SYSTEM IMPLICATIONS OF INFORMATION PRIVACY

H. E. Petersen
R. Turn

April 1967

P-3504

SYSTEM IMPLICATIONS OF INFORMATION PRIVACY

H. E. Petersen*
R. Turn*

The RAND Corporation, Santa Monica, California

ABSTRACT

Various questions of providing information privacy for remotely accessible on-line, time-shared information systems are explored. Such systems, especially the remote terminals and the communication network, are vulnerable to threats to privacy ranging from accidental dumping of information as a result of hardware or software failures to deliberate penetration using sophisticated equipment. Deliberate attacks are to be expected since payoff from obtained, altered, or erased information could be high. The resources required vary from the cost of a tape recorder to a large investment in equipment and knowhow.

A range of protective countermeasures is discussed, and their choice and implication considered. It appears possible to counter a given level of threat without unreasonable expenditures of resources. The protective techniques discussed

---

*Any views expressed in this Paper are those of the authors. They should not be interpreted as reflecting the views of The RAND Corporation or the official opinion or policy of any of its governmental or private research sponsors. Papers are reproduced by The RAND Corporation as a courtesy to members of its staff.

This Paper was prepared for presentation at the Spring Joint Computer Conference, Atlantic City, April 17-19, 1967.

in this Paper include: shielding to reduce electro-magnetic emanations; use of once-only passwords for access control; application of privacy transformations to conceal information in user-processor communications and in data files; recording of attempted penetrations; and systematic verification of the hardware and software integrity. It appears possible to engineer various privacy protection techniques into information systems so that the cost of protection is proportional to the amount received, and is borne largely by those users who desire privacy for their communications and/or files.

SYSTEM IMPLICATIONS OF INFORMATION PRIVACY

H. E. Petersen
R. Turn
The RAND Corporation, Santa Monica, California

I. INTRODUCTION

Recent advances in computer time-sharing technology promise information systems which will permit simultaneous on-line access to many users at remotely located terminals. In such systems, the question naturally arises of protecting one user's stored programs and data against unauthorized access by others. Considerable work has already been done in providing protection against accidental access due to hardware malfunctions or undebugged programs. Protection against deliberate attempts to gain access to private information, although recently discussed from a philosophical point of view,[1-4] has attracted only fragmentary technical attention. This paper presents a discussion of the threat to information privacy in non-military information systems, applicable countermeasures, and system implications of providing privacy protection.

The discussion is based on the following model of an information system: a central processing facility of one or more processors and an associated memory hierarchy; a set of information files--some private, others shared by a number of users; a set of public or private query terminals at geographically remote locations; and a communication network of common carrier, leased, or private

lines. This time-shared, on-line system is referred to throughout as "the system."

II. THREATS TO INFORMATION PRIVACY

Privacy of information in the system is lost either by accident or deliberately induced disclosure. The most common causes of accidental disclosures are failures of hardware and use of partially debugged programs. Improvements in hardware reliability and various memory protection schemes are generally suggested as countermeasures. Deliberate efforts to infiltrate an on-line, time-shared system can be classified as either passive or active.

Passive infiltration may be accomplished by wiretapping or by electromagnetic pickup of the traffic at any point in the system. Although considerable effort has been applied to counter such threats to defense communications, non-governmental approaches to information privacy usually assume that communication lines are secure, when in fact they are the most vulnerable part of the system. Techniques for penetrating communication networks may be borrowed from the well-developed art of listening in on voice conversations.[5,6] (While the minimum investment in equipment is higher than that required to obtain a pair of headsets and a capacitor, it is still very low since a one-hundred-dollar

tape recorder and a code conversion table suffice.) Clearly, digital transmission of information does not provide any more privacy than, for example, Morse code. Nevertheless, some users seem willing to entrust to digital systems valuable information that they would not communicate over a telephone.

