

Achieving Electronic Privacy

A cryptographic invention known as a blind signature permits numbers to serve as electronic cash or to replace conventional identification. The author hopes it may return control of personal information to the individual

by David Chaum

Every time you make a telephone call, purchase goods using a credit card, subscribe to a magazine or pay your taxes, that information goes into a data base somewhere. Furthermore, all these records can be linked so that they constitute in effect a single dossier on your life—not only your medical and financial history but also what you buy, where you travel and whom you communicate with. It is almost impossible to learn the full extent of the files that various organizations keep on you, much less to assure their accuracy or to control who may gain access to them.

Organizations link records from different sources for their own protection. Certainly it is in the interest of a bank looking at a loan application to know that John Doe has defaulted on four similar loans in the past two years. The bank's possession of that information also helps its other customers, to whom the bank passes on the cost of bad loans. In addition, these records permit Jane Roe, whose payment history is impeccable, to establish a charge account at a shop that has never seen her before.

That same information in the wrong hands, however, provides neither protection for businesses nor better service for consumers. Thieves routinely use a stolen credit card number to trade on their victims' good payment records;

murderers have tracked down their targets by consulting government-maintained address records. On another level, the U.S. Internal Revenue Service has attempted to single out taxpayers for audits based on estimates of household income compiled by mailing-list companies.

The growing amounts of information that different organizations collect about a person can be linked because all of them use the same key—in the U.S. the social security number—to identify the individual in question. This identifier-based approach perforce trades off security against individual liberties. The more information that organizations have (whether the intent is to protect them from fraud or simply to target marketing efforts), the less privacy and control people retain.

Over the past eight years, my colleagues and I at CWI (the Dutch nationally funded Center for Mathematics and Computer Science in Amsterdam) have developed a new approach, based on fundamental theoretical and practical advances in cryptography, that makes this trade-off unnecessary. Transactions employing these techniques avoid the possibility of fraud while maintaining the privacy of those who use them.

In our system, people would in effect give a different (but definitely verifiable) pseudonym to every organization they do business with and so make dossiers impossible. They could pay for goods in untraceable electronic cash or present digital credentials that serve the function of a banking passbook, driver's license or voter registration card without revealing their identity. At the same time, organizations would benefit from increased security and lower record-keeping costs.

Recent innovations in microelectronics make this vision practical by providing personal "representatives" that store and manage their owners' pseudonyms, credentials and cash. Micropro-

cessors capable of carrying out the necessary algorithms have already been embedded in pocket computers the size and thickness of a credit card. Such systems have been tested on a small scale and could be in widespread use by the middle of this decade.

The starting point for this approach is the digital signature, first proposed in 1976 by Whitfield Diffie, then at Stanford University. A digital signature transforms the message that is signed so that anyone who reads it can be sure of who sent it [see "The Mathematics of Public-Key Cryptography," by Martin E. Hellman; SCIENTIFIC AMERICAN, August 1979]. These signatures employ a secret key used to sign messages and a public one used to verify them. Only a message signed with the private key can be verified by means of the public one. Thus, if Alice wants to send a signed message to Bob (these two are the cryptographic community's favorite hypothetical characters), she transforms it using her private key, and he applies her public key to make sure that it was she who sent it. The best methods known for producing forged signatures would require many years, even using computers billions of times faster than those now available.

To see how digital signatures can provide all manner of unforgeable credentials and other services, consider how they might be used to provide an electronic replacement for cash. The First Digital Bank would offer electronic bank notes: messages signed using a particular private key. All messages bearing one key might be worth a dollar, all those bearing a different key five dollars, and so on for whatever denominations were needed. These electronic bank notes could be authenticated using the corresponding public key, which the bank has made a matter of record. First Digital would also make public a key to authenticate electronic documents

DAVID CHAUM is head of the Cryptography Group at the Center for Mathematics and Computer Science (CWI) in Amsterdam. He is also a founder of Digi-Cash, which develops electronic payment systems. Chaum received his Ph.D. in computer science from the University of California, Berkeley, in 1982 and joined CWI in 1984. He helped to found the International Association for Cryptologic Research and remains active on its board; he also consults internationally on cryptology.

