such as fingerprint readers, face-recognition systems, voice-spectrum analyzers, etc. The cost and reliability of these devices vary widely, but some of them (notably fingerprint readers) are becoming increasingly available on low-end hardware. They could therefore become common enough to attain the ubiquitous access that is one of the desiderata of the PING project. While biometric solutions in principle offer the best possible form of authentication, they suffer from a serious weakness: the physiological features they measure (e.g. the shape of fingerprints, the iris pattern) cannot be changed; if the corresponding data are intercepted, the authentication system becomes useless. A possible compromise is represented by devices such as smart-cards or hardware keys; some of them plug directly into one of the I/O ports of a personal computer, while the others require some kind of reader. Possibly coupled with passwords to increase security, these devices represent an effective and affordable solution to the problem of personal identification.

```
<Ping-object name="Example">
  <Header>
    <Author alias="agent1" />
    <Creation-date time="944003712498" />
    <Privileges>
      <Privilege role="owner" read="t" annotate="t" delete="t" />
      <Privilege role="author" read="t" modify="t" />
      <Privilege role="other" read="t" />
    </Privileges>
  </Header>
  <Data type="text/xml" url="data1.xml">
  </Data>
</Ping-object>
```

Fig. 1. XML representation of a PING object.

In the course of another project, the Health Information Identification and DeIdentification Toolkit (HIIDIT) [11], we have also explored various alternatives for cryptographic identification systems that prevent, without the patient's express consent, the tracing of an authenticating identity to a recognizable human identity. Whether or not this technology is used for PING will be determined to a large degree by the confidence that is placed in the encryption of the PING record.

As an example of a scheme to prove an agent's identity, we describe a procedure based on public-key cryptography [12]. We assume that the server possesses, or is able to obtain, a copy of the public key of each known agent. In order to test whether an agent really owns the identity it is claiming, the following steps are taken:

- The server generates a random sequence of bytes of arbitrary length and sends it to the agent.
- The agent computes a digital signature of the sequence using its private key, according to one of the several algorithms available for this purpose (every agent should implement at least one signing algorithm, while the PING server should implement the largest possible number of signature verification algorithms).
- The agent then sends the signature back to the server that can test its validity using the agent's public key.

In order to illustrate the authorization procedure in PING, let us assume that agent A has contacted the PING server and requested to perform the action P on an object O and that it has also communicated to the server its identity and (optionally) the credentials it wishes to use to authenticate itself. The PING server performs the following steps:

1. It creates an initially empty set of *valid roles* RV.
2. It verifies whether the supplied identity is known. If it is unknown, it adds the role 'Other' to RV, and skips to step 4.
3. It verifies the agent's identity by evaluating the supplied credentials, or requesting additional credentials in case they are not suficient. (This step might require some form of negotiation between the server and the agent, in order to determine the form of authentication to use. The server can refuse to authenticate an agent using an authentication scheme it believes to be too weak, even if the agent supplies the correct credentials.) If authentication does not succeed, access is denied, and the procedure terminates.
4. It adds the agent's identity and all the roles that were assigned to it by the database owner to the set RV.
5. It accesses the object O and determines the set RP of roles whose privileges include the action P. If 'Owner' is a member of RP, it is replaced by the identity of the owner of the PING database. If 'Author' is a member of RP, it is replaced by the identity of the agent that created the object O.
6. The PING server determines the intersection between RP and RV. If the intersection is empty, the request is denied, and the procedure terminates. Otherwise, the request is granted.

Note that the above-described procedure decouples authentication from authorization: the ability to perform an action on an object depends on the roles of the requesting agent and only indirectly on its identity.

Finally, encryption techniques are employed by PING to protect the data from unauthorized access and modification. In particular, encryption is used in two contexts:

- The objects constituting the PING database are stored as encrypted files on a traditional web server; in this way, although accessible to everybody (including the PING server), they are not readable. Moreover, the use of cryptographic hash

functions makes it possible to detect modifications to the data.
- Data are sent from the PING server to the agents over a secure, encrypted network link (e.g. SSL, Secure Socket Layer).

