



SECURITY WITHOUT IDENTIFICATION: TRANSACTION SYSTEMS TO MAKE BIG BROTHER OBSOLETE

The large-scale automated transaction systems of the near future can be designed to protect the privacy and maintain the security of both individuals and organizations.

DAVID CHAUM

Computerization is robbing individuals of the ability to monitor and control the ways information about them is used. As organizations in both the private and the public sectors routinely exchange such information, individuals have no way of knowing if the information is inaccurate, obsolete, or otherwise inappropriate. The foundation is being laid for a dossier society, in which computers could be used to infer individuals' life-styles, habits, whereabouts, and associations from data collected in ordinary consumer transactions. Uncertainty about whether data will remain secure against abuse by those maintaining or tapping it can have a "chilling effect," causing people to alter their observable activities. As computerization becomes more pervasive, the potential for these problems will grow dramatically.

On the other hand, organizations are vulnerable to abuses by individuals. Everyone pays indirectly when cash, checks, consumer credit, insurance, and social services are misused. The obvious solution for organizations is to devise more pervasive, efficient, and interlinked computerized record-keeping systems, perhaps in combination with national identity cards or even fingerprints. However, this would exacerbate the problem of individuals' loss of monitorability and control, and would likely be unacceptable to many.

The new approach presented here offers an effective and practical solution to these problems.

The New Approach and How It Differs

Three major differences define the new approach. First is the way identifying information is used. Currently, many Western countries require citizens to carry documents bearing universal identification numbers. Driver's licenses are being upgraded to perform a similar function in the United States, and international efforts

for machine-readable national identity documents are gaining momentum. But organizations already use such essentially identifying data as name, date, and place of birth or name and address to match or link their records on individuals with those maintained by other organizations.

With the new approach, an individual uses a different account number or "digital pseudonym" with each organization. Individuals will create all such pseudonyms by a special random process. Information further identifying the individual is not used. A purchase at a shop, for example, might be made under a one-time-use pseudonym; for a series of transactions comprising an ongoing relationship, such as a bank account, a single pseudonym could be used repeatedly. Although the pseudonyms cannot be linked, organizations will be able to ensure that the pseudonyms are not used improperly by such measures as limiting individuals to one pseudonym per organization and ensuring that individuals are held accountable for abuses created under any of their pseudonyms. Individuals will be able to authenticate ownership of their pseudonyms and use them while ensuring that they are not improperly used by others.

A second difference is in who provides the mechanisms used to conduct transactions. Today, individuals hold a variety of "tokens" issued them by organizations, such as paper documents and plastic cards with magnetic or optical stripes or even embedded microcomputers. These tokens are usually owned by the issuing organization and contain information inscrutable to and unmodifiable by the individual holding them. Increasingly, individuals are being asked to perform transactions directly using computer-controlled equipment, such as automatic teller and point-of-sale terminals. Such equipment and chip cards are tamper resistant and contain secret numeric keys to allow secure communication with central computer facilities. Individuals

derive little direct benefit from these security provisions, however, since they must reveal their own secrets to the organization-provided mechanism and take the information provided to them by that mechanism on faith.

Individuals conduct transactions under the new approach using personal card computers that might take a form similar to a credit-card-sized calculator, and include a character display, keyboard, and a limited distance communication capability (like that of a television remote control). Such card computers could be purchased or constructed just like any other personal computer, and would have no secrets from or structures unmodifiable by their owners. They would be as simple to use as automatic teller machines. During a purchase at a shop, for example, a description of the goods and cost would be communicated to the card computer, which would display this information to the card owner, who would allow each transaction by entering a secret authorizing number on the card computer's keyboard. The same authorizing number originally programmed into the card computer by its owner is used to allow all transactions. Without this number, a lost or stolen card computer would be of very little use. However, the full capabilities of a lost card computer could be readily installed in a replacement card computer using backup data saved at home or elsewhere. The saved data would be in a safely encoded form that could only be decoded by a replacement card computer once the owner or some trustees supplied other sufficient secret numbers. These card computers are already technically feasible.

The nature of the security provided under the new approach also differs substantially: Current systems emphasize the one-sided security of organizations attempting to protect themselves from individuals; the new approach allows all parties to protect their own interests. The new approach relies on individuals keeping secret keys from organizations and organizations devising other secret keys that are kept from individuals. During transactions, parties use these keys to provide each other with specially coded confirmation of the transaction details, which can be used as evidence of improper actions sufficient to resolve disputes.

The systems presented in the new approach rely on currently used coding techniques to provide organizations with security against abuses by individuals. Consequently, if the underlying codes could be broken, individuals could breach the security of the systems. These codes are "cryptographic" and can be broken, in principle, by trying enough guessed keys, though such guessing is infeasible because of the enormous number of possible keys. No feasible attack or any proofs of security are known for these codes. In contrast, the security provided for individuals against organizations being able to link the pseudonyms in the systems presented here is "unconditional": Simple mathematical proofs show that, with appropriate use of the systems, even conspiracy of all organizations and tapping of all communication lines cannot yield enough information to link the pseudonyms—regardless of how clever the

approach is or how much computation is expended.

The feasibility of the new approach can be demonstrated for a comprehensive set of three kinds of consumer transactions: communication, payment, and credential. Each of these kinds of transactions raises its own special problems.

COMMUNICATION TRANSACTIONS

As more communication travels in electromagnetic and digital form, it becomes easier to learn more about individuals from their communication. Exposure of message content is one obvious danger that is already addressed by well-known cryptographic coding techniques. A more subtle and difficult problem with current communication systems, however, is the exposure of "tracing information." Individuals' addresses, which are often required by organizations and are commonly sold freely by them as mailing lists, are one kind of tracing information. The trend is toward greater use of such information. Comprehensive and computerized information on who telephones whom and when, for instance, is increasingly being collected and maintained by phone companies. Emerging electronic mail systems, other computer networks, and even some new phone systems automatically deliver tracing information with each message. When this information is available on a mass basis, associations, their structure, and even their relation to events are laid bare. Furthermore, tracing information can be used to link together all the records related to an individual that are held by organizations with whom the individual communicates. So long as communication systems allow system providers, organizations, or eavesdroppers to obtain tracing information, they are a growing threat to individuals' ability to determine how information about themselves is used. They are also unsuitable for the new approach.