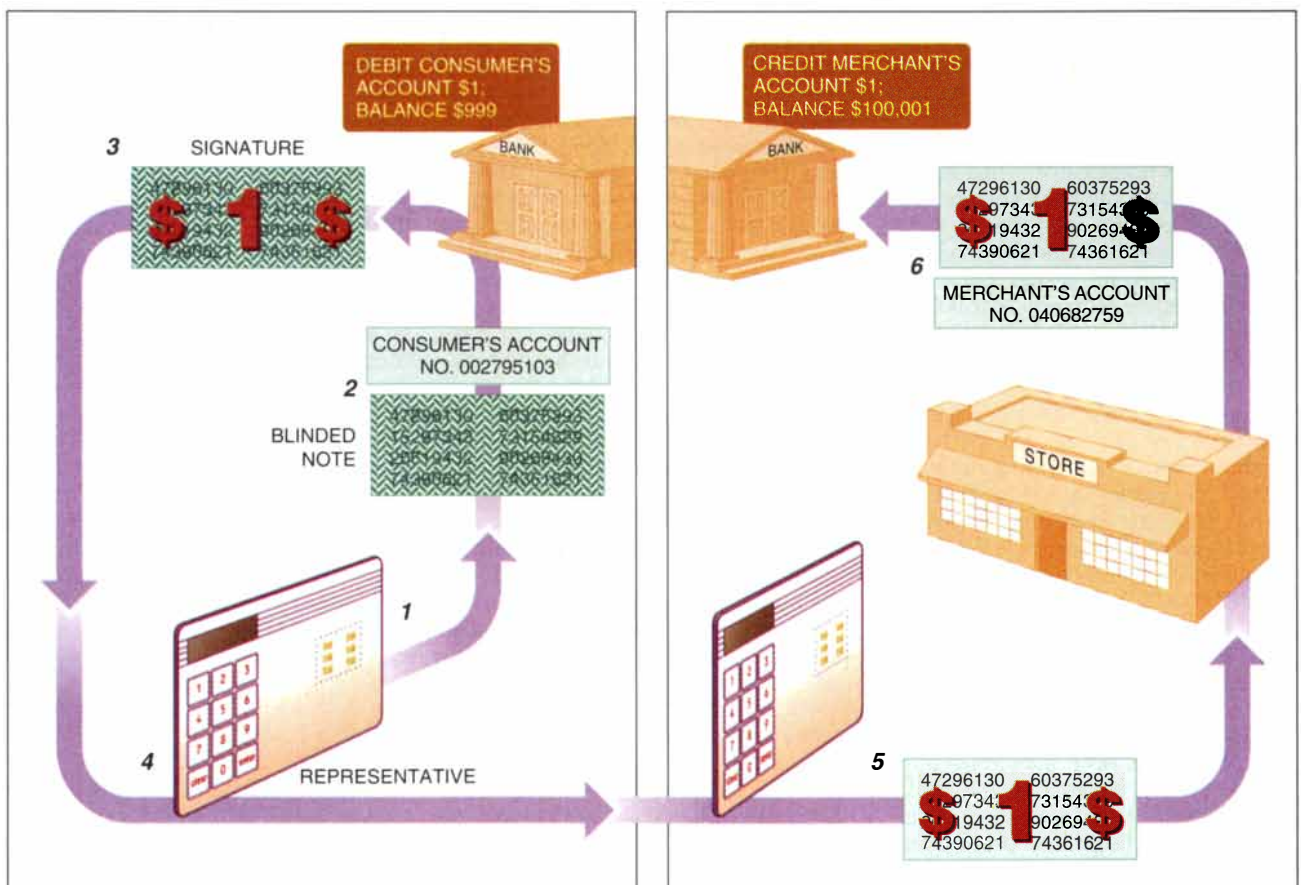
sent from the bank to its customers. To withdraw a dollar from the bank, Alice generates a note number (each note bears a different number, akin to the serial number on a bill); she chooses a 100-digit number at random so that the chance anyone else would generate the same one is negligible. She signs the number with the private key corresponding to her "digital pseudonym" (the public key that she has previously established for use with her account). The bank verifies Alice's signature and removes it from the note number, signs the note number with its worth-one-dollar signature and debits her account. It then returns the signed note along with a digitally signed withdrawal receipt for Alice's records. In practice, the creation, signing and transfer of note numbers would be carried out by Alice's card computer. The power of the cryptographic protocols, however, lies in the fact that they are secure regardless of physical medium: the same transactions could be carried out using only pencil and paper.

