Trust in the New Information Age

David P. Maher

In promoting new ways of communicating and doing business, AT&T and other creators of new communications systems need to address some issues of fundamental importance to our customers. People need to have a sense of security about the information that they store, transmit, and receive. They must have confidence in its integrity, source, and destination. Commitments implied by an electronic transaction must be backed by authority that is clearly defined and dependable. People and organizations must also be confident that their correspondence and transactions will not be electronically collected and pieced together to form an image that may someday be used to haunt them. These issues touch the very essence of AT&T's business, and they will affect AT&T's success in carrying out its mission. This paper highlights some ways in which security technology can be used to address fundamental problems that arise with the introduction of new information and communications systems.

Introduction

The inexorable force of economics is thrusting us into the new information age. Business efficiency, rather than the availability of technology, is now motivating the use of multimedia teleconferencing, telecollaboration, telecommuting, telepublishing, the paperless office, electronic transactions, etc. The term "information superhighway" seems inadequate for what is now being constructed. The new information systems do allow information to travel at high speeds, but they are also enabling people to be virtually everywhere, and to have information instantly at hand no matter where that information is stored. AT&T and all other companies that are constructing and selling the technology supporting these systems have a stake in ensuring that people feel at ease with them. But information technology, like nuclear energy and biotechnology, can make people feel very uncomfortable. There are several sources of discomfort:

- Information is broadcast over airwaves, transmitted over public and private networks with multiple-access protocols, and stored in places that may be intentionally or unintentionally accessible to the public. It is understandable that the proprietors of that information may need assurance that their assets are being protected.
- As people become aware of how much information is being gathered about them, as they realize that every interaction with a computer system or communications network can be, and often is, recorded and added to their personal profiles, as they watch this "datashadow" that follows them grow beyond their control, their discomfort with the new information age may tend to increase. Indeed, what we buy, what we watch, where we go, our education, employment, credit and medical histories, and much more can be amalgamated into a virtual image of us. If misused, or just misunderstood, this could haunt not only us, but also our progeny.
- As more business is conducted impersonally, using computer networks, trust becomes a significant issue. More and more people are asked to trust things that they do not understand. In place of a warm handshake with a real person, in place of business letters with "recognizable" signatures, people

Panel 1. Abbreviations, Acronyms, and Terms

CRC—cyclic redundancy check

cyberspace—an abstraction within which we imagine that events, associated with electronic computations and communications, happen.

DSA—Digital Signature Algorithm

EES—Escrowed Encryption Standard

hashing—A process whereby a document or message is condensed into a brief digest that can be used as an integrity check.

IVES—Information Vending Encryption System

KCA—key certificate authority

KCS—key certification system

MAC—message authentication code

plaintext—unencrypted text

SHA—Secure Hash Algorithm

spoofing—In cryptographic systems, the process of masquerading as the originator or intended recipient of a message, or the process of substituting a bogus message for part or all of a legitimate message.

are now asked to deal with electronic agents and computerized systems that are hard to comprehend, let alone trust. Moreover, people's natural instincts to wariness appear to be justified; both analysis and experience show that many of these systems can be easily manipulated to perform undesirable actions.

Trust, fundamental to the very concept of commerce, is at the heart of all these issues. As AT&T offers people new ways of doing business, it must assure them that they can proceed with confidence that is justified. Researchers and developers of new information systems need to deal effectively with the issue of trust. How this issue arises, how AT&T is addressing it, and what questions remain to be addressed are the topics of this paper. As one might expect, the issue of trust in the new information age encompasses social and political science, as well as information science.

Information Systems' Properties

Some commercial information systems' properties that relate to trust and are important to electronic commerce are:

- Authenticity of identity — Any entity can reliably ascertain the identity of those involved in a transaction, and of the originator of a document, message, or other piece of information;
- Confidentiality — Information is accessible only to those for whom it is intended;
- Integrity — Information is not modifiable, either accidentally or maliciously, without detection;
- Validity of remittances — Any entity can ascertain that the tender for a transaction is legitimate and reliable; and
- Service availability — The system is protected against denial of service.