The other side of the issue is that current systems provide inadequate protection against individuals who forge messages, or falsely disavow having sent or received messages. With paper communication, handwritten signatures are easily forged well enough to pass routine checking against signature samples and cannot be verified with certainty, even by expert witnesses. Also, paper receipts for delivery are too costly for most transactions, are often based only on handwritten signatures, and usually do not indicate message content. Emerging electronic mail and similar systems address these problems under the current approach in several obvious ways: by attempting to guarantee recipients the correct address from which each message is sent; by installing tamper-resistant identity card readers or the like at public points of entry to the communication system; and by keeping records of messages delivered, to provide certification of delivery. As computerized systems come into wider use, potential for such abuse by individuals will increase, but such solutions under the current approach rely on tracing information and thus are in fundamental conflict with individuals' ability to control access to information about themselves.

The nature of the solution is such that messages are untraceable, except for the recipient's ability to authen-

ticate them as having been sent by the owner of a particular pseudonym. The concepts of untraceability and pseudonymous authentication, presented separately in the following, are intertwined in the payment and credential transaction systems to be presented.

Unconditional Untraceability


The problem of preventing messages from being traced to the sender is now considered. The essential concept of the solution can be illustrated by a hypothetical situation. Suppose you were invited to dine at a restaurant by two of your friends. After dinner, the waiter comes to your table and mentions that one of the three of you has already paid for the dinner—but he does not say which one. If you paid, your friends want to know since they invited you, but if one of them paid, they do not want you to be able to learn which of the two of them *has* paid.

The problem is solved at the table in the following simple way: Your friends flip a coin behind a menu so that they can see the outcome, but you cannot. It is

agreed that each of them will say aloud which side the coin falls on, but that if one of them paid that one should say the opposite side. The uninteresting case is when they both say heads or both say tails: Then everyone knows you paid. If one of them says heads and the other says tails, however, then you know that one of the two of them paid—but you have absolutely no information as to which one. You do know that the one you heard say tails paid if the coin was heads, and that the other one paid if the coin was tails. But since each outcome of the coin toss is equally likely, you learn nothing from their utterances about which of the two of them has paid.


The system described allows the friend who paid to send you an unconditionally untraceable message; even though you know who says what, you cannot trace the “I paid” message, no matter how clever or time consuming your analysis.

Converting this two-sender single-recipient system to a more general system requires several extensions (presented and fully detailed in [2]). Increasing the number



Universally identifying numbers or other equivalent identifying information is presented by the individual cardholder to each organization—in the current approach. Unrelated generic examples are shown of three kinds of transactions: *communication*, in which the individual sends an authorizing message and receives a notifying message; *payment*, in which the

individual pays an organization or receives a payment; and *credential*, in which a certification that an individual has some credential is transferred from an organization B to an organization C. The identifying information—845—allows all transaction records to be linked and collected together into a dossier on the individual.



Different numbers or digital pseudonyms are used with each organization by a personal card computer held and trusted only by the individual—under the new approach. The credential transfer is no longer just between organizations: It must now go through the card where the pseudonym—451—used with the issuing organization B is transformed to the

pseudonym—314—used with the receiving organization C. Systems using this approach can provide organizations with improved protection against abuses by individuals, and also allow individuals to ensure that pseudonyms cannot be traced across the dashed boundary lines, thereby preventing dossier compilation.

of potential senders beyond two can prevent even cooperating subsets of potential senders from tracing transmissions to particular senders. Just as many other people may overhear the statements made at the table; actual systems would, in effect, broadcast each transmission to all participants, preventing anyone from knowing who receives which message. Because real messages are digitally coded, further coding (detailed later) can prevent all but the intended recipient from decoding confidential messages.

Digital Signatures

Now consider the problem of preventing senders of messages from later disavowing their messages. The solution is based on the concept of *digital signatures*, first proposed by Diffie and Hellman [5]. To see how this concept works, consider an old-fashioned codebook divided into two halves, like an English–French and French–English dictionary, except that only English words are used. Thus, if you look up an English word in the front half of the codebook, you find the corre-

sponding (but usually semantically unrelated) English code word; if you then look this code word up in the back half, you find your original English word. Codebooks are constructed by pairing off words at random: In the front half of the book, the pairs are ordered by their first words, and in the back half by their second words.

If you construct such a codebook, you can use it in your communication with an organization. You keep the front half as your *private key*, and you give the back half to the organization as your *digital pseudonym* with that organization. Before sending a message to the organization, you encode the message by translating each word into code using your private key; this encoded message is called a digital signature. When the organization receives the digital signature from you, it translates it back to the original English message using your digital pseudonym.

The immensely useful property of digital signatures is their resistance to “forgery.” No one—not even the organization that has your digital pseudonym—can eas-



Unconditionally untraceable messages are illustrated by a hypothetical situation (see text). The "I paid" message is unconditionally untraceable, since the guest (right) cannot trace it to a particular host—no matter how much computation or which approach is used.

ily forge a digital signature of yours. Such forgery would entail creating something that decodes to a sensible English message using your digital pseudonym. In the codebook analogy, forgery, of course, merely requires searching through or completely inverting the half of the book that is your digital pseudonym, but with actual digital-signature cryptographic techniques currently in use, forgery is thought to require so much computation as to be infeasible even for the fastest computers working for millions of years. If an organization cannot forge a digital signature of yours, then it cannot successfully claim that you sent it a message that you in fact did not send. A third-party arbiter would decide in favor of the organization only if that organization could show a digital signature that yielded the disputed message when translated with your digital pseudonym. But, because forgery is infeasible, the organization can only show such a message if you created it. Naturally, organizations would save copies of all digital signatures in anticipation of such disputes.

An organization could create its own private key/digital-pseudonym pair, and widely disseminate the digital pseudonym while keeping the corresponding private key to itself. It would use this private key to form digital signatures on all messages before sending them to individuals. The organization, however, would create only a single pair, which it would use for all digital signatures it issues. Anyone getting a message from the organization would first decode it using the organization's disseminated digital pseudonym. This would allow individuals to convince the organization, or anyone else if necessary, that the message had in fact been sent by the organization. In the payment and credential systems introduced in the following two sections, such digital signatures, as issued by organizations, play an important role.