When Alice wants to pay for a purchase at Bob's shop, she connects her "smart" card with his card reader and transfers one of the signed note numbers the bank has given her. After verifying the bank's digital signature, Bob transmits the note to the bank, much as a merchant verifies a credit card transaction today. The bank re-verifies its signature, checks the note against a list of those already spent and credits Bob's account. It then transmits a "deposit slip," once again unforgeably signed with the appropriate key. Bob hands the merchandise to Alice along with his own digitally signed receipt, completing the transaction.

This system provides security for all three parties. The signatures at each stage prevent any one from cheating either of the others: the shop cannot deny that it received payment, the bank cannot deny that it issued the notes or that it accepted them from the shop for deposit, and the customer can neither deny withdrawing the notes from her account nor spend them twice.

This system is secure, but it has no privacy. If the bank keeps track of note numbers, it can link each shop's deposit to the corresponding withdrawal and so determine precisely where and when Alice (or any other account holder) spends her money. The resulting dossier is far more intrusive than those now being compiled. Furthermore, records based on digital signatures are more vulnerable to abuse than conventional files. Not only are they self-authenticating (even if they are copied, the information they contain can be verified by anyone), but they also permit a person who has a particular kind of information to prove its existence without either giving the information away or revealing its source. For example, someone might be able to prove incontrovertibly that Bob had telephoned Alice on 12 separate occasions without having to reveal the time and place of any of the calls.

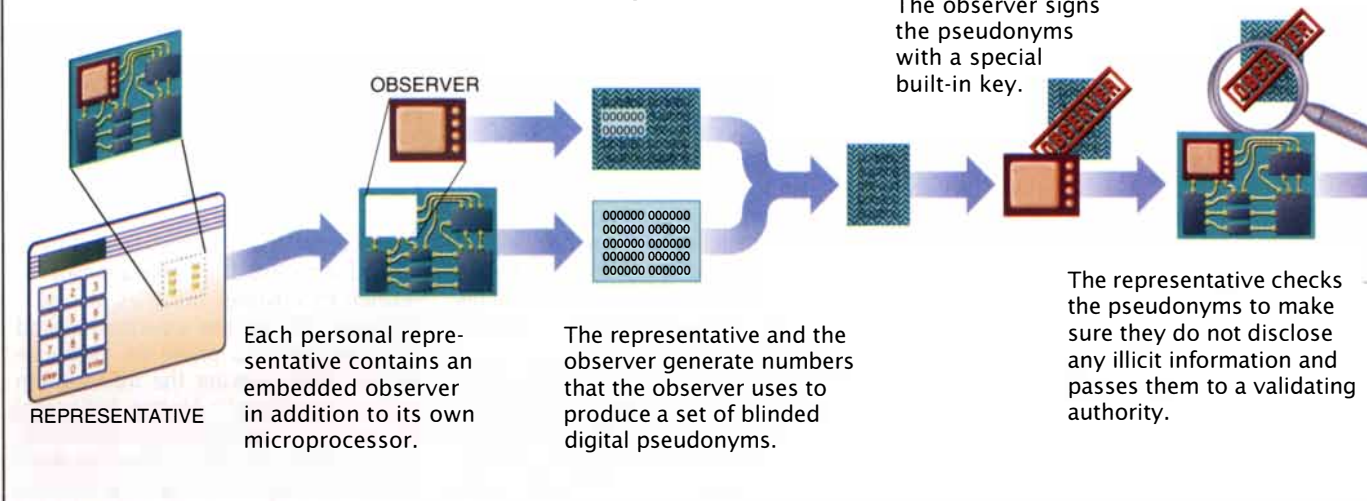
I have developed an extension of digital signatures, called blind signatures, that can restore privacy. Before send-