Other important properties are also desirable, but those described here are central to the issues of trust raised in this paper.

Problems of Identity As one conducts commerce over a computer network, it is imperative to know the identity and authority of the entities encountered. This is especially true when legal commitments are involved. It is straightforward enough for an electronic entity — such as a computerized ordering system, or a system that collects bids and proposals — to declare its identity. But how do we ensure that this entity represents whom it claims to be? How can we ensure that any transaction that we agreed to electronically cannot later be disputed on the basis of incorrect identity? How can we make such transactions auditable?

Almost every aspect of commerce, except for the actual delivery of tangible goods, can be carried out through multimedia messaging over computer networks. Messages can be created directly by people, as well as by computer programs called electronic agents. If a message is created by an electronic agent, whom does the agent represent, and what authority can be attached to it? Agents are commonly designed to gather information and to ask questions. Other agents are designed to answer questions. By what method can agents trust other agents? In cyberspace, an individual can control several electronic "personae" — electronic agents acting under his or her direction. (See Panel 1 for definitions of abbreviations, acronyms, and terms.) What legal authority do these proxies have? Given that these agents may be moving through networks seeking access to many different kinds of systems, and given that electronic commerce will require equal access, the concept of equal rights for electronic agents is not frivolous. But then, what constitutes legal identity and parentage? What binds the agent to the authority of the entity who created it?

On terra firma, we use both explicit and subtle means to verify identity. Often the means are implicit. If, during a phone conversation, I invite someone to my office, the person who shows up at the appointed time is probably the same person I spoke to on the phone, especially if the voice sounds the same. Several cues in our conversation will probably give me confidence in at least some aspects of my visitor's identity. Certainly I can be fooled, but a complete masquerade takes considerable effort and risk.

Today, masquerading in cyberspace is easy and often entails little risk. Recently, on the Internet, people have been forging names on electronic mail messages threatening President Clinton's life. This was done to disrupt the lives of those whose names were forged. Almost no effort is required to make such messages untraceable.

What are the solutions? Robust identification methods are available today to help to identify people and electronic agents acting through computer networks. These methods use a concept called a digital signature, explained in Andrew Odlyzko's paper in this issue of the *Journal*.[1] See Panel 2 for an explanation of digital signatures.

AT&T's Gretag Data Systems subsidiary has developed a digital signature system for Swiss securities' clearinghouses. It allows brokers to digitally sign orders to buy and sell securities. When an agreement is reached on a given security, the selling broker signs a sell order and forwards it to a clearinghouse; the buying broker signs a buy order and sends it to the same clearinghouse. This clearinghouse then checks the identities of the two parties, records the new ownership, and saves the digitally signed orders as proof of the transaction. This type of system eliminates the expense of issuing and keeping paper stock certificates, which speeds up the entire process.

The ephemeral nature of computer messages will make digital signatures commonplace in open electronic commerce systems. To solve the identity problem completely, however, we need a complex infrastructure to certify and revoke signatures. Because of the subtleties involved in building this infrastructure, a practical system, used and recognized worldwide, does not yet exist. AT&T researchers are currently examining the problems that need to be solved to provide the support necessary for a universal digital signature system.

Panel 2. Digital Signatures

A document or message can be electronically signed by an individual in such a way that the signature can be publicly verified. The sender can electronically sign a document and send it to the recipient, along with a certificate that allows the recipient to verify the accompanying signature. This certificate is nothing more than another electronic document that binds the sender's name, address, title, and other pertinent information to the signature verification function that he or she includes. This latter document is itself signed by someone whom the recipient recognizes and trusts. These certificates can be used to certify roles and to explicitly identify authority and limits of authority. A certificate can say, in effect, "the included mathematical function can be used to verify that a given document was signed by a person who represents ABC corporation as its purchasing agent. This person has the authority to sign contracts of value up to $1,000,000, but no more." If a digitally signed document is altered after it is signed, the value of the verification function applied to the document will indicate that fact. However, one cannot determine whether the document was intentionally altered, or was changed as a result of a transmission or memory error.

