Policy Forum ■

Against Simple Universal Health-care Identifiers

PETER SZOLOVITS, PhD, ISAAC KOHANE, MD, PhD

Medical records are becoming fully computerized. Technical, administrative, and economic forces are pushing toward standardization on a single identifier, such as the Social Security number (SSN), to index all records. Consequently, the privacy and security of our medical histories will be severely compromised. We argue that there are sensible and effective technologic means available to reduce the risks of such compromise, and that it is time to *design* the characteristics we want in our record-keeping systems.

Current Trends and Their Dangers

Over a year ago, a meeting of the American College of Medical Informatics took up the question of how to identify patient records and concluded that the simplest, most expedient solution was to adopt the SSN (extended by a check digit) as the universal health-care identifier.¹ The advantages of this proposal are that virtually everyone already has an SSN and that an organization exists that issues new ones as needed. Recognized disadvantages include relatively frequent cases where more than one person was issued the same SSN, people with multiple SSNs, and newborns and "marginal" people who have health care needs but no SSNs. Other potential problems include: eventual insufficiency of a nine-digit "addressing" scheme, lack of consistency checks leading to easy misidentification, and privacy considerations. Our major concerns grow mainly out of the considerations of privacy.

In any society that values privacy but also keeps records, there must be concern about possible invasions of privacy. It has always been possible to violate the privacy of selected individuals by devoting enough resources to looking into their affairs. Examination of public records, interviews with associates, the press, illicit access to private files, and physical surveillance all offer the dedicated snoop ways of invading the privacy of a particular person. Fortunately, this research is cumbersome and expensive, and can therefore be applied only in relatively rare situations.

If we organize our records in such a way that the indexing of information is routine, then we make the job of the snoop much simpler and less expensive. Credit bureaus, for example, with their voluminous data sets on our consumer behavior, make it easy to violate the privacy of any specific individual.² The reporter J. Rothfeder, for example, was able to use a simple ruse to gain access to such national information repositories and to track down specific details of the life of former Vice President Dan Quayle.³ The existence and penetrability of such databases currently provides a quantitative, not qualitative, deterioration of individual privacy. A public figure such as the vice president could always be selected as the target of investigation; the conveniently collected data, used as Rothfeder demonstrated, merely make this less costly.

A more threatening consequence of large, insecure databases is the ability to search them easily and cheaply for groups of previously anonymous people with certain characteristics.^{4,5} With this ability, the snoop can not only invade the privacy of someone already targeted, but can also develop new lists of interesting targets who meet specific criteria. The techniques are already widely (and legitimately) used by marketing organizations that select prospective catalog recipients by sorting through records of past purchases. Current use tends to be limited to searches over specific, isolated data sets, not over the lifetime accumulation of information about everyone. Yet, under current proposals, we are making it simpler

Affiliations of the authors: MIT Laboratory for Computer Science (PS), Cambridge, and Children's Hospital, Harvard Medical School (IK), Boston, MA.

Supported in part by a grant from IBM Corporation and by grant RO1 LM05296 from the National Library of Medicine.

Correspondence and reprints: Peter Szolovits, PhD, MIT Laboratory for Computer Science, 545 Technology Square, Room 416, Cambridge, MA 02139. e-mail: PSZ@lcs.mit.edu

Received for publication: 3/28/94; accepted for publication: 3/28/94.

to collate information from very different sources by indexing all transactions pertaining to an individual under his or her SSN. Future snoops may be able to develop lists of people with certain educational and job backgrounds who suffer from specific maladies and who like to spend money on certain kinds of entertainment. The opportunities for abuse are enormous. More widespread adoption of a single identifier facilitates and encourages just this situation.

The SSN was created by the government's Social Security program in the 1930s. Originally limited to recording individual contributions to the Social Security plan, its approved federal use has been broadened to identifying taxpayers and their tax transactions, civil service employees, Defense Department personnel, and recipients of some forms of public assistance. In addition, states use SSNs for their own tax-related records, and many also index drivers' licenses, motor vehicle registrations, and criminal histories to the same identifier. Nongovernment uses include records holding an individual's history of employment, insurance, credit, and education. If current trends continue, health records will join this list.

With growing interoperability of database systems, we are getting close to the time when a single query can, at low cost, find a selection of individuals based on any or all of the characteristics indexed by the SSN listed above. To anyone who values privacy even slightly, this is a frightening prospect.

The Falsely Perceived Technologic Imperative

When ACMI's recommendation first appeared, one of us (PS) began a discussion of these issues with ACMI members on a mailing list. We were surprised that many of our colleagues did not share our concerns. Our impression of the electronic discussion is that the responses split into three camps with the following (caricature) positions:

1. The universal use of the SSN is bad, because it makes the collation of large databases easy, and is likely to lead to intolerable abuses. (These are the people who agreed with us.)
2. It's too bad that using a universal identification number (ID) has these undesirable consequences, but the horse is out of the barn already, so we may as well just learn to live with it.