DIGITAL CASH flows tracelessly from bank through consumer and merchant before returning to the bank. Using a small computer "representative," a person creates a random number to serve as a bank note. The bank debits the appropriate account and signs the note with an unforgeable digital

signature indicating its value. The bank credits the merchant's account when the note is presented for payment. A technique known as a blind signature prevents the bank from seeing the note number so the bank will be unable to correlate withdrawals from one account with deposits to another.

How to Create Secure Digital Pseudonyms



ing a note number to the bank for signing, Alice in essence multiplies it by a random factor. Consequently, the bank knows nothing about what it is signing except that it carries Alice's digital signature. After receiving the blinded note signed by the bank, Alice divides out the blinding factor and uses the note as before.

The blinded note numbers are "unconditionally untraceable"—that is, even if the shop and the bank collude, they cannot determine who spent which notes. Because the bank has no idea of the blinding factor, it has no way of linking the note numbers that Bob deposits with Alice's withdrawals. Whereas the security of digital signatures is dependent on the difficulty of particular computations, the anonymity of blinded notes is limited only by the unpredictability of Alice's random numbers. If she wishes, however, Alice can reveal these numbers and permit the notes to be stopped or traced.

Blinded electronic bank notes protect an individual's privacy, but because each note is simply a number, it can be copied easily. To prevent double spending, each note must be checked on-line against a central list when it is spent. Such a verification procedure might be acceptable when large amounts of money are at stake, but it is far too expensive to use when someone is just buying a newspaper. To solve this problem, my colleagues Amos Fiat and Moni Naor and I have proposed a method for generating blinded notes that requires the payer to answer a random numeric query about each note when making a payment. Spending such a note once does not compromise unconditional untrace-

ability, but spending it twice reveals enough information to make the payer's account easily traceable. In fact, it can yield a digitally signed confession that cannot be forged even by the bank.

Cards capable of such anonymous payments already exist. Indeed, DigiCash, a company with which I am associated, has installed equipment in two office buildings in Amsterdam that permits copiers, fax machines, cafeteria cash registers and even coffee vending machines to accept digital "bank notes." We have also demonstrated a system for automatic toll collection in which automobiles carry a card that responds to radioed requests for payment even as they are traveling at highway speeds.

My colleagues and I call a computer that handles such cryptographic transactions a "representative." A person might use different computers as representatives depending on which was convenient: Bob might purchase software (transmitted to him over a network) by using his home computer to produce the requisite digital signatures, go shopping with a "palm-top" personal computer and carry a smart credit card to the beach to pay for a drink or crab cakes. Any of these machines could represent Bob in a transaction as long as the digital signatures each generates are under his control.

Indeed, such computers can act as representatives for their owners in virtually any kind of transaction. Bob can trust his representative and Alice hers because they have each chosen their own machine and can reprogram it

The observer signs the pseudonyms with a special built-in key.