Digitally signed messages are also illustrated by a hypothetical situation (see text). Actual computerized digital signature systems now in use are not unconditionally secure, although the amount of computation required for forgery is thought to be unobtainable in practice.

Digital Signatures in Practice

Actual digital signatures are realized using numbers, and can be extended to ensure confidentiality of message content and provide certification of delivery.

Practical computerized digital-signature techniques work like the codebook analogy above, except that everything is done with numbers. Private keys and digital pseudonyms are represented as two-hundred-digit numbers, instead of as halves of codebooks; messages and signatures are also represented as two-hundred-digit numbers, instead of as strings of English words. A standard public mathematical procedure allows anyone with a private key to form a corresponding digital signature from a message, and a similar procedure allows recovery of the original message using the corresponding digital pseudonym (just as the simple procedure for looking words up in either half of the codebook can be public, so long as the private key is not). Another public mathematical procedure allows anyone to create a private-key/digital-pseudonym pair from a random starting point (just as a simple procedure allowed the two halves of a codebook to be generated from a random pairing of words). Rivest, Shamir, and Adleman [6] proposed such a numeric digital-signature technique, which seems to be highly secure against forgery.

Message confidentiality during transmission is obtained by using digital pseudonyms and private keys in a different way: After signing a message, but before transmitting it, the sender encodes it using the digital pseudonym of the intended recipient. Thus, the signed message can be recovered only by decoding the transmission using the intended recipient's private key.


Currently, there are two strategies for preventing false disavowal of message receipt. Both of these strategies can be adapted for digital signatures. One imitates the approach currently used to certify paper mail: Mes-

sages are only given to the recipient if the recipient provides a digitally signed receipt of delivery. The other holds all potential recipients responsible for messages made available as a matter of public record. This allows either party to present the signed message and point to the corresponding doubly encoded transmission in the public record as evidence that the message was available for receipt, since decoding the signed message with the digital pseudonym of the sender yields the message content, and encoding it with the pseudonym of the recipient yields the transmission in the public record.


PAYMENT TRANSACTIONS

Automation of payment systems is giving the providers of these systems and others easy access to revealing and extensive information about individuals through payments for things like travel, entertainment, purchases


from shops, subscriptions, donations, etc. Today, many paper transaction records of when, how much, and to whom payment was made are translated into electronic form. The trend is toward initial capture of payment data in electronic form, such as at the point of sale, facilitating the electronic capture of the potentially more revealing details of what was purchased. Computerization is extending the data capture potential of payment systems in other ways, such as by the variety of emerging informational services proper, like pay television and videotex, and also by new systems that directly connect central billing computers to things like electric-utility meters and automobile-identification sensors buried in toll roads. Just as tracing data in communication systems allows all of an individual's records with organizations to be linked because they all use the same address, payment data allow linking of records that involve payments with the same account.



Unconditionally untraceable messages with numbers are sent essentially as with words, except that everything is represented as zeros and ones. Only the exclusive-or operation \oplus is used (defined as $1 \oplus 0 = 0 \oplus 1 = 1$ and $0 \oplus 0 = 1 \oplus 1 = 0$). The 0 or 1 outcome of the coin toss is shown as k . A host wishing to send the "I paid" message, which is represented as 1, transmits $k \oplus 1$; a host not wishing to send the message transmits only k . When the guest forms the exclusive-or of the two transmissions, $[1]$ and $[2]$, the result is 1 if one host sent the message and 0 if no host sent it—because k appears twice and cancels (since $k \oplus k = 0$ and $0 \oplus m = m$). If there are more hosts at the table, each flips a coin and shares the outcome with the host to the left, skipping the guest. Each host then forms a transmission as the exclusive-or of the two outcomes he or she shares, exclusive-ored with an additional 1 if the "I paid" message is being sent. Every coin toss appears twice and is canceled in the exclusive-or that the guest forms from all the transmissions, and the result is again 1 if a host paid and 0 if no host paid. In actual computerized systems, real messages are encoded as sequences of zeros and ones, and the whole protocol is repeated with new k s for each digit to be sent. Senders noticing that their messages are being garbled by collision with other messages wait a randomly chosen interval before attempting to resend.



Digital signatures with numbers use special arithmetic systems, in which raising a number to a power scrambles it, and raising it to a corresponding power unscrambles it. (One power acts as the private half of the codebook, and the other power as the corresponding half, called a digital pseudonym.) First the message is encoded as a one-hundred-digit number, and then the digits are repeated to form a two-hundred-digit number with this special repeated halves property. Next the signer raises the special number to a private power \bar{x} and makes the result known to others in transmission $[1]$. Someone obtaining this digitally signed message merely raises it to the corresponding digital pseudonym power x and checks that the result has the special repeated halves property. If it does, then the recipient knows that the message was signed by the holder of the private power.



Untraceable payments are illustrated by an analogy to envelopes and carbon paper. The individual (actually the card in the computerized analog) seals a blank slip of paper and a facing piece of carbon paper in an envelope, and supplies it to the bank. The bank deducts one dollar from the individual's account, applies a "worth one dollar" signature (stamp) on the outside of the envelope, and returns the unopened envelope to the individual. Upon receiving this, the individual verifies the bank's validating signature. Before making payment some time later, the individual removes the envelope and carbon, leaving only the signed slip of paper. When the shop receives the slip, it verifies the carbon image of the

validating signature on it and supplies it to the bank for deposit. After also verifying the slip's validating signature, the bank honors the deposit since it knows the slip must have been in an envelope that it signed. The bank does not, however, know which of the many envelopes that it signed contained the note, and thus cannot trace it to the individual's account. In actual computerized systems, unless the individual allows tracing, withdrawals on one side of the dashed boundary and payments on the other side are unconditionally untraceable to each other—even if the bank and all other organizations cooperate.

Abuses of payment systems by individuals, as well as abuses facilitated by payment systems, are also substantial and growing problems. Uncollectible payments made by consumers, such as checks drawn against insufficient funds and credit-card misuse, cost society billions of dollars each year. Paper-currency-based systems are vulnerable to such things as counterfeiting and theft. Lack of auditability also allows paper currency to be conveniently used for illicit payments such as bribes, extortion, and black-market purchases. Protecting against these various kinds of abuse while computerizing under the current approach seems to call for highly pervasive and interlinked systems capturing and retaining account identifiers as well as other payment data, which is naturally in conflict with the interests of individuals.