Active infiltration--an attempt to enter the system to directly obtain or alter information in the files--can be overtly accomplished through normal access procedures by:

- Using legitimate access to a part of the system to ask unauthorized questions (e.g., requesting payroll information or trying to associate an individual with certain data), or to "browse" in unauthorized files.
- "Masquerading" as a legitimate user after having obtained proper identifications through wiretapping or other means.
- Having access to the system by virtue of a position with the information center or the communication network but without a "need to know" (e.g., system programmer, operator, maintenance, and management personnel).

Or an active infiltrator may attempt to enter the system covertly (i.e., avoiding the control and protection programs) by:

- o Using entry points planted in the system by unscrupulous programmers or maintenance engineers, or probing for and discovering "trap doors" which may exist by virtue of the combinatorial aspects of the many system control variables (similar to the search for "new and useful" operation codes--a favorite pastime of machine-language programmers of the early digital computers).
- o Employing special terminals tapped into communication channels to effect:
  - "piggy back" entry into the system by selective interception of communications between a user and the processor, and then releasing these with modifications or substituting entirely new messages while returning an "error" message;
  - "between lines" entry to the system when a legitimate user is inactive but still holds the communication channel;
  - cancellation of the user's sign-off signals, so as to continue operating in his name.

In all of these variations the legitimate user provides procedures for obtaining proper access. The infiltrator is limited, however, to the legitimate user's authorized files.

More than an inexpensive tape recorder is required for active infiltration, since an appropriate terminal and entry into the communication link are essential. In fact, considerable equipment and know-how are required to launch sophisticated infiltration attempts.

The objectives of infiltration attempts against information systems have been discussed by a number of authors[1-3,7] from the point of view of potential payoff. We will merely indicate the types of activities that an infiltrator may wish to undertake:

- Gaining access to desired information in the files, or discovering the information interests of a particular user.
- Changing information in the files (including destruction of entire files).
- Obtaining free computing time or use of proprietory programs.

Depending on the nature of the filed information, a penetration attempt may cause no more damage than satisfying

the curiosity of a potentially larcenous programmer. Or it may cause great damage and result in great payoffs; e.g., illicit "change your dossier for a fee," or industrial espionage activities. (See Table I for a summary of threats to information privacy.)

More sophisticated infiltration scenarios can be conceived as the stakes of penetration increase. The threat to information privacy should not be taken lightly or brushed aside by underestimating the resources and ingenuity of would-be infiltrators.

III. COUNTERMEASURES

The spectrum of threats discussed in the previous section can be countered by a number of techniques and procedures. Some of these were originally introduced into time-shared, multi-user systems to prevent users from inadvertently disturbing each other's programs,[8] and then expanded to protect against accidental or deliberately induced disclosures of information.[8,9] Others found their beginning in requirements to protect privacy in communication networks.[10] In the following discussion, we have organized the various countermeasures into several classes.

Table I

SUMMARY OF THREATS TO INFORMATION PRIVACY

Nature of Infiltration	Means	Effects
Accidental	Computer malfunctioning; user errors; undebugged programs	Privileged information dumped at wrong terminals, printouts, etc.
Deliberate Passive	Wiretapping, electromagnetic pickup, examining carbon papers, etc.	User's interest in information revealed; content of communications revealed
Deliberate Active	Entering files by: "Browsing"; "Masquerading"; "Between lines"; "Piggy-back" penetration	Specific information revealed or modified as a result of infiltrator's actions

ACCESS MANAGEMENT

These techniques are aimed at preventing unauthorized users from obtaining services from the system or gaining access to its files. The procedures involved are authorization, identification, and authentication. Authorization is given for certain users to enter the system, gain access to certain files, and request certain types of information. For example, a researcher may be permitted to compile earnings statistics from payroll files but not to associate names with salaries. Any user attempting to enter the system must first identify himself and his location (i.e., the remote terminal he is using), and then authenticate his identification. The latter is essential if information files with limited access are requested; and is desirable to avoid mis-charging of the computing costs. The identification-authentication steps may be repeated any number of times (e.g., when particularly sensitive files are requested).

Requirements for authentication may also arise in reverse, i.e., the processor identifies and authenticates itself to the user with suitable techniques. Applied to certain messages from the processor to the user (e.g., error messages or requests for repetition) these could be authenticated as coming from the processor and not,

for example, from a piggy-back penetrator. (Since various specific procedures are applied to different parts of the system, a detailed discussion of their structures and effectiveness is presented in Sec. IV.)