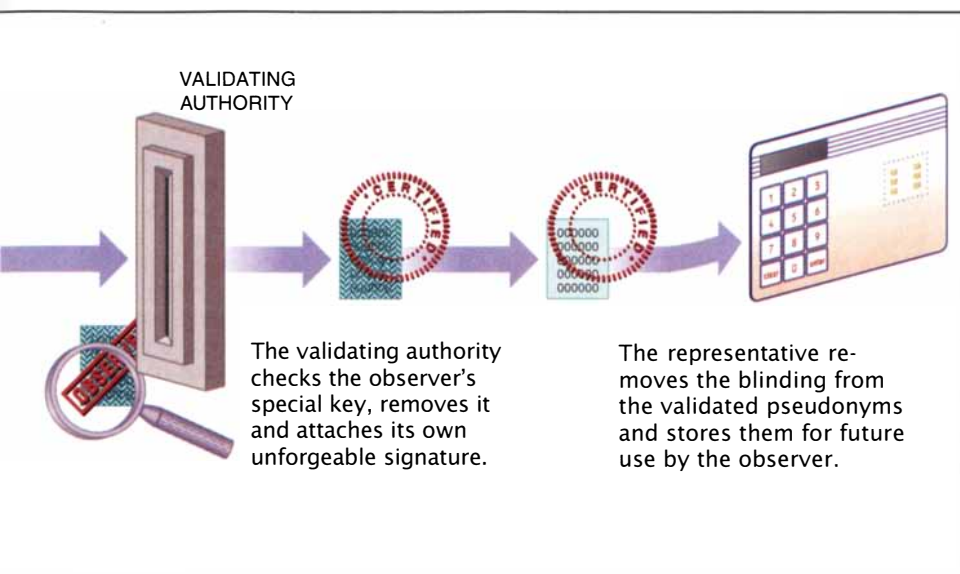
The representative checks the pseudonyms to make sure they do not disclose any illicit information and passes them to a validating authority.

at will (or, in principle, build it from scratch). Organizations are protected by the cryptographic protocol and so do not have to trust the representatives.

The prototypical representative is a smart credit-card-size computer containing memory and a microprocessor. It also incorporates its own keypad and display so that its owner can control the data that are stored and exchanged. If a shop provided the keypad and display, it could intercept passwords on their way to the card or show one price to the customer and another to the card. Ideally, the card would communicate with terminals in banks and shops by a short-range communications link such as an infrared transceiver and so need never leave its owner's hands.

When asked to make a payment, the representative would present a summary of the particulars and await approval before releasing funds. It would also insist on electronic receipts from organizations at each stage of all transactions to substantiate its owner's position in case of dispute. By requiring a password akin to the PIN (personal identifying number) now used for bank cards, the representative could safeguard itself from abuse by thieves. Indeed, most people would probably keep backup copies of their keys, electronic bank notes and other data; they could recover their funds if a representative were lost or stolen.

Personal representatives offer excellent protection for individual privacy, but organizations might prefer a mechanism to protect their interests as strongly as possible. For example, a bank might want to prevent double spending of bank notes altogether rather than



simply detecting it after the fact. Some organizations might also want to ensure that certain digital signatures are not copied and widely disseminated (even though the copying could be detected afterward).

Organizations have already begun issuing tamperproof cards (in effect, their own representatives) programmed to prevent undesirable behavior. But these cards can act as "Little Brothers" in everyone's pocket.

We have developed a system that satisfies both sides. An observer—a tamper-resistant computer chip, issued by some entity that organizations can trust—acts like a notary and certifies the behavior of a representative in which it is embedded. Philips Industries has recently introduced a tamper-resistant chip that has enough computing power to generate and verify digital signatures. Since then, Siemens, Thomson CSF and Motorola have announced plans for similar circuits, any of which could easily serve as an observer.

The central idea behind the protocol for observers is that the observer does not trust the representative in which it resides, nor does the representative trust the observer. Indeed, the representative must be able to control all data passing to or from the observer; otherwise the tamperproof chip might be able to leak information to the world at large.

When Alice first acquires an observer, she places it in her smart-card representative and takes it to a validating authority. The observer generates a batch of public and private key pairs from a combination of its own random numbers and numbers supplied by the

card. The observer does not reveal its numbers but reveals enough information about them so that the card can later check whether its numbers were in fact used to produce the resulting keys. The card also produces random data that the observer will use to blind each key.

Then the observer blinds the public keys, signs them with a special built-in key and gives them to the card. The card verifies the blinding and the signature and checks the keys to make sure they were correctly generated. It passes the blinded, signed keys to the validating authority, which recognizes the observer's built-in signature, removes it and signs the blinded keys with its own key. The authority passes the keys back to the card, which unblinds them. These keys, bearing the signature of the validating authority, serve as digital pseudonyms for future transactions; Alice can draw on them as needed.