It is important to note that the only location where the objects exist in unencrypted form is inside the memory of the PING server, and that the server also stores the key used to decrypt them. The PING architecture, therefore, concentrates all vulnerabilities in a single point, the PING server, making it easier to defend the entire system from possible attacks.

4.3. Performance

Speed, although important, is not a fundamental requisite of the PING architecture, since, in a web-based application, speed is usually determined by the network latency time. Accordingly, we decided to develop the system using a programming language that privileges portability and reusability over speed of execution. We will nevertheless try to identify the components of the software architecture that could reduce the overall response time of an application based on PING.
- Encryption. Public-key encryption algorithms, that are known to be relatively slow, are only used in the agent authentication phase. Data encryption is performed using a symmetric cypher that is much faster. Hardware solutions are also available to encrypt the contents of a file before writing it to long-term storage.
- XML parsing. XML was designed to be easy to parse and to generate, and we therefore believe that the choice to encode data in XML brings about no performance penalties. Moreover, as XML gains widespread adoption, the number of readily available optimized XML parsers will keep increasing, as will the advantages of using a common data representation format.
- Access control. The authorization procedure described above relies on comparing simple identifiers and enjoys low computation complexity. Authentication procedures can introduce delays in certain circumstances (e.g. when random numbers have to be generated), but authentication only needs to be performed once for each session.

5. Research and public health

There is an ever-growing list of parties that claim some right of access to the patient's record. Apart from the substantial fiscal interests involved, there are also significant concerns about being able to perform clinical research and public-health surveillance. For this reason, many health-care systems are developing data warehouses for analytic purposes, the contents of which are obtained under blanket consent from patients who received care within those systems. There is some controversy about the secondary uses of data obtained under these blanket consents, and furthermore, the data typically just capture the parts of the health history obtained during contact with the health-care system. And typically only of a single health-care system, even if the care over the patient's lifetime was distributed over several systems.

PING provides an alternative means for conducting research and public-health surveillance using both data obtained from institutional sources with patient annotations as well as the patient's own original entries (from a variety of patient-directed inputs). We have designed a program called a PING-poller, which broadcasts queries from a PING server to all of a specified population

of PING records. The PING poller then dynamically builds a relational database using the results returned by the PING server. Each user can specify those queries to which they will respond, and the subset of their data they want to make visible. For example, a patient may choose to allow public-health authorities to poll their immunization history (to allow the generation of region-wide immunization registries) and a pharmaceutical company only, all events that might be classified as adverse events during a constrained period of a drug trial to which the patient has consented. Several challenges are apparent and may worry both public health authorities and researchers: the patients may not choose to release all or any of their pertinent data. Furthermore, the data obtained will be biased by the nature of those patients who are agreeable to using web-based technology and furthermore agreeable to using PING as a reporting and record system. As PING complements and does not replace existing data sources, the first challenge is not as significant as it might seem. As for the second challenge, there are several statistical techniques, including independent randomized polling, that can be used to correct for the biased characteristics of this data collection and reporting system.

Perhaps the most promising feature of PING for large-scale population-based research is that it overcomes the often petty institutional obstacles that prevent data sharing. It is hardly a secret that several institutions have limited efforts for data sharing ostensibly in the cause of patient privacy when often it has been market competition, institutional rivalries and notions of intellectual property of the aggregated patient data that have been the underlying cause of resistance to integration. In contrast, patients may be strongly motivated to share their data with specific duly authenticated parties for various reasons (their own care, interest in furthering research in their own disease, general altruism). By placing a patient-augmented copy of these institutional records in the patient's control, it may be that the availability of this information for numerous researchers will increase.

6. Applications

If successful, PING will represent a major shift in the way medical information is collected, managed, and used. We believe that the key to the adoption of the PING technology will be represented by the ability to provide a set of applications that take advantage of it. The first implementation of the PING server, written in the Java™ programming language, is currently being used as the development platform for several such applications. The following are examples of applications currently under development or being planned for the near future.