These problems are solved with the new systems since no organization, not even the payment system provider who maintains the accounts, is able to trace the flow of money between accounts. The system provider naturally knows the balance of each account, and

if funds were to transfer between accounts instantaneously, the simultaneous but opposite changes in balance would make tracing easy. The new system prevents such tracing in practice by allowing funds to be withdrawn and held as multidenomination notes, in some ways like "unmarked bills," before they are deposited to other accounts. The systems differ from paper currency, however, in part because individuals, but not organizations, can allow transfers to be traced and audited whenever needed, making stolen funds unusable and these systems unattractive for many kinds of illicit payments. The fully computerized systems introduced here offer practical yet highly secure replacements for most current and proposed consumer payment systems (as detailed in [3]).

Blind Signatures for Untraceable Payments

The payment system introduced is based on an extension of digital signatures known as *blind signatures*. This concept is easily understood by an analogy to carbon-paper-lined envelopes. If you put a piece of paper inside


such an envelope and a signature mark is later made on the outside of the envelope, the carbon paper in the envelope transfers the signature onto the slip.

Consider how you might use such envelopes to make payments. Suppose a bank had a special signature mark that it guaranteed to be worth one dollar, in the sense that the bank would pay one dollar for any piece of paper with that mark on it. You take a carbon-lined envelope containing a plain slip of paper to the bank and ask to withdraw one dollar from your account. The bank then deducts one dollar from your account, makes the signature mark on the outside of your envelope, and returns it to you. The signature is "blind" since the bank cannot see the slip through the envelope. Upon getting the unopened envelope back, you verify that the proper signature mark has been made on it. When you remove the slip from the envelope, it bears the carbon image of the bank's signature mark. You can then go out and buy something for one dollar from a shop, using the signed slip to make payment. The shop verifies the carbon image of the bank's signature on the

slip before accepting it as payment.

Now consider the position of the bank when a slip is received for deposit from a shop. The bank verifies the signature on the slip submitted for deposit, just as the shop did, and puts a dollar on the shop's account. Because the signature checked out, the bank knows that the slip must have been in an envelope that it signed. But of course the bank uses exactly the same signature mark to sign many such envelopes each day for all its account holders, and since all slips were hidden in envelopes during signing, the bank cannot know which envelope the slip was in. Therefore it cannot learn which account the funds were withdrawn from. More generally, the bank cannot determine which withdrawal corresponds with which deposit—the payments are untraceable.

In actual computerized systems, the envelopes and slips of paper are replaced by numbers, the bank's signature mark by a digital blind signature, and payments are unconditionally untraceable (as detailed in *Leaving the Analogy*, below). The protocol for transacting a



Untraceable payments with numbers are made much as in the paper analogy. First the individual's card computer chooses half the digits of n by a physical random process, and repeats these digits to form the note number n with this special repeated halves property (which is equivalent to choosing a suitable slip of paper at random in the analogy). The card also forms a totally random number r (which is equivalent to choosing an envelope and carbon). The card then raises the random number r to the bank's "worth one dollar" public power b , multiplies this by the note number n (which is equivalent to sealing the slip in the envelope), and supplies the result to the bank in transmission [1]. The bank deducts from the account, uses the corresponding private

power \bar{b} to sign the transmission, and returns the result to the card in [2]. The card verifies that the bank returned exactly the right thing, and obtains the signed note by dividing out the random r (which is equivalent to removing the envelope and carbon). When a payment is made, the shop checks that transmission [3] is a signed special number, and then forwards a copy [4] to the bank for deposit. The bank checks the signature just as the shop did, and accepts the deposit if the particular note has not already been deposited. If individuals do not divulge the random r s that their cards create, the [1]s can be unconditionally untraceable to the [4]s because there is exactly one r that would make any [1] correspond with any [4].

withdrawal from a bank or making a payment would of course be carried out automatically by the card computer; the card computer's owner would only have to allow transactions by entering the secret authorizing number.

Extending the Envelope Analogy

Note numbers can provide much the same kind of protection as check numbers do today. Since the bank is unable to look into the envelopes, nothing is revealed to the bank by a random number written on the slip before it is signed. (In fact, each slip has a unique random paper fiber pattern that might serve as just such a note number.) Stolen notes should not be accepted by the bank if the individual who withdrew the funds reports the note numbers. Also, the bank can attest to the account to which funds have been deposited if the individual payer provides the note number. Such traceability by the payer would discourage use of these systems for payment of bribes, extortion, and other illicit payments: Receivers of such payments risk having their accounts traced if they deposit the notes, and being apprehended or just finding that what they have is worthless if they try to spend them.

A variation on this system prevents organizations (even in cooperation with banks) from tracing the accounts of individuals to whom they pay such things as wages, refunds, settlements, and rebates. The individual places the slip in the envelope as before. This blinded slip is then provided to the payer organization (instead of the bank), which then supplies the blinded slip to the bank for signing and withdrawal from its own account. The signed but still blinded slip is then returned by the organization to the individual, who verifies the signature, removes the envelope, and later provides the contained slip to the bank for deposit.

Other extensions to the basic concept, not considered here, can offer replacements for today's payment systems attractive to both financial institutions and consumers. Different signatures would be used for different denominations. Clearing centers could handle most of the work and responsibility, while allowing banks to offer their own customized services with reduced investment and risk. Further variations allow the payment system to be used just as credit and debit cards are used today, with interest charges for use of credit and interest earnings on unspent funds. Generic receipts indicating only the denominations and type of expense could be used for tax reporting and the like. (When using credential systems presented in the next section, such receipts obtained under pseudonyms of one individual cannot be shown on pseudonyms of other individuals.)

Leaving the Analogy

Actual payment systems would work very much along the lines of the paper analogy, except that they would use numbers (as detailed more fully in the figure on page 1037). A note number is first created by a physical random process within the individual's card computer (like the note number chosen at random and written on

the slip of paper by the payer). Next, the card computer transforms this note number into the numeric equivalent of the message "this is note number: 416 . . ." The card computer then "blinds" this numeric note by combining it with a second random number (corresponding to the payer choosing an envelope at random and placing the slip in it). During withdrawal, the bank uses the private key of the desired denomination to form a digital signature on the numeric note (like the signature mark formed on envelopes by the bank). When the signed blinded note is ultimately returned, the card computer is able to unblind the note by a process that removes the random blinding number from the digital signature while leaving the signature on the note (like the payer removing the envelope). The organization receiving payment uses the digital pseudonym of the bank to decode the signature and verify that the numeric note contains an appropriate message and is thus a valid digital signature.