An observer could easily prevent (rather than merely detect) double spending of electronic bank notes. When Alice withdraws money from her bank, the observer witnesses the process and so knows what notes she received. At Bob's shop, when Alice hands over a note from the bank, she also hands over a digital pseudonym (which she need use only once) signed by the validating authority. Then the observer, using the secret key corresponding to the validated pseudonym, signs a statement certifying that the note will be spent only once, at Bob's shop and at this particular time and date. Alice's card verifies the signed statement to make sure that the observer does not

leak any information and passes it to Bob. The observer is programmed to sign only one such statement for any given note.

Many transactions do not simply require a transfer of money. Instead they involve credentials—information about an individual's relationship to some organization. In today's identifier-based world, all of a person's credentials are easily linked. If Alice is deciding whether to sell Bob insurance, for example, she can use his name and date of birth to gain access to his credit status, medical records, motor vehicle file and criminal record, if any.

Using a representative, however, Bob would establish relationships with different organizations under different digital pseudonyms. Each of them can recognize him unambiguously, but none of their records can be linked.

In order to be of use, a digital credential must serve the same function as a paper-based credential such as a driver's license or a credit report. It must convince someone that the person attached to it stands in a particular relation to some issuing authority. The name, photograph, address, physical description and code number on a driver's license, for example, serve merely to link it to a particular person and to the corresponding record in a data base. Just as a bank can issue unforgeable, untraceable electronic cash, so too could a university issue signed digital diplomas or a credit-reporting bureau issue signatures indicating a person's ability to repay a loan.

When the young Bob graduates with honors in medieval literature, for example, the university registrar gives his representative a digitally signed message asserting his academic credentials. When Bob applies to graduate school, however, he does not show the admissions committee that message. Instead his representative asks its observer to sign a statement that he has a B.A. cum laude and that he qualifies for financial aid based on at least one of the university's criteria (but without revealing which ones). The observer, which has verified and stored each of Bob's credentials as they come in, simply checks its memory and signs the statement if it is true.

In addition to answering just the right question and being more reliable than paper ones, digital credentials would be both easier for individuals to obtain and to show and cheaper for organizations to issue and to authenticate. People would no longer need to fill out long and revealing forms. In-

stead their representatives would convince organizations that they meet particular requirements without disclosing any more than the simple fact of qualification. Because such credentials reveal no unnecessary information, people would be willing to use them even in contexts where they would not willingly show identification, thus enhancing security and giving the organization more useful data than it would otherwise acquire.

Positive credentials, however, are not the only kind that people acquire. They may also acquire negative credentials, which they would prefer to conceal: felony convictions, license suspensions or statements of pending bankruptcy. In many cases, individuals will give organizations the right to inflict negative credentials on them in return for some service. For instance, when Alice borrows books from a library, her observer would be instructed to register an overdue notice unless it had received a receipt for the books' return within some fixed time.

Once the observer has registered a negative credential, an organization can find out about it simply by asking the observer (through the representative) to sign a message attesting to its presence or absence. Although a representative could muzzle the observer, it could not forge an assertion about the state of its credentials. In other cases, organizations

might simply take the lack of a positive credential as a negative one. If Bob signs up for skydiving lessons, his instructors may assume that he is medically unfit unless they see a credential to the contrary.