In some systems, certification of signatures has a deep hierarchy. Person A in my department might sign my certificate, but Person B at AT&T headquarters might sign A's credentials to issue certificates for my department. Person C at some nationally recognized agency might certify B's role as the certificate issuer for AT&T, and so on. When I check someone's credentials, I can go up the hierarchy as far as necessary.

Meanwhile, smaller systems can use the "speakeasy" method of identification: Messages can be encrypted using a "secret word." Anyone who knows the secret word well enough to encrypt a message using it is identified as a member of the club. Unlike the digital signature method alluded to earlier, this may not guarantee the integrity of the message itself, but the message does bear some sort of signature. The major drawback to this approach is the logistics of getting the secret words distributed.

Figure 1. A cheap attack on public-key exchanges.

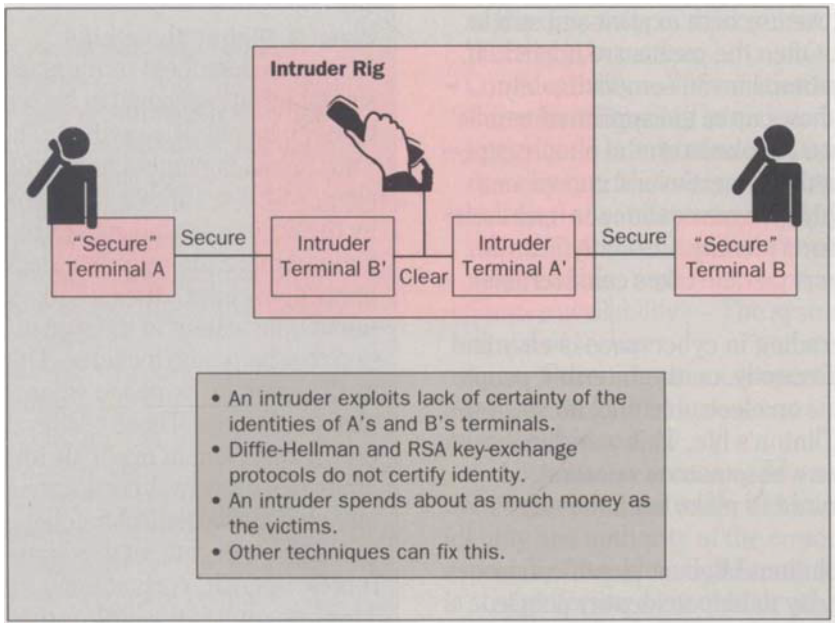

Digital signatures do not solve all identity problems. A good example is the cloning problem, which plagues the cable television and cellular radio industries. Cable converter boxes often use encryption to ensure that premium services are delivered only to paying customers. The correct cryptographic keys are required to enable a customer to tune into a premium channel, or to watch a "pay-per-view" program. Because the converter boxes are readily available, clone boxes can be created using traditional reverse engineering techniques, complete with all cryptographic information and any signatures or signing functions. These clone boxes react to broadcast signals precisely the way the original cable converter boxes do, and they have access to the same information. Unfortunately, the user of the clone box is not easy to identify, and is typically not billed. Current solutions to the cloning problem do not appear to be as robust as solutions to other identity problems. Converters are designed to make cloning difficult, but judging from the number and variety of converter box clones, a successful approach to this problem has yet to be devised.

AT&T's Information Vending Encryption System (IVES), under development in the AT&T Secure Communications Systems Development Laboratory, uses a system-level solution that:

- Minimizes the payoff for successful cloning,
- Uses clone discovery methods (detecting the existence of perfect clones),
- Can disable clones, and
- Makes it costly to clone information receiver circuitry.