There might seem to be danger in that numbers, unlike paper, can be copied easily and exactly. Banks must be sure that the same numeric notes cannot be deposited more than once. A solution is for the bank to maintain a list of note numbers accepted, and to consult the list before accepting a note for deposit. The cost of maintaining such a list can be far less per transaction than the actual transaction cost of current payment systems, especially since expiration dates in note numbers can allow old numbers to be discarded.

Another conceivable danger is that the bank's digital signature could be forged, which would allow counterfeiting. The security against this kind of threat derives from the underlying digital-signature cryptographic technique, which is currently being proposed as an international standard and being used by banks and even to protect nuclear materials. The odds of someone guessing a valid signed numeric note or of any two independently chosen two-hundred-digit note numbers being the same are on the order of 1 in 10 to the 75th power.

The numeric notes are unconditionally untraceable: The correspondence between withdrawals and deposits cannot be learned by the bank from the numbers. In the untraceable communication system described in the last section, the possible outcomes of the coin tosses were both equally likely, which meant that every correspondence between senders and messages was equally likely. Similarly, because all suitable numbers are equally likely to be used for the independent blinding of each note, all correspondences between withdrawals and deposits are equally likely. More specifically, a unique random blinding number is implied by the correspondence between any particular blinded note and any particular signed note.

CREDENTIAL TRANSACTIONS

There are legitimate needs for individuals to show credentials in relationships with many organizations. Problems arise when unnecessary data are revealed in the process. As used here, "credentials" are statements based on an individual's relationship with organiza-

tions that are, in general, provided to other organizations. Some credentials, such as passports, drivers' licenses, and membership cards, are commonly shown by individuals in the form of certificates. Individuals control access to these certificates, but not to the irrelevant or unnecessary information—such as address, date of birth, and universally identifying numbers—that they usually also contain. Individuals are also often asked to provide credential information without substantiating certificates, as when they fill out applications or tax forms. Even when the credential needed is simple, such forms often request much unnecessary or unnecessarily detailed data, presumably to allow confirmation. But confirmation can link irrelevant information and ultimately link back to information too old to be appropriate. The trend today is toward taking monitorability and control of the credentials process completely away from individuals by allowing organizations to be the repositories of all credential data. Individuals would merely provide the identifying information that allows linking to their own credentials.

The countervailing problem is that credential systems are subject to widespread abuse by individuals, such as the modification and the copying of many kinds of paper and plastic certificates that are made easy by today's technology. This is one reason why certificates are falling into disuse and organizations are maintaining credentials themselves. Information provided without substantiating documents is, of course, the easiest kind to falsify, which may account for the rapid deployment of so-called matching techniques that allow organizations to use identifying information to link and share records. Special problems are raised by credentials that an individual might not care to show; these will be called "negative" credentials. Assuring the absence of negative credentials is often impractical with certificates or even matching. Today, this problem is addressed by centralized information maintainers who attempt to collect reports of negative credentials from all possible issuers. Use of multiple complete identities by sophisticated criminals is a related problem. As with communication and payments, the obvious measures under the current approach for preventing abuse of credentials by individuals—widespread use of highly secure identity documents providing links to centrally maintained credentials—are antithetical to the ability of individuals to determine how information about themselves is used.

The solution is based on an individual's ability to take a specially coded credential issued under one pseudonym and to transform it into a similarly coded form of the same credential that can be shown under the individual's other pseudonyms. Since these coded credentials are maintained and shown only by individuals, they provide control similar to that provided by certificates. Individuals can also tailor the coded form shown so that it provides only the necessary information and can ensure that obsolete information becomes unlinkable to current pseudonyms. Abuses by individuals, such as forgery, improper modification, and sharing, are prevented by the crypto-


graphic coding and by the protocols for such coding. Because these coded credentials are conveniently issued and shown, they can be widely used, obviating the need for unsubstantiated credentials and matching. Centralized maintainers can still determine and issue credentials on the absence of negative credentials, but cannot link to other information. Each person is able to use at most one pseudonym with any organization requiring such protection, thereby effectively preventing use of multiple complete identities. Extensions ensure accountability for abuses created under any of an individual's pseudonyms.

The Basic Credential System

The essential concept again is presented by analogy to carbon-lined envelopes, only this time the envelopes would have windows. First, you make up your pseudonyms at random and write them on a plain slip of paper. When you want to get a credential from an organization, you put the slip of paper in a carbon-lined envelope with a window exposing only the part of the slip bearing the pseudonym you will use with that organization. Upon receiving the envelope from you, the organization makes a special signature in a repeating pattern across the outside of the envelope. The kind of signature pattern indicates the kind of credential the issuing organization decides to give the person whose pseudonym they see through the window; the signature pattern serves as the credential. When you get the envelope back from the issuing organization, you verify the signature pattern. Before showing the credential to an organization, you place the slip in an envelope with a window position exposing only the pseudonym you use with that organization and some of the adjacent credential signature pattern. The receiving organization checks that the appropriate pseudonym and credential signature pattern are recognizable through the window. This approach naturally allows a variety of credentials to be obtained and shown.

You need not show all of your credentials to every organization: You can restrict that which is revealed to only what is necessary. Because of the way the signature patterns repeat across the slip, a recognizable part of every signature pattern appears adjacent to each pseudonym. In providing an envelope to an organization, though, you can limit the view through the window so that only necessary signatures are visible. The credentials visible could simply be limited by blacking out parts of the window, but more flexible restriction is possible in actual systems. You might have a credential that represents your income, for instance. You could transform this credential into a more limited credential indicating only that your income falls within a particular range. An even more powerful kind of restriction allows an organization only to verify that a combination of credentials meeting some requirement is held, without revealing anything to the organization about which sufficient combination is actually held.