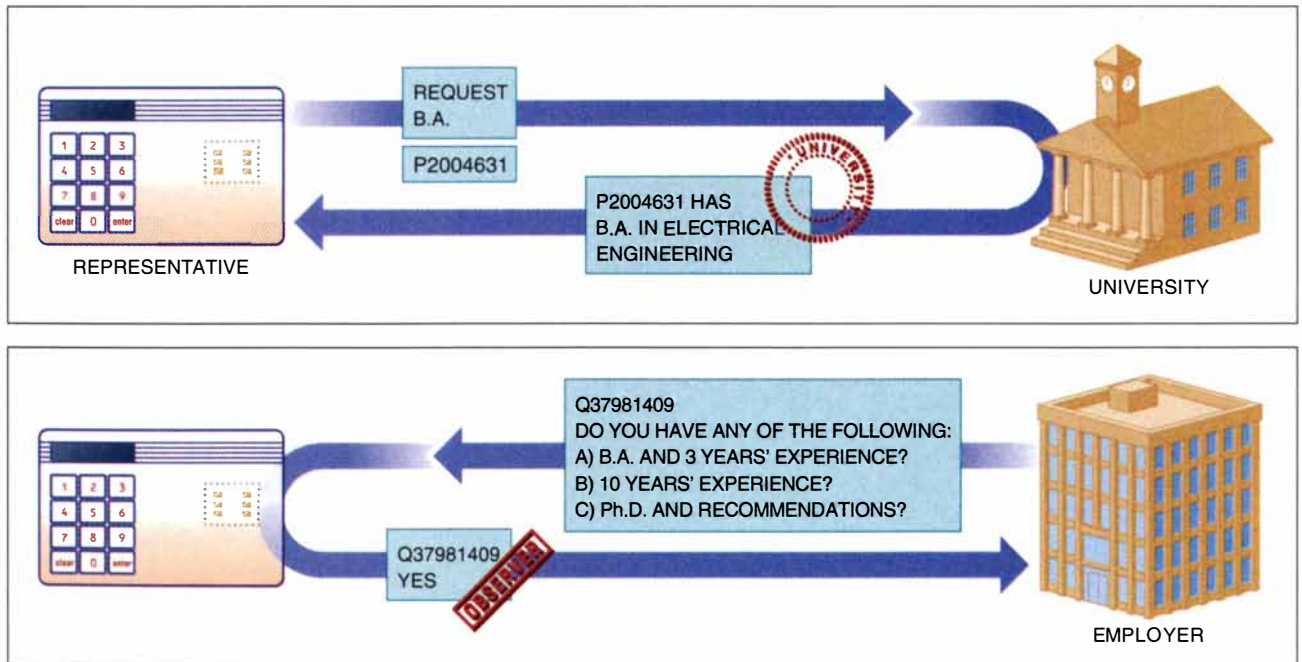
For most credentials, the digital signature of an observer is sufficient to convince anyone of its authenticity. Under some circumstances, however, an organization might insist that an observer demonstrate its physical presence. Otherwise, for example, any number of people might be able to gain access to nontransferable credentials (perhaps a health club membership) by using representatives connected by concealed communications links to another representative containing the desired credential.

Moreover, the observer must carry out this persuasion while its input and output are under the control of the representative that contains it. When Alice arrives at her gym, the card reader at the door sends her observer a series of single-bit challenges. The observer immediately responds to each challenge with a random bit that is encoded by the card on its way back to the organization. The speed of the observer's response establishes that it is inside the card (since processing a single bit introduces almost no delay compared with the time that signals take to traverse a wire). After a few dozen iter-

ations the card reveals to the observer how it encoded the responses; the observer signs a statement including the challenges and encoded responses only if it has been a party to that challenge-response sequence. This process convinces the organization of the observer's presence without allowing the observer to leak information.

Organizations can also issue credentials using methods that depend on cryptography alone rather than on observers. Although currently practical approaches can handle only relatively simple queries, Gilles Brassard of the University of Montreal, Claude Crépeau of the École Normale Supérieure and I have shown how to answer arbitrary combinations of questions about even the most complex credentials while maintaining unconditional unlinkability. The concealment of purely cryptographic negative credentials could be detected by the same kinds of techniques that detect double spending of electronic bank notes. And a combination of these cryptographic methods with observers would offer accountability after the fact even if the observer chip were somehow compromised.

The improved security and privacy of digital pseudonyms exact a price: responsibility. At present, for example, people can disavow credit card purchases made over the tele-



DIGITAL CREDENTIALS put personal information under the control of an individual's representative and its observer. When Alice (one of the author's two hypothetical characters) finishes her undergraduate work, the university gives

her a digitally signed degree. Later, her observer can use its knowledge of the degree to answer questions about her qualifications without revealing any more information about her than absolutely necessary.

phone or cash withdrawals from an automatic teller machine (ATM). The burden of proof is on the bank to show that no one else could have made the purchase or withdrawal. If computerized representatives become widespread, owners will establish all their own passwords and so control access to their representatives. They will be unable to disavow a representative's actions.