Combined, these points make the system solution effective. IVES is being developed to ensure that when information or entertainment is sold, only paying customers have access to it, protecting the investment of the information vendor, as well as the value of the information itself.

Problems of Confidentiality There are many significant reasons why confidentiality is so important in the new information age. Free-market economies are becoming more prevalent, and access to information can determine success in the marketplace. Companies within this environment are motivated to protect their intellectual assets, and to gather as much intelligence as they can. Both the means and the opportunities for intelligence gathering are growing.

With faster computers, high-capacity storage systems, and new, sophisticated data analysis techniques, much more information can now be assimilated and used. This has tremendous implications for personal privacy, as well as for commercial competition.

Information, including intellectual and creative property, is more routinely being sold in electronic form,

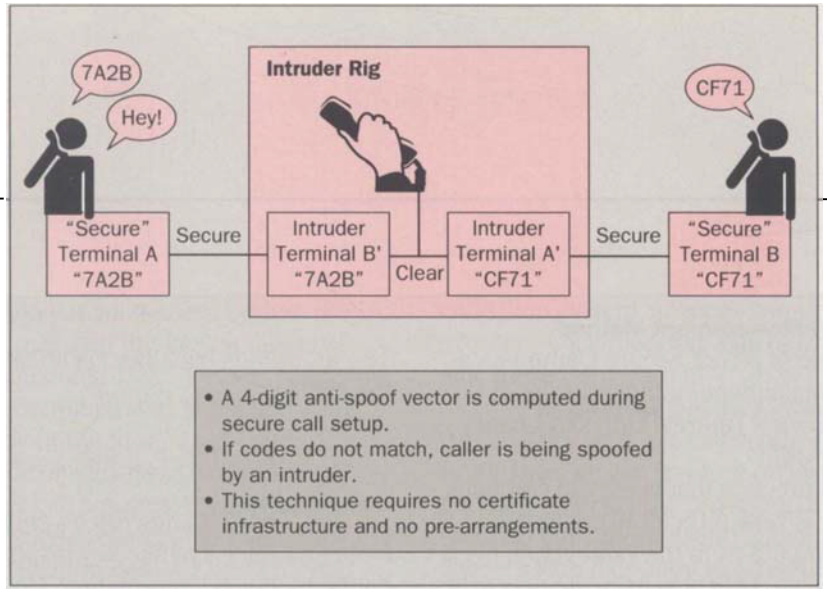

Figure 2. Protection against public-key attacks.

not just as an accessory to a business, but as the business itself. Theft of information in this context equates to avoidance of payment. Software piracy and theft of cable television service are common examples.

A system supports confidentiality if it has effective ways of controlling access to information. Various means may be employed. Physical barriers may be used, but they have no effect on most networks, especially worldwide computer networks designed to support open electronic commerce. Such techniques as the use of passwords to control access to a network are notoriously ineffective. Several types of access control devices use challenge–response protocols and/or biometric functions, such as retinal scans and hand-geometry analysis. Though these can be effective, they can also be costly. Often they do not provide fine control of access to information, because the information in computers may not be properly segregated. For example, the Computer Emergency Response Team at Carnegie Mellon University recently reported that the microphones attached to many workstations are connected directly to voice-sample buffers, which, by default, have no access control (i.e., they have "world access"). As a result, the voice information is readable by anyone on the network. It is not likely that people holding conversations near these workstations are aware that their conversations are being broadcast, sometimes internationally!

In contrast, access to information that is encrypted using a strong encryption algorithm is limited to those who have access to the decryption key. Everything from a complete file system down to individual fields in database records can be encrypted. Keys can be shared by many or by a few, increasing the span and precision of access control using encryption. When the cryptographic algorithm (hiding function) is effective, the access control problem is reduced to one of effective key management.