An organization can ensure that no individual is able to transact with it under more than one pseudonym. One way an individual could attempt to use more than



Untraceable credential transfers between pseudonyms are illustrated by an analogy to window envelopes and carbon paper. The individual (actually the card in the computerized analog) writes the pseudonyms on a slip and seals it, along with a facing piece of carbon paper, in an envelope the window of which exposes only the pseudonym—523—used with organization X. Organization X then applies a signature (stamp) on the outside of the envelope received, with the choice of “C” as the repeating pattern that indicates the kind of credential issued. The individual verifies the signature re-

turned. When the individual later wishes to show the credential to organization Y, the original envelope and carbon are discarded, and the slip is placed in a new envelope the window of which exposes only the pseudonym—965—used by the individual with Y. Now Y verifies the signature through the window of the envelope and knows that 965 has been issued credential C. Organization Y cannot, however, learn the other pseudonyms written on the slip. Actual computerized systems maintain the unconditional untraceability of pseudonyms across the dashed boundary lines.

one pseudonym with an organization is to use different pseudonyms on the same slip of paper. This is prevented by a standard division of the slip into zones, where each zone is assigned to a particular organization; envelopes are accepted by an organization only if the window exposes the organization's zone, which bears a single indelibly written pseudonym. A second way of attempting to use more than one pseudonym per organization is to use more than one slip. This is prevented by the establishment of an “is-a-person” organization that restricts each person to at most one is-a-person signature. Other organizations only accept envelopes with this signature recognizable through the windows. This is-a-person organization might ensure that it issues no more than one signature per person by taking a thumbprint and checking before giving a signature that the print is not already on file. The collection of thumbprints poses little danger to individuals, since the is-a-person organization cannot link the prints with anything.

The pseudonyms used by individuals are untraceable, in the sense that envelopes give no clue, apart from the signatures shown, about the other randomly chosen

pseudonyms they contain. Of course, the computerization of these systems would provide unconditional untraceability using digital blind signatures on numbers. (More complete details on such systems are presented in [4].)

Credential Clearinghouses

When individuals have similar relationships with many organizations, there is often need for the centralized control provided by a *credential clearinghouse*, an organization that develops credential information about individuals' relationships with its member organizations and provides this information to these organizations. In current practice, clearinghouse functions are performed by such major organizations as credit agencies, bank associations, insurance industry associations, national criminal information systems, and tax authorities. Member organizations typically exchange information with clearinghouses during initiations and terminations of relationships.


For concreteness, consider how a credit clearinghouse might control the use of consumer credit using an extended form of the credential system. The clear-

inghouse gives you a number of *enabling* credentials that in effect say "This person is authorized for \$100 worth of credit. If no resolution credential is returned to us within a year, we will assume that the individual has not repaid." You could provide one such credential to a shop, which then gives you credit worth up to \$100. When you settle your account with the shop some time later, they give you the corresponding *resolution* credential, which you ultimately return to the clearinghouse. An important property of this approach is that the clearinghouse and shops cannot link the credentials; the clearinghouse with the cooperation of all the shops cannot learn which shop you went to, any more than the shop can learn your pseudonym with the clearinghouse, since the enabling and resolution credentials are unconditionally untraceable.

Security against abuse by individuals requires that the enabling credential be prevented from being shown to more than one shop. Otherwise someone could obtain too much credit from a single enabling credential. Similarly, it should not be possible to show a single resolution credential more than once to the clearinghouse, since otherwise someone could convince the

clearinghouse that more debt had been repaid than was in fact repaid.

Now consider how the credentials in the credit-agency example might be handled in terms of the envelope analogy. First, you get an open padlock from the shop you want credit from, lock it through a hole punched in your slip, and provide the locked slip, in the appropriate window envelope, to the clearinghouse. The clearinghouse checks the envelope from the outside to assure that a padlock is locked through the slip and that the window exposes part of a slip bearing your credit-worthy pseudonym, makes the enabling signature on the envelope, and returns the envelope to you. When you provide the slip to the shop in the appropriate envelope, the shop is able to see its padlock through the window (the clearinghouse's window did not allow the padlock to be seen, but only allowed the fact that the padlock was locked through the slip to be felt) and can extend you the credit. Upon settling your account, the shop gives you the key to their padlock. In effect, the key is the resolution credential. It allows you to remove the padlock and return the slip to the clearinghouse in the appropriate envelope. The clearinghouse



Untraceable credentials with numbers also follow the paper analogy. The way transmissions [1.1] and [1.2] are developed is detailed below. The so-called one-way function f is easily computed by a publicly known technique, but its inverse is thought to be infeasible to compute. Organization X determines the validity of both transmissions received by verifying that the first is a signature on the one-way function of the second. Later, X provides the signature for the desired credential on [1.2]. The card verifies the signature and replaces s_x by s_y . Organization Y verifies [3.1] and [3.2] just as X did. When the credential [3.3] is received by Y, it is verified as a signed copy of [3.2]. A special organization Z ensures that the [1.1]s (and the [3.1]s) are of the proper form, but

does not obtain information useful in tracing. First the card supplies many candidates to Z, each of the form $q_n = f(u \cdot s_n^c) \cdot t_n^2$, where u is the special pseudonym used by the individual with Z, and $s_n = f(s_n')$ and $t_n = f(t_n')$, with s_n' and t_n' created at random by the card. When the card later learns which candidates q_n have been selected at random for inspection by Z, the card supplies the corresponding s_n' and t_n' for each. This allows Z to verify that $q_n = f(u \cdot f(s_n')^c) \cdot f(t_n')^2$. If all inspected candidates verify, then Z supplies the signed form of all uninspected candidates. (Extensions further reduce the chance of an improper candidate being signed.) The card transforms a signed candidate into [1.1] by dividing out t_n .

checks that the intact slip is returned without the lock and thus knows that you repaid, though it cannot determine which shop was involved.

Structuring Clearinghouses

Further restrictions on the information available to clearinghouses, as well as better control of abuses by individuals, can be achieved by a partially hierarchical structuring of clearinghouses. There might be clearinghouses for each of the dozen or two major areas of consumer interaction with organizations—areas such as credit, education, social services, tax, insurance, voter registration, licenses, employment, criminal, and even military service. Each such clearinghouse might have a number of subclearinghouses below it. An education clearinghouse, for example, might have subclearinghouses below it for primary, secondary, university, and professional education. An organization interacts with organizations immediately below it in such a hierarchy just as a clearinghouse interacts with its member organizations. Only an initial enabling and final resolution credential were transferred in the previous example of credit clearinghouses, but more generally credentials can be transferred between an organization and the organization hierarchically below it in either direction and at any point during a relationship. Subclearinghouses reduce the amount of detail obtainable by clearinghouses, which reduces the information that can be linked by the combined structure of clearinghouses and subclearinghouses.