PROCESSING RESTRICTIONS

Although access control procedures can eliminate the simple threats from external sources, they cannot stop sophisticated efforts nor completely counter legitimate users or system personnel inclined to browse. An infiltrator, once in the system will attempt to extract, alter, or destroy information in the files. Therefore, some processing restrictions (in addition to the normal memory protection features) need to be imposed on files of sensitive information. For example, certain removable files may be mounted on drives with disabled writing circuits, and alterations of data performed only after requests are authenticated by the controller of each file. Copying complete files (or large parts of files) is another activity where processing controls need to be imposed--again in the form of authentication by file controllers.

In systems where very sensitive information is handled, processing restrictions could be imposed on specific users in instances of "suspicious" behavior. For example, total

cancellation of any program attempting to enter unauthorized files may be an effective countermeasure against browsing.[9]

THREAT MONITORING

Threat monitoring concerns detection of attempted or actual penetrations of the system or files either to provide a real-time response (e.g., invoking job cancellation, or starting tracing procedures) or to permit post facto analysis. Threat monitoring may include recording of all rejected attempts to enter the system or specific files, use of illegal access procedures, unusual activity involving a certain file, attempts to write into protected files, attempts to perform restricted operations such as copying files, excessively long periods of use, etc. Periodic reports to users on file activity may reveal possible misuse or tampering, and prompt stepped-up auditing along with a possible real-time response. Such reports may range from a page by page synopsis of activity during the user session, to a monthly analysis and summary.

PRIVACY TRANSFORMATIONS

Privacy transformations[10,11] are techniques for coding the data in user-processor communications or in files to conceal information. They could be directed against passive (e.g., wiretapping) as well as sophisticated active threats

(e.g., sharing a user's identification and communication link by a "piggy-back" infiltrator), and also afford protection to data in removable files against unauthorized access or physical loss.

A privacy transformation consists of a set of reversible logical operations on the individual characters of an information record, or on sets of such records. Reversibility is required to permit recovery (decoding) of the original information from the encoded form. Classes of privacy transformations include:

- o Substitution--replacement of message characters with characters or groups of characters in the same or a different alphabet in a one-to-one manner (e.g., replacing alphanumeric characters with groups of binary numerals).
- o Transposition--rearrangement of the ordering of characters in a message.
- o Addition--using appropriate "algebra" to combine characters in the message with encoding sequences of characters (the "key") supplied by the user.

Well-known among a number of privacy transformations of the "additive" type[11] are the "Vigenere cipher," where a short sequence of characters is repeatedly used

combine with the characters of the message, and the "Vernam system," where the user-provided sequence is at least as long as the message. Successive applications of several transformations may be used to increase the complexity.

In general, the user of a particular type of privacy transformation (say, the substitution of characters in the same alphabet) has a very large number of choices of transformations in that class (e.g., there are 26 factorial--4×10^{26}--possible substitution schemes of the 26 letters of the English alphabet). The identification of a particular transformation is the "key" chosen by the user for encoding the message.

Any infiltrator would naturally attempt to discover the key--if necessary, by analyzing an intercepted encoded message. The effort required measures the "work factor" of the privacy transformation, and indicates the amount of protection provided, assuming the key cannot be stolen, etc. The work factor depends on the type of privacy transformations used, the statistical characteristics of the message language, the size of the key space, etc. A word of caution against depending too much on the large key space: Shannon[11] points out that about 3×10^{12} years would be required, on the average, to discover the key used in the aforementioned

substitution cipher with 26 letters by an exhaustive trial-and-error method (eliminating one possible key every microsecond). However, according to Shannon, repeatedly dividing the key space into two sets of roughly an equal number of keys and eliminating one set each trial (similar to coin-weighing problems) would require only 88 trials.

The level of work factor which is critical for a given information system depends, of course, on an estimate of the magnitude of threats and of the value of the information. For example, an estimated work factor of one day of continuous computation to break a single key may be an adequate deterrent against a low-level threat.