Distribution of keys over a computer network presupposes a solution to the identity problem, because spoofing legitimate recipients in a key distribution system is a common cheap attack. The Diffie-Hellman key distribution system[2] is commonly used, but it is highly susceptible to a cheap spoofing attack, necessitating supplementary anti-spoof measures (see Figures 1 and 2). AT&T has designed a certificate-based key distribution system for use in messaging systems and cryptographic file systems. It can support many of AT&T's secure voice and data products (see Panel 3). When a user receives credentials from a key certification system (KCS), he or she automatically has a distinct encryption key for everyone else in the system.

Suppose that AT&T's key management system had a million users. Implicit in the few hundred bytes in a given user's credentials is the ability to construct 1 million symmetric encryption keys, one for every user in the system (including the originator). Once two users, Alice and Bob, get one another's public certificates, each of them can generate a unique, symmetric cryptographic key without further communication with one another, or with the KCS. When the KCS produces and distributes 1 million sets of credentials, it implicitly generates and distributes about a half-trillion distinct symmetric cryptographic keys. In this system, one can encrypt a message just once, and then send it to many users, each of whom can decrypt it. The system also has a means of "cryptographic backup." If a user loses his key, or is unavailable, a designated backup agent (but no one else) can reconstruct the key.

Problems of Information Integrity In electronic commerce systems, it is important to ensure that documents and messages cannot be altered, either unintentionally or as part of an attempt to deceive someone. A system lacking this capability is not auditable. Often, physical protection of transaction records can suffice. As more informa-

Panel 3. A Sample Key Management Method

Here's how one of AT&T Secure Communications Systems' key management methods works. It may be useful to first review Andrew Odlyzko's paper in this issue of the *Journal*.

Let $f(x,r)$ be a function that is easy to compute, but whose inverse in r is very difficult to compute. That is, r is hard to compute from the value $f(x,r)$, even when the value x is known. Let f also have the property that for two values, r_1 and r_2, $f(f(x,r_1), r_2) = f(f(x,r_2), r_1)$. An example of such a function with these two properties is exponentiation of the number x to the rth power, mod p, where p is a large prime number. The Diffie-Hellman key exchange method[2] is based on this function and these two properties.

Let $h()$ be a one-way hash function, such as the secure hash algorithm (SHA), specified by the National Institute of Standards and Technology in FIPS PUB 180.[3] This algorithm is suitable for producing message digests.

Let $s(ks,)$ be a digital signature signing function, such as specified in the Digital Signature Algorithm (DSA), part of FIPS PUB 186.[4] Here k_s is the secret signing key for a given entity. Let $v(k_p,)$ be the corresponding verification function, based on the public key k_p.

When Alice registers with the key management system, she obtains system parameters x and p, and she generates her own random secret key, r_A. She fills out a form entering her ID information (name, address, title, etc.) and her public encryption key, and sends the form, notarized, to the key certificate authority (KCA). Symbolically, Alice sends:

$$(ID_A, f(x,r_A))$$

Note that Alice does not reveal her secret key, r_A. The KCA sends back her certificate, which consists of the triple

$$CA = (ID_A, f(x, r_A), s(k_{s,KCA}, h((ID_A, f(x,r_A))))).$$

The KCA adds Alice's certificate to a directory. Note that Alice's public key is hashed together, and bound to, her ID information. Now Bob, who already has a certificate from the KCA, wants to send a secret message to Alice. He can do so by looking into the KCA's directory (or any other appropriate directory) for Alice's certificate, or by getting it from Alice directly (which happens automatically when he calls Alice's secure telephone). Bob then validates Alice's certificate by computing $v(k_{p,KCA}, C_A)$. He can then use the symmetric encryption key, $f(f(x,r_A),r_B)$, because Alice is the only other person who can compute that key. In general, if there are n users of this system, there are $n(n+1)/2$ unique, symmetric encryption keys that can be derived from the key information in the system. Each person can generate n unique, certified keys, one for everyone else in the system, including himself or herself. If Alice receives a message from Bob that he has encrypted using the symmetric key derived earlier, she can verify that it came from Bob, but she cannot demonstrate to someone else that it came from Bob, unless she reveals her secret key, r_A, or Bob reveals his. By adding public signature verification functions to the certificate, Bob can sign the message, using a digital signature function certified by the KCA, and Alice can then prove to anyone who trusts the KCA that the message was sent by Bob.