- *Newborn medical record.* This application has been funded by the National Library of Medicine as part of Phase II of its Next Generation Internet project. The application domain of the postpartum period involves the collaboration of several institutions (birth hospital, primary care providers, tertiary care institution, state screening laboratory) and of the families, that often already maintain personal medical records (e.g. the baby booklets, immunization and growth histories). In addition, it does not require significant investments in entering antecedent data to bring the record up to date. In prior projects, we have already built some of the integration infrastructure required for this test-bed. Specifically, the birth hospitals (Brigham and Women's Hospital and Beth Israel), in which most of the infants seen

by Children's Hospital affiliated pediatricians are born, are linked securely over the Internet through the W3-EMRS data-sharing architecture [13]. This linkage was implemented as part of the BiliLIGHT project to implement automated guidelines for the management of infants with jaundice [14].

- *Secure and reliable lab-test result reporting to patients*. We developed a system based on PING to report automatically the results of throat culture tests to patients. The system is composed of two modules. The first module is a *puller* that extracts the test results from Children's Hospital database when they become available, and stores them in the patient's own PING record. The second module is a web-based interface through which the patients can retrieve the test results and acknowledge them. The whole system is designed to automate the process by which test results are communicated to patients, a process that currently relies on telephone calls. An additional benefit is that the test result, represented as an HL7 message, becomes a permanent part of the patient's PING record. The application is currently undergoing internal testing and will be tested with real users in the fall of 2000.
- *Immunization record*. Prior work in immunization systems contrasts with the PING approach. Linkins et al. [15] reported that 36 immunization registry projects were operational in the United States, with several hundred others in various stages of development or deployment. Unfortunately, each state has its own legislation governing the privacy and distribution of immunization records [16]. Even without considering the legal restrictions, there currently exist no standards for the interchange of immunization records between immunization registries. Because there is no standard way to query registries in an aggregate manner, agencies such as state departments of health and the Centers for Disease Control and Prevention thus have no consistent method to obtain aggregate regional measures of immunization rates. Twelve birthing hospitals in the state of Massachusetts now prompt parents to allow their newborn's immunization information to be entered into the state registry for sharing. This suggests that parents are now being exposed very early to the concept of the computerized immunization registry and that their child's vaccination history will be available there. However, with centralized state or regional registries, parents do not own their child's computerized immunization history, cannot easily view the history themselves, and cannot have the history follow them or their child. PING will allow cryptographically implemented, role-based access to authorized researchers and public-health authorities. Additionally, for particular queries, only anonymized data will be returned, whereas, for other applications, full identified data will be obtained (only upon the explicit, irrepudiable authorization by the patient). In this, too, the experience of the HIIDIT project will be particularly helpful. For example, a state or national organization with proper authorization could use such queries with PING to obtain aggregate measures of immunization rates without compromising a patient's privacy.
- *A personal genomic record*. The availability of high-throughput, massively parallel techniques to genotype and to assess RNA expression levels in various tissues is growing rapidly [17]. It is highly likely that with minimal tissue or blood, the entire genome of every individual will be quickly profiled within the coming decade. The prognostic and diagnostic quality of this information

is hard to overstate and so, in the same measure, will be the issue of confidentiality and patient autonomy. Consequently, one of the major outgrowths of both the PING and HIIDIT projects is the development of an extension of PING engineered specifically for the data modeling and confidentiality needs of the personal genomic record. Although there has been a lot of work on storing sequence data [18] and microarray data [19], there has been very little research on what it would take to store an individual's entire genomic information along with lifelong protectable access privileges.

7. Conclusions

We are now at a juncture in the development of health information systems in which the crucial role of the patient as primary informant and autonomous 'customer' is being increasingly recognized. There are several directions that the architectures of these systems can take as well as the technologies that are used to implement them. The decisions taken in this regard will have lasting properties in maintaining and changing the role of the patient vis-à-vis the health-care system and its providers. In particular, the confidentiality and control of the patient's record are widely seen as valuable commodities with multiple parties demanding access rights. In this paper, we have described a highly distributed architecture that depends on the least common denominator as its storage medium: generic web pages. Furthermore, the architecture allows for a high degree of patient control and autonomy in the management of their own annotated copy of the medical record while allowing for multiple public health and research applications. Preliminary prototypes have been implemented; actual deployment started in Spring 2000, and the first fully fledged testbed is expected to be in operation by the fall of 2000.