Other criteria which influence the selection of a class of privacy transformations are:[11]

- o Length of the key--Keys require storage space, must be protected, have to be communicated to remote locations and entered into the system, and may even require memorization. Though generally a short key length seems desirable, better protection can be obtained by using a key as long as the message itself.
- o Size of the key space--The number of different privacy transformations available should be as

large as possible to discourage trial-and-error approaches, and to permit assignment of unique keys to large numbers of users and changing of keys at frequent intervals.

- o Complexity--Affects the cost of implementation of the privacy system by requiring more hardware or processing time, but may also improve the work factor.
- o Error sensitivity--The effect of transmission errors or processor malfunctioning may make decoding impossible.

Other criteria are, of course, the cost of implementation and processing time requirements which depend, in part, on whether the communication channel or the files of the system are involved (see Sec. IV).

INTEGRITY MANAGEMENT

Important in providing privacy to an information system is verification that the system software and hardware perform as specified--including an exhaustive initial verification of the programs and hardware, and later, periodic checks. Between checks, strict controls should be placed on modifications of software and hardware. For example, the latter may be kept in locked cabinets equipped with alarm devices.

Verification of the hardware integrity after each modification or repair should be a standard procedure, and inspection to detect changes of the emissive characteristics performed periodically.

Integrity of the communication channel is a far more serious problem and, if common carrier connections are employed, it would be extremely difficult to guarantee absence of wiretaps.

Personnel integrity (the essential element of privacy protection) poses some fundamental questions which are outside the scope of this paper. The assumptions must be made, in the interest of realism, that not everyone can be trusted. System privacy should depend on the integrity of as few people as possible.

IV. SYSTEM ASPECTS OF INFORMATION PRIVACY

As pointed out previously, not all parts of an information system are equally vulnerable to threats to information privacy, and different countermeasures may be required in each part to counter the same level of threat. The structure and functions of the information processor, the files, and the communication network with terminals are, in particular, sufficiently different to warrant separate discussion of information privacy in these subsystems.

COMMUNICATION LINES AND TERMINALS

Since terminals and communication channels are the principal user-to-processor links, privacy of information in this most vulnerable part of the system is essential.

Wiretapping

Many users spread over a wide area provide many opportunities for wiretapping. Since the cost of physically protected cables is prohibitive, there are no practical means available to prevent this form of entry. As a result, only through protective techniques applied at the terminals and at the processor can the range of threats from simple eavesdropping to sophisticated entry through special terminals be countered. While a properly designed password identification-authentication procedure is effective against some active threats, it does not provide any protection against the simplest threat--eavesdropping--nor against sophisticated "piggy-back" entry. The only broadly effective countermeasure is the use of privacy transformations.

Radiation

In addition to the spectrum of threats arising from wiretapping, electromagnetic radiation from terminals must be considered.[12]

Electromagnetic radiation characteristics will depend heavily on the type of terminal, and may in some cases pose serious shielding and electrical-filtering problems. More advanced terminals using cathode ray tubes for information display may create even greater problems in trying to prevent what has been called "tuning in the terminal on Channel 4."

Use of privacy transformations also helps to reduce some of the problems of controlling radiation. In fact, applying the transformation as close to the electro-mechanical converters of the terminal as possible minimizes the volume that must be protected, and reduces the extent of vulnerable radiation characteristics.

Obviously, the severity of these problems depends upon the physical security of the building or room in which the terminal is housed. Finally, proper handling and disposal of typewriter ribbons, carbon papers, etc., are essential.

Operating Modes

Whether it would be economical to combine both private and public modes of operation into a single standard terminal is yet to be determined; but it appears desirable or even essential to permit a private terminal to operate in the public mode, although the possibility of compromising the privacy system must be considered. For example, one can

easily bypass any special purpose privacy hardware by throwing a switch manually or by computer, but these controls may become vulnerable to tampering. The engineering of the terminal must, therefore, assure reasonable physical, logical, and electrical integrity for a broad range of users and their privacy requirements.