These techniques have been extended to support cryptographic file systems, secure fax protocols, secure message broadcasting systems, and other applications.

tion is sent over publicly accessible wide-area networks, or stored in common backup facilities and distributed files systems, this becomes a difficult, if not impossible, task. Encrypting documents can sometimes help to detect unauthorized alteration, but it is possible to alter the underlying *plaintext* (the original text) without decrypting the document. Techniques typically used to detect random errors (such as cyclic redundancy checks [CRCs]) are vulnerable to an alteration attack. However, if a document is digitally signed or includes a message authentication code (MAC), alterations can be detected.

MACs, commonly found in banking systems, use symmetric key cryptographic techniques to provide a 64-bit digest of a document that cannot be constructed without access to the key. When a digital signature is used, a digest produced by a hash function (without a crypto-

graphic key) is signed. A key management system is required to produce and maintain the key elements of the MAC and signature functions.

As Woody Allen demonstrated in his film *Zelig*, and as seen more recently in the film *Forest Gump*, it is possible to convincingly edit one image into other images, placing a fictional character into newsreels of important past events. Digital editing techniques continue to make this easier to do and more difficult to detect. Direct digital photography makes it almost impossible to detect many types of alterations. Because photographs are used as evidence in court, soon it may be necessary to have cameras sign electronic photographs as they are produced.

Validity of Remittances Remittance, or payment for goods or services is, of course, fundamental to commerce, electronic or not. In electronic commerce, transactions can occur so fast and money can be whipped around the world so quickly that purely electronic systems have unique problems. Traditional electronic payment systems that use credit card and debit card numbers are vulnerable to number-snatching attacks. Service terminals (e.g., cellular handsets) can be cloned. There is not just a *potential* for fraud, there *is* a lot of fraud, and it is growing.

Identifying a customer reliably can go a long way towards solving remittance problems. Electronic payment systems can properly identify a buyer that provides digitally signed, certified checks using methods such as those discussed earlier. For many people, however, solving the problem using robust identity techniques is not the answer, because they feel that computers are already gathering too much information about us. We can understand the concerns of people who are being persuaded to use a smart card for toll collection in a smart highway system. They do not want the computer to be able to track their whereabouts, because that is an obvious invasion of privacy. Consequently, many companies are proposing anonymous electronic payment systems, which reliably solve the remittance problems of vendors, but keep the identity of customers secret, even from the banks who hold the cash. Such systems seem "just the ticket" for smart highways.

On the other hand, free availability of anonymous remittance systems, especially electronic cash systems, in which the payee's identity can also be kept secret, can support organized crime by hiding massive money-laundering schemes. It can also threaten our tax collection system, as more commerce moves "underground." On whichever side of the political fence one sits, these issues are more than technical.

Problems of Service Availability Electronic services can be disrupted. We can create garbage in cyberspace faster than we can create it physically. Not only can that garbage obstruct the provision of services to others, it can also bury more useful information. For example, two lawyers who flagrantly violated the rules of the Internet against advertising have been subjected to "e-mail bombs," constructed by Internet vigilantes. Their intent is to overload the lawyers' computers, fax machines, and telephones with hundreds of thousands of messages. Such attacks, categorized as denial of service, can be extremely effective on the Internet, where messaging is free, and anonymity is supported.

On any network where anonymity is supported as an absolute right, this kind of abuse, justified or not, is possible because perpetrators are free from fear of reprisal, and from attempts at halting their behavior. Since entrepreneurs are now developing the Internet for its commercial potential, the possibility of Internet regulation is being raised. While such regulation would solve the problem and make service more reliable, it would also destroy the free and open nature of the network. This problem provokes the question: In this context, can we trust freedom? How much, and what type of, network regulation is necessary?