Current tamper-resistant systems such as ATMs and their associated cards typically rely on weak, inflexible security procedures because they must be used by people who are neither highly competent nor overly concerned about security. If people supply their own representatives, they can program them for varying levels of security as they see fit. (Those who wish to trust their assets to a single four-digit code are free to do so, of course.) Bob might use a short PIN (or none at all) to authorize minor transactions and a longer password for major ones. To protect himself from a robber who might force him to give up his passwords at gunpoint, he could use a "duress code" that would cause the card to appear to operate normally while hiding its more important assets or credentials or perhaps alerting the authorities that it had been stolen.

A personal representative could also recognize its owner by methods that most people would consider unreasonably intrusive in an identifier-based system; a notebook computer, for example, might verify its owner's voice or even fingerprints. A supermarket check-out scanner capable of recognizing a person's thumbprint and debiting the cost of groceries from their savings account is Orwellian at best. In contrast, a smart credit card that knows its owner's touch and doles out electronic bank notes is both anonymous and safer than cash. In addition, incorporating some essential part of such identification technology into the tamper-proof observer would make such a card suitable even for very high security applications.

Computerized transactions of all kinds are becoming ever more pervasive. More than half a dozen countries have developed or are testing chip cards that would replace cash. In Denmark, a consortium of banking, utility and transport companies has announced a card that would replace coins and small bills; in France, the telecommunications authorities have proposed general use of the smart cards now used at pay telephones. The government of Singapore has requested



COMPUTERIZED CREDIT CARD developed by Toshiba and Visa International contains a microprocessor, memory, keypad and display. Although this card identifies its user during transactions, the same hardware could be reprogrammed as a personal representative for spending digital cash.

bids for a system that would communicate with cars and charge their smart cards as they pass various points on a road (as opposed to the simple vehicle identification systems already in use in the U.S. and elsewhere). And cable and satellite broadcasters are experimenting with smart cards for delivering pay-per-view television. All these systems, however, are based on cards that identify themselves during every transaction.

If the trend toward identifier-based smart cards continues, personal privacy will be increasingly eroded. But in this conflict between organizational security and individual liberty, neither side emerges as a clear winner. Each round of improved identification techniques, sophisticated data analysis or extended linking can be frustrated by widespread noncompliance or even legislated limits, which in turn may engender attempts at further control.

Meanwhile, in a system based on representatives and observers, organizations stand to gain competitive and political advantages from increased public confidence (in addition to the lower costs of pseudonymous record-keeping). And individuals, by maintaining their own cryptographically guaranteed records and making only necessary disclosures, will be able to protect their privacy without infringing on the legiti-

mate needs of those with whom they do business.

The choice between keeping information in the hands of individuals or of organizations is being made each time any government or business decides to automate another set of transactions. In one direction lies unprecedented scrutiny and control of people's lives, in the other, secure parity between individuals and organizations. The shape of society in the next century may depend on which approach predominates.

FURTHER READING

- SECURITY WITHOUT IDENTIFICATION: TRANSACTION SYSTEMS TO MAKE BIG BROTHER OBSOLETE. David Chaum in *Communications of the ACM*, Vol. 28, No. 10, pages 1030-1044; October 1985.
- THE DINING CRYPTOGRAPHERS PROBLEM: UNCONDITIONAL SENDER AND RECIPIENT UNTRACEABILITY. David Chaum in *Journal of Cryptology*, Vol. 1, No. 1, pages 65-75; 1988.
- MODERN CRYPTOLOGY: A TUTORIAL. Gilles Brassard in *Lecture Notes in Computer Science*, Vol. 325. Springer-Verlag, 1988.
- PRIVACY PROTECTED PAYMENTS: UNCONDITIONAL PAYER AND/OR PAYEE UNTRACEABILITY. David Chaum in *Smart Card 2000: The Future of IC Cards*. Edited by David Chaum and Ingrid Schauer-müller-Bichl. North-Holland, 1989.