Terminal Identification

An unambiguous and authenticated identification of a terminal is required for log-in, billing, and to permit system initiated call-back for restarting or for "hang-up and redial"[13] access control procedures. The need for authentication mainly arises when the terminal is connected to the processor via the common-carrier communication lines, where tracing of connections through switching centers is difficult. If directly wired connections are used, neither authentication nor identification may be required, since (excluding wiretaps) only one terminal can be on a line.

Identification of a terminal could involve transmission of an internally (to the terminal) generated or user-entered code word consisting, for example, of two parts: one containing a description or name of the terminal; the other, a password (more about these later) which authenticates that the particular terminal is indeed the one claimed in

the first part of the code word. Another method suggested for authenticating the identity of a terminal is to use computer hang-up and call-back procedures.[13]

After terminal identification has been satisfactorily established, the processor may consult tables to determine the privacy level of the terminal; i.e., the users admitted to the terminal, the protection techniques required, etc.

User Identification

As with a terminal, identifying a user may require stating the user's name and account number, and then authenticating these with a password from a list, etc.

If the security of this identification process is adequate, the normal terminal input mechanisms may be used; otherwise, special features will be required. For example, hardware to accept and interpret coded cards might be employed, or sets of special dials or buttons provided. Procedures using the latter might consist of operating these devices in the correct sequence.

In some instances, if physical access to a terminal is appropriately controlled, terminal identification may be substituted for user identification (and vice versa).

Passwords

Clearly, a password authenticating a user or a terminal would not remain secure indefinitely. In fact, in an environment of potential wiretapping or radiative pickup, a password might be compromised by a single use. Employing lists of randomly selected passwords in an "one-time-use" manner where a new word is taken from the list each time authentication is needed has been suggested as a countermeasure under such circumstances.[9] One copy of such a list would be stored in the processor, the other maintained in the terminal or carried by the user. After signing in, the user takes the next word on the list, transmits it to the processor and then crosses it off. The processor compares the received password with the next word in its own list and permits access only when the two agree. Such password lists could be stored in the terminal on punched paper tape, generated internally by special circuits, or printed on a strip of paper. The latter could be kept in a secure housing with only a single password visible. A special key lock would be used to advance the list. Since this method of password storage precludes automatic reading, the password must be entered using an appropriate input mechanism.

The protection provided by use of once-only passwords during sign-in procedures only is not adequate against more sophisticated "between lines" entry by an infiltrator who has attached a terminal to the legitimate user's line. Here the infiltrator can use his terminal to enter the system between communications from the legitimate user. In this situation the use of once-only passwords must be extended to each message generated by the user. Automatic generation and inclusion of authenticating passwords by the terminal would now be essential for smoothness of operation; and lists in the processor may have to be replaced by program or hardware implemented password generators.

Privacy Transformations

The identification procedures discussed above do not provide protection against passive threats through wiretapping, or against sophisticated "piggy-back" entry into the communication link. An infiltrator using the latter technique would simply intercept messages--password and all--and alter these or insert his own (e.g., after sending the user an error indication). Authentication of each message, however, will only prevent a "piggy-back" infiltrator from using the system after cancelling the sign-off statement of the legitimate user.

Although it may be conceivable that directly wired connections could be secured against wiretapping, it would be nearly impossible to secure common-carrier circuits. Therefore, the use of privacy transformations may be the only effective countermeasure against wiretapping and "piggy-back" entry, as they are designed to render encoded messages unintelligible to all but holders of the correct key. Discovering the key, therefore, is essential for an infiltrator. The effort required to do this by analyzing intercepted encoded messages (rather than by trying to steal or buy the key) is the "work factor" of a privacy transformation. It depends greatly on the type of privacy transformations used, as well as on the knowledge and ingenuity of the infiltrator.