Hierarchical structuring can also be used to enforce sanctions against individuals perpetrating abuses with even a single organization. Within a hierarchy of clearinghouses, each would expect to learn of serious abuses against organizations below it by a lack of special periodic “no serious abuse” credentials (or by a lack of resolution credentials); if a clearinghouse receives a complete set of such credentials, it also periodically issues a “no serious abuse” credential. Someone lacking such a credential from the highest level clearinghouses might be refused service by member organizations. A more practical variation allows the same transfers of credentials to be conducted only once in advance, with each organization attaching, in terms of the envelope analogy, a locked padlock. Only when the individual receives the corresponding key to the lock from every organization that attached a lock can all the locks be removed and the credential be shown in the required form without locks. If the keys were required to be made available by organizations at a set interval before the credential is required, time might be provided for clearing up errors and misunderstandings, or even for more formal grievance procedures if needed.

Preventing the Use of Obsolete Information

If individuals change pseudonyms periodically, they cannot be linked to obsolete information. Pseudonyms might be changed on a yearly basis. The initial information associated with new pseudonyms would be provided through the transfer of credentials from previous pseudonyms. The changeovers might be staggered to

allow time for completion of pending business.

There are additional benefits to changing pseudonyms aside from the weeding out of obsolete information. The periodic reduction to essential information also prevents organizations from gradually accumulating information that might ultimately be used to link pseudonyms. Another consequence of individuals transferring all the initial information for a period is that they must then know the requirements for information by each organization, must know where each piece of information comes from, and must consent to each such transfer. Thus, such arrangements ensure that information linkable by each organization is known to and agreed on—that is, that it can be monitored and controlled by individuals.

BROADER ISSUES

Advantages to Individuals

As the public becomes more aware of and familiar with the extent and possibilities of emerging information technology, there should be a growing demand for the kinds of systems described here. Individuals stand to gain in increased convenience and reliability, improved protection against abuses by individuals and organizations, a kind of equal parity with organizations, and, of course, monitorability and control over how information about themselves is used.

Individuals will be free to obtain their card computers from any source, to use whatever other hardware or software they choose, and to interface into the communication system wherever they please.

The techniques already touched on for saving encoded backup copies of a card computer's data are relevant in terms of advantages to individuals. The card computer would create a key to encode the backup copies it issues. A replacement card computer needs only this key and a backup copy to obtain the full capabilities of the original card. The key might be impractical for an owner to remember, since it should be at least 40 digits long. A convenient and reliable arrangement for maintaining the key involves dividing it into parts and giving different parts to various trustees. Unconditionally secure techniques allow various designated subsets of trustees to completely recover the key; other insufficient subsets would thus be unable to learn anything about the key. A sufficient subset of trustees could provide the key to its owner, if so requested. Other subsets might be sufficient to recover the key, the backup data, and the owner's secret authorizing number, enabling the trustees in such subsets to take over the owner's affairs when needed. More generally, such an approach illustrates how an individual's right to designate proxies, a right that is of course enjoyed by organizations, is ensured.

It has been stated that a lost or stolen card computer is of very little use to anyone other than its owner. This is because only the owner need know the secret authorizing number that the card computer requires before allowing a transaction. This number might typically be about six digits long. A reasonably tamper-

resistant part of a card computer might make the card useless as a replacement for a thief's own card and could even make use of physical identification techniques such as fingerprints to prevent anyone but its owner from using it to conduct transactions. Even assuming that sophisticated criminals could extract the information content of tamper-resistant parts of the card, a great many actual trial uses of guessed authorizing numbers with organizations might still be required before the actual number could be determined, making such attacks quite likely to be detected and to fail.

Individuals can always sacrifice their protection by revealing linking information. Of course, the systems discussed here can provide secure relationships without requiring such disclosures. It is even possible under exceptional circumstances for persons accused of abuses under pseudonyms to demonstrate that the pseudonyms are not theirs, without revealing linking information. For example, in communication transactions, people could show that their physical entry to the system was not used for a particular message; in payment transactions, they could show that a payment did not involve their account; and in credential transactions, they could show that a pseudonym was not among the set obtainable under their thumbprint.

Pseudonyms would be used only for the computerized part of ordinary consumer transactions, in a way that would provide acceptable protection against linking. Pseudonym use might be transparent to anyone conducting transactions: People never need to actually see pseudonyms and could usually forget that they were being used. Of course, the scope of the separated relationships enjoyed by individuals need not depend on the actual legal or administrative structure of organizations. But some linking of separate relationships might occur, for example, in the case of a consumer who actually wanted to be recognized, as part of an investigation, or in other exceptional situations. But linking of some relationships does not, in general, allow others to be linked, and the regular changing of pseudonyms already described allows linkings to be shed over time. Naturally, the scope of relationships, as well as such things as the granularity and timing requirements of the transaction systems, must be adjusted to provide the desired kind of separation.

Security under the new approach need not restrict individuals from enjoying the same protections as organizations, and an equal opportunity to use the systems. A payment, for example, could be made between two friends using their card computers without involving any other computer system. A small business would even be able to handle all customer transactions with a card computer.

Advantages to Organizations

Organizations also have much to gain: Transaction systems under the new approach will bring all the advantages of advanced computerization, improve security, and be a force for improved goodwill from the public. Not only do organizations generally have an interest in

maintaining good relations with individuals—in making transactions, they have many of the same interests and concerns as individuals. Thus, the advantages to individuals considered above apply in part to organizations as well.

The mechanisms that the systems described here would use also compare favorably, from the economic stance of organizations, with systems based on the logical extension of the current approach, which require widely trusted tamper-resistant devices at all points of entry to transaction systems. Such requirements also mean substantial agreement, outlay, and commitment to design before widespread use can begin. In addition to the substantial costs and risks of such an approach, early commitment to design usually leads to obsolete technology once systems come into use. Tamper-resistant techniques currently available also require substantial compromise between cost and security. Since mutually trusted and tamper-resistant equipment is not required with the systems described here, any entry point to a system can be freely used; users can even supply their own terminal equipment and take advantage of the latest technology.