The Limits of Technology

As the discussion of remittance illustrates, new technology is engendering even more technology to solve the problems it creates, so that we can enjoy the original technology's benefits. But the implications of the solutions have side effects that can shake up the social equilibrium.

Yet another example is the concept of an electronic health ID card. It has many advantages, including simplification of payment systems, and amalgamation of extensive, detailed health records. These cards, which can be used to save a person's life in an emergency, and to optimize treatment of many disorders, will contain intimate details of an individual's background. For example, these cards will be able to store the results of many types of genetic tests, identifying gene configurations that, owing to data that identify behavioral correlations, can subject an individual, as well as his or her progeny, to

unfair, subtle, but effective discrimination. Will we be able to design a system that permits legitimate access to this information, but denies availability to those who could misuse it? Can a control system be designed to ensure quick, reliable access by whatever number of people could potentially need the information during an emergency?

With the recent introduction by the U.S. Government of an encryption device known as the Clipper chip , a national debate has ensued on the concept of cryptographic key escrow systems. The new Escrowed Encryption Standard (EES)[5] specifies how a law enforcement agency with a legal wiretap order can get access to a cryptographic key used to encrypt conversations, faxes, or messages. The public has been told by some that they have nothing to fear because these systems do not change the balance of power between the Government and its citizens. Others insist that the system will permit the Government to monitor all our telephone conversations and invade our privacy. Neither of these contradictory positions is supportable, yet both persist. The fact remains that cryptography is rapidly advancing, and is coming of age.

Some powerful capabilities are becoming much more accessible to the public, allowing individuals to protect their privacy just when it appears that it is vitally needed. Yet criminals can use these same capabilities as invulnerable shields. Criminals who roam cyberspace are opportunists, snatching advantage away from law-abiding citizens who may momentarily drop their guard. These criminals can dart behind their cryptographic protective structures, safely awaiting the arrival of their next prey. Are we to provide law-enforcement officials with no means of defending us? Once again, it seems to be a matter of trust.

Conclusion

The information age is pushing us into the new world of cyberspace. People meet, work, and transact business there. Workflow is managed there. Data is stored and organized there. People get their information there. Our "datashadows" are there, and virtual reality is there. What will it take for us to feel comfortable there? As we rapidly provide new ways of sharing information, will we be able to adequately protect our secrets, as well as our intellectual and creative property? Can we construct recognizable, controllable protective structures that shelter us from the adversity in cyberspace? Can we know whom and what to trust? Some possibilities are raised in this paper, but we have no definitive answers yet. These questions are explored elsewhere in this issue.

In this paper, we have stressed that information technology, like nuclear energy and bioengineering, can be dangerous. It has a dark side, and it can be misused. Cryptographic techniques have enormous potential, but only recently has this technology been used extensively outside military organizations. We are just beginning to learn how to use it to efficiently solve the problems discussed here.

References

1. A. Odlyzko, "Public Key Cryptography," *AT&T Technical Journal*, Vol. 73, No. 5, pp. 17–23.
2. Bruce Schneier, *Applied Cryptography*, John Wiley and Sons, 1994.
3. *FIPS Pub 180, Federal Information Processing Standard*, National Institute of Standards and Technology, 1994.
4. *FIPS Pub 186, Federal Information Processing Standard*, National Institute of Standards and Technology, 1994.
5. *FIPS Pub 185, Federal Information Processing Standard*, National Institute of Standards and Technology, 1994.

(Manuscript approved August 1994)

***David P. Maher**, chief scientist for AT&T Secure Communications Systems, is located in Largo, Florida. He is responsible for the architecture of new security products, systems, and services. Mr. Maher received a B.A., M.S., and Ph.D. in mathematics from Lehigh University, Bethlehem, Pennsylvania. He joined AT&T in 1981 and, in 1991, was made a Bell Laboratories Fellow for his work on platforms for secure communications products.*