Although we believe that PING will prove to be a scaleable and useful personal healthcare infrastructure, we expect its principal contribution to be in setting a high standard for interoperability, patient control and autonomy in the use of their own clinical record. Consequently, we welcome collaborations in this enterprise, as well as comments and discussions regarding our particular design choices.

Acknowledgements

We wish to thank David Clark for advice and critique. This work was supported by grant N01LM-9-3536 from the National Library of Medicine.

References

[1] K.D. Mandl, I.S. Kohane, Healthconnect: Clinical Grade Patient–Physician Communication, in: Proceedings, AMIA Annual Symposium, 1999, pp. 849–853.
[2] Mandl KD, Szolovits P, and Kohane IS. Public standards and patient control: How to keep electronic medical records accessible but private. Br. Med. J. (in press).
[3] H.J. Tange, A. Hasman, P.F.dV. amd Robb, H.C. Schouten, Medical narratives in electronic medical records, Int. J. Med. Inform. 46 (1997) 7–29.
[4] I.S. Kohane, F.J.v. Wingerde, J. Fackler, et al., Sharing electronic medical records across multiple heterogeneous and competing institutions, in: Proceedings of the AMIA Annual Fall Symposium, Washington, DC, 1996, pp. 608–612.
[5] Health level seven standards. http://www.hl7.org/.
[6] G. Hripcsak, P.D. Clayton, T.A. Pryor, P. Haug, O.B. Wigerz, J.vd. Lei, The Arden Syntax for medical logic modules, in: R.A. Miller (Ed.), Proceedings, Symposium on Computer Applications in Medical Care, IEEE Computer Society Press, New Brunswick, New Jersey, 1990, pp. 200–204.

[7] The Guideline Interchange Format. http://www.glif.org/.
[8] E.R. Harold, XML Bible, IDG Books, Foster City, CA, 1999.
[9] HL7 SGML/XML Special Interest Group. http://www.mcis.duke.edu/standards/HL7/committees/sgml/index.html.
[10] New York Times Magazine. The time capsule. http://www.nytimes.com/library/magazine/millennium/m6/capsule-panel.html.
[11] I.S. Kohane, H. Dong, P. Szolovits, Health information identification and deidentification toolkit, in: C. Chute (Ed.), Proceedings, Annual Fall Symposium of the American Medical Informatics Association, Hanley & Belfus, Philadelphia, PA, 1998, pp. 356–360.
[12] B. Schneier, Applied Cryptography, Wiley, New York, 1995, p. 19.
[13] I.S. Kohane, P. Greenspun, J. Fackler, C. Cimino, P. Szolovits, W3-EMRS: Access to multi-institutional electronic medical records via the World Wide Web, in: Spring Congress of the American Medical Informatics Association. Boston, MA, 1995.
[14] Y. Sun, F.J.v Wingerde, I.S. Kohane, O. Harary, K.D. Mandl, S.R. SalemSchatz, et al., The challenges of automating a real-time clinical practice guideline, Clinical Performance and Quality Health Care 1 (7) (1999) 28–35.
[15] R.W. Linkins, S.M. Feikema, Immunization registries: the cornerstone of childhood immunization in the 21st century, Pediatr. Ann. 27 (6) (1998) 349–354.
[16] L. Gostin, Z. Lazzarini, Childhood immunization registries. A national review of public health information systems and the protection of privacy, J. Am. Med. Assoc. 274 (22) (1995) 1793–1799.
[17] D.D. Bowtell, Options available — from start to finish — for obtaining expression data by microarray, Nat. Genet. 21 (1999) 25–32.
[18] D.A. Benson, M.S. Boguski, D.J. Lipman, J. Ostell, B.F.F. Ouellette, B.A. Rapp, D.L. Wheeler, GenBank, Nucleic Acids Res. 27 (1) (1999) 12–17.
[19] A. Wang, A. Pierce, K. Judson-Kremer, S. Gaddis, C.M. Aldaz, D.G. Johnson, M.C. MacLeod, Rapid analysis of gene expression (RAGE) facilitates universal expression profiling, Nucleic Acids Res. 27 (23) (1999) 4609–4618.