3. "No problem." And anyway, any reasonable alternatives would be inordinately costly.*

In addition, we were struck by the dearth of suggestions (from a community normally brimming with technical ideas) about how one might use technologic means to ameliorate these problems.

The ideas we present below focus on one specific, narrow issue: How to make it more difficult for unauthorized individuals to perform massive searches across large databases, collating information from multiple sources. We do not specifically address the "targeted" invasion of privacy where the individual to be investigated is already identified, but mainly address the search for interesting populations of people who share certain health (and other) characteristics but who are not initially known.

Technologic Alternatives

One of the major technical advances in computer science occurred in the late 1970s and is just now beginning to influence commercial practice: the development of *public-key cryptography*. In any cryptographic method, we try to ensure that the sender of a message can encode it in some way such that the recipient can decode it, but such that anyone else who intercepts the coded message cannot make sense of it. Simple codes have been in use since Roman times, and very effective coding schemes have been a routine part of military communications in our century. Such schemes require, however, that both the sender and the recipient of a message agree, ahead of time, on the coding and decoding method they will use. This makes spontaneous encrypted communication impossible, and leaves the difficult question of how to agree on a secret encoding/decoding method before there is an effective encrypted means of communicating. (In most cases, this has been done by physical exchange of codes, which requires significant advance planning.)

Public-key cryptography does away with the need for prearranged cryptographic codes in a clever way that relies on the complexity of certain mathematical computations. Anyone who wants to use this method needs to acquire two keys that allow arbitrary messages to be encoded and decoded. The keys define functions that are mathematical inverses of each other,

*Even so, verifying the accuracy of all SSNs and ensuring future accuracy would require a major organization and investment estimated at \$1.0 to \$2.5 billion.⁶

but with the characteristic that it is phenomenally costly to calculate one of them if given only the other.^{7†} For Alice, a typical user, we will call these A and A' . Alice keeps A a well-guarded secret, but publishes A' so that anyone interested in communicating a secret message to her may find it. If Bob, for example, wishes to send her a message m , he can compute $A'(m)$ and transmit it to Alice, who then applies her secret key A to it to recover the original message, because $A(A'(m)) = m$. A potential snoop, Charley, cannot reconstruct m even if he intercepts $A'(m)$ because he does not know A , and because it is very costly to compute A even knowing A' .

An interesting corollary of the fact that public and private keys are inverses is that the same keys form the basis for a reliable means of authenticating the sender of a message. Thus, if Bob's keys are B and B' , Bob can assure Alice that a message he sends her really came from him by encoding part of it with his private key and asking her to decode it with his public key, which she can easily find. Thus, Bob sends $A'(B(m))$, and Alice computes $B'(A(A'(B(m)))) = m$. If this computation yields a sensible message, Alice can conclude that its sender must have known B , and must therefore be Bob. The system is robust; individuals can change their keys simply by publishing a new public key, and elaborations are possible to create escrow agents, trusted intermediaries, messages that require cooperation among several parties to decode, and a broad range of other interesting secure communication mechanisms.

This approach is becoming commercially popular and is the basis of *digital signatures*, which now form an essential part of various computer-based authentication systems such as the Apple Macintosh PowerTalk Signers, PGP, a widespread public-domain digital signature scheme for Internet e-mail, General Magic's Telescript communication language, and MIT Project Athena's Kerberos user authentication system.

Cryptography-based Health-care Identifiers

We believe that some method based on the ideas of public-key cryptography can be fruitfully applied to

[†]Currently, the most popular method uses the computational difficulty of factoring very large (hundreds of digits) numbers to make it impossible to calculate one key from its inverse. The best currently known algorithms for factoring require a time that grows exponentially in the size of the number being factored. Thus, breaking one of these codes can be made practically impossible simply by choosing keys that are large enough. It has not been proven that this is the most efficient factoring algorithm, but since none that is more efficient has been found in many years of research, the technique is thought to be safe.

the problem of collecting and keeping comprehensive medical records that are easy to integrate when it is appropriate to do so, but that are difficult to collate in service of widespread searches that can invade an individual's privacy. We provide two example designs here to illustrate feasible alternative approaches.