The new systems would make more sophisticated use of cryptographic techniques than many proposals under the current approach. But even the difference between the simplest current proposals and the mechanisms required by the systems presented here is just a fraction of a chip in the technologies of the near future. Ordinary microcomputers are already capable of conducting the required protocols for individuals using the communication, payment, and limited versions of the credential systems described here.

Since the sensitivity and the quantity of consumer data in the hands of organizations are reduced, so is their exposure to incidents that might impair public perception or incur legal liability. Reductions in data could also streamline operations, and the increased appropriateness of the remaining data could provide more effective policy and decision procedures. Also, obtaining information needed for decision making by surveys and the like might be more successful in the future under systems ensuring untraceability.

Detected abuses can be dealt with to an extent acceptably close to the limits of any transaction system. Individuals defaulting on requirements or perpetrating serious abuses can always step outside the controls of any transaction system by going "underground." Transaction systems are thus limited to preventing further transactions once an abuse or default threshold is reached. The new approach stops short of approaching this limit because, as has been mentioned, it ensures individuals some time delay—hopefully enough for due process if needed—before all transactions are prevented. The new approach restricts the amount of default possible by providing a desired balance between prior restraint, as in the basic payment system, and accountability after the fact, as with credit and other clearinghouse functions.

Undetected abuses can be restrained to an extent also acceptably close to the limits of any transaction

system. The communication, payment, and credential systems described here seem quite able to prevent undetected abuse by individuals. But no transaction system is able to detect or prevent abuse that results from an individual obtaining something through legitimate use of a system and then transferring it, outside the system, to another person. Transferring the ability to use a communication system to others is an instance of the proxy right already discussed. When such transfers occur in the context of payment transactions, they can be treated as illicit payments, which, as has already been pointed out in the section on payment transactions, can be deterred. The credential system directly prevents the transfer of credentials from the pseudonyms of one person to those of another. Currently, "in-person" proxy is prevented by certificates bearing photos. Such photo tokens could still be used with the new approach, but they would bear only a photo, an indication of the kind of credential, and possibly a pseudonym.

It is too easy to step outside current transaction systems by dealing in cash, using coin phones, sending anonymous letters, and using false credentials. Significant security improvements can only be obtained with comprehensive systems. But such security under the current approach may meet with substantial and broad-based resistance from individuals—particularly with awareness of the alternatives posed by the new approach.

Future Implications

Large-scale automated transaction systems are imminent. The architecture chosen for these systems may have a long-term impact on the centralization of our economic system, on some of our basic liberties, and even on our democracy. The initial choice of direction will gather economic and societal momentum, making reversal increasingly less likely.

Economic centralization may be furthered under the current approach. Computerization has already allowed organizations to grow to unprecedented size and influence. Further computerization could increase aggregation and centralization by allowing service providers and other major actors to obtain far-reaching information about individuals. If this information were partitioned into separate unlinkable relationships, such aggregation and centralization might be reversed. The information age has other significant potential for decentralization: Information gathering and dissemination lack the inherent centralization of earlier technologies, and payment systems integrated with communication systems allow potentially unrestricted access to distribution channels. The new approach offers individuals and small organizations the same access to such services as large organizations.

Some of our basic liberties may be threatened by computerization under the current approach. The interlinking of relationships and the surveillance required just for practical security may become unacceptable. Such surveillance and linking are unnecessary

when information can be made public, scanned, or bought and sold pseudonymously.

The chilling effect of a growing surveillance potential could also decrease expression and participation. The loss of monitorability and control could increase alienation and also decrease participation. Sophisticated marketing techniques that rely on profiles of individuals are now being used to manipulate public opinion and elections. The potential exists not only for reversing these problems, but for increasing democratization. For instance, with the kinds of systems presented here, multiparty secure election and polling could be conveniently conducted without centralized coordination and with those expressing their views able to show relevant credentials.

Advances in information technology have always been accompanied by major changes in society: The transition from tribal to larger hierarchical forms, for example, was accompanied by written language, and printing technology helped to foster the emergence of large-scale democracies. Coupling computers to telecommunications technologies creates what has been called the ultimate medium—it certainly is a big step up from paper. One might ask, To what forms of society could this new technology lead? The two approaches appear to hold quite different answers.

Acknowledgments. The author would like to thank Jan-Hendrick Evertse, Wiebren de Jonge, and Ronald L. Rivest for discussions during the early development of the ideas herein presented, as well as everyone who showed interest in and commented on this work.

REFERENCES

1. Burnham, D. *The Rise of the Computer State: The Threat to Our Freedoms, Our Ethics and Our Democratic Process*. Foreword by Walter Cronkite. Random House, New York, 1983. (Not cited in text.)
2. Chaum, D. The dining cryptographers problem: Unconditional sender and recipient untraceability. Available from the author.
3. Chaum, D. Privacy protected payments: Unconditional payer and/or payee untraceability. Available from the author.
4. Chaum, D. Showing credentials without identification: Transferring signatures between unconditionally unlinkable pseudonyms. Available from the author.
5. Diffie, W., and Hellman, M.E. New directions in cryptography. *IEEE Trans. Inf. Theory* T-IT76 (Nov. 1976), 644–654.
6. Rivest, R.L., Shamir, A., and Adleman, L. A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM* 21, 2 (Feb. 1978), 120–126.

CR Categories and Subject Descriptors: C.2.0 [Computer-Communication Networks]: General—security and protection; E.3 [Data Encryption]: public key cryptosystems; J.1 [Administrative Data Processing]: financial (e.g., EFTS); K.4.1 [Computers and Society]: Public Policy Issues—privacy; K.4.2 [Computers and Society]: Social Issues—abuse and crime involving computers; K.8 [Personal Computing]

General Terms: Security

Additional Key Words and Phrases: blind signatures, card computers, digital pseudonyms, multiparty security, unconditional security, untraceability

Author's Present Address: David Chaum, Centre for Mathematics and Computer Science, Kruislaan 413, 1098 SJ Amsterdam, The Netherlands. Internet: chaum@mcvax.uucp; Telex: 12571 mactrl nl.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.