The first is a completely decentralized scheme in which the individual patient has ultimate control over the degree to which the lifetime collection of medical information about him or her is made available to others. Instead of an SSN, every individual is issued, at birth, a unique private key, K . Every institution that maintains health records is issued a public key, H' . Thus, each hospital, clinic, doctor's office, HMO, insurance company, other third-party payer, government regulatory agency, epidemiologic data center, research study, etc., has its unique H' . When an individual first needs to deal with an institution, a universal cryptographic function f is applied to K and H' to compute the ID number under which that individual's records will be kept at that institution. The computation of ID does not, however, reveal the patient's private key, K , just as Alice's sending $A(m)$ did not reveal A . Therefore, any specific institution can have access only to those records it has itself collected about a patient, to those the patient has asked other institutions to forward to this one, and to those at other institutions for which the patient has provided this one their old ID .⁵

Because $ID = f(K, H')$, it depends on both the individual and the institution. Therefore, no two people will have the same ID at the same institution, and no two institutions will have the same ID for the same individual. Furthermore, a person can validate his or her identity by demonstrating the ability to compute his or her ID given the institution's H' . Under this scheme, if one institution needs to gain access to a patient's records at another, it must ask the patient to compute his or her ID at that other institution, by applying f to the patient's secret K and the public H' of the other institution. The same mechanism is used to tell a hospital, for example, the appropriate ID under which to report reimbursement claims to an insurance company or epidemiologic data to the Centers for Disease Control. Any institution holding an unauthorized patient record could be exposed by demonstrating that $f(K, H')$ does not correspond to the identifier used by that institution.

This scheme would require patients to carry something like a "smart card" instead of the current Social

Security card, because the computation of f is not practical without a small computer, and the keys involved are lengthy. Longitudinal records can be assembled by chaining records back from those of current health care providers to those of past providers, and a past provider could not find our newer information without the patient's cooperation. Mechanisms would also need to be established to deal with people losing their cards and keys, though the simplest expedient is just to issue someone a new K and then include in his or her new records references to the old records.

A second possible scheme centers on the institution rather than the individual, and assumes that all patients' private keys are held by a centrally organized *ID server*. For identification purposes, each patient has something that serves as his or her public key and can be as widely known as the SSN is today. The ID server must be a trusted institution, perhaps established and administered by the government or some semipublic consortium. Any institution may request (by authenticating its identity) an ID for a particular patient to be used at that institution, and the ID server returns it, protected by encrypted communications. In this case, neither the patient nor any health care provider knows the patient's private key or the patient's ID at any other institution. The ID server is responsible for arranging all transfers of data from one institution to another, according to an agreed-upon and authenticated protocol.

Neither of the two simple schemes mentioned above is likely to be satisfactory, and the design of a coherent scheme that meets the data communication and privacy needs of the health care system is a difficult task. For example, we have treated the patient's records at a single institution as indivisible, but it might be necessary to provide different protections for different parts of a medical record. Psychiatric records and HIV status may not need to be available to all accessors of the record.

We have focused on the use of technology to prevent unauthorized access. In some settings it may be more appropriate to give presumptive access to anyone who claims a need to know, but to use authentication methods to ensure an unforgeable "signed" log of everyone who has actually accessed the data. Improper access is then not prevented, but a record of

unjustifiable access implies an *a posteriori* liability for invasion of privacy.

The technology of public-key cryptography is a very powerful tool, and its creative application can lead to many interesting systems that provide different degrees of privacy, convenience, and flexibility. Some of these capabilities will certainly be developed piecemeal in any case, because American industry, including the health care industry, will adopt useful and easily available methods to enhance the security and reliability of its communications. Adopting technically naive solutions such as universal use of the SSN as a health-care identifier will simply lock us into a system that sacrifices personal privacy because of the mistaken impression that nothing better is feasible. Our peers on the "Information Superhighway" will laugh at our missed opportunities, at least until they realize that we have also compromised their privacy.

We call for a serious and informed discussion of the needs of health care record keeping. Let us decide what we want, define the trade-offs that will arise, and use the rich technologies available to design solutions that best achieve our desires.

References ■

1. Board of Directors of the American Medical Informatics Association. Standards for medical identifiers, codes, and messages needed to create an efficient computer-stored medical record. *Am Med Informatics Assoc.* 1993;1:1-7.
2. Kratka J. For their eyes only: The insurance industry and consumer privacy. Massachusetts Public Interest Research Group, 1990.
3. Rothfeder J. Privacy for sale: How computerization has made everyone's private life an open secret. New York: Simon and Schuster, 1992.
4. Curran WJ, Stearns B, Kaplan H. Privacy, confidentiality and other legal considerations in the establishment of a centralized health-data system. *N Engl J Med.* 1969;281:241-8.
5. Gostin LO, Turek-Brezina J, Powers M, Kozloff R, Faden R, Steinauer DD. Privacy and security of personal information in a new health care system. *JAMA.* 1993;270:2487-93.
6. Statement of Gwendolyn S. King, commissioner of Social Security. Hearings before the House Subcommittee on Social Security, Committee on Ways and Means. 102nd Congress, 1st session; 1991.
7. Rivest R, Shamir A, Adleman L. A method for obtaining digital signatures and public-key cryptosystems. *Comm ACM.* 1978;21:120-6.