The type of privacy transformation suitable for a particular communication network and terminals depends on the electrical nature of the communication links, restrictions on the character set, structure and vocabulary of the query language and data, and on the engineering aspects and cost of the terminals. For example, noisy communication links may rule out using "auto key" transformations[11] (i.e., those where the message itself is the key for encoding); and highly structured query languages may preclude

direct substitution schemes, since their statistics would not be obscured by substitution and would permit easy discovery of the substitution rules. Effects of character-set restrictions on privacy transformations become evident where certain characters are reserved for control of the communication net, terminals, or the processor (e.g., "end of message," "carriage return," and other control characters). These characters should be sent in the clear and appear only in their proper locations in the message, hence imposing restrictions on the privacy transformation.

A number of other implications of using privacy transformations arise. For example, in the private mode of terminal operation the system may provide an end-to-end (terminal-to-processor) privacy transformation with a given work factor, independently of the user's privacy requirements. If this work factor is not acceptable to a user, he should be permitted to introduce a preliminary privacy transformation using his own key in order to increase the combined work factor. If this capability is to be provided at the terminal, there should be provisions for inserting additional privacy-transformation circuitry.

Another (possibly necessary) feature might allow the user to specify by appropriate statements whether privacy

transforms are to be used or not. This would be part of a general set of "privacy instructions" provided in the information-system operating programs. Each change from private to public mode, especially when initiated from the terminal, should be authenticated.

FILES

While the above privacy-protection techniques and access-control procedures for external terminals and the communication network may greatly reduce the threat of infiltration by those with no legitimate access, they do not protect information against 1) legitimate users attempting to browse in unauthorized files, 2) access by operating and maintenance personnel, or 3) physical acquisition of files by infiltrators.

A basic aspect of providing information privacy to files is the right of a user to total privacy of his files--even the system manager of the information center should not have access. Further, it should be possible to establish different levels of privacy in files. That is, it should be feasible to permit certain of a group of users to have access to all of the company's files, while allowing others limited access to only some of these files.

In this context certain standard file operations--such as file copying--would seem inappropriate, if permitted in an uncontrolled manner, since it would be easy to prepare a copy of a sensitive file and maintain it under one's own control for purposes other than authorized. Similarly, writing into files should be adequately controlled. For example, additions and deletions to certain files should be authorized only after proper authentication. It may even be desirable to mount some files on drives with physically disabled writing circuits.

Access Control

Control of access to the files would be based on maintaining a list of authorized users for each file, where identification and authentication of identity (at the simplest), is established by the initial sign-in procedure. If additional protection is desired for a particular file, either another sign-in password or a specific file-access password is requested to reauthenticate the user's identity. The file-access passwords may be maintained in a separate list for each authorized user, or in a single list. If the latter, the system would ask the user for a password in a specific location in the list (e.g., the tenth password). Although a single list requires

less storage and bookkeeping, it is inherently less secure. Protection of files is thus based on repeated use of the same requirements--identification and authentication--as for initial access to the system during sign-in. This protection may be inadequate, however, in systems where privacy transformations are not used in the communication net (i.e., "piggy-back" infiltration is still possible).

Physical Vulnerability

An additional threat arises from possible physical access to files. In particular, the usual practice of maintaining backup files (copies of critical files for recovery from drastic system failures) compounds this problem. Storage, transport, and preparation of these files all represent points of vulnerability for copying, theft, or an off-line print out. Clearly, possession of a reel of tape, for example, provides an interloper the opportunity to peruse the information at his leisure. Applicable countermeasures are careful storage and transport, maintaining the physical integrity of files throughout the system, and the use of privacy transformations.

Privacy Transformations

As at the terminals and in the communication network, privacy transformations could be used to protect files against failure of normal access control or physical protection procedures. However, the engineering of privacy transformations for files differs considerably:

- Both the activity level and record lengths are considerably greater in files.
- Many users, rather than one, may share a file.
- Errors in file operations are more amenable to detection and control than those in communication links, and the uncorrected error rates are lower;
- More processing capability is available for the files, hence more sophisticated privacy transformations can be used.
- Many of the files may be relatively permanent and sufficiently large, so that frequent changes of keys are impractical due to the large amount of processing required. Hence, transformations with higher work factors could be used and keys changed less frequently.

It follows that the type of privacy transformation adequate for user-processor communications may be entirely

unacceptable for the protection of files.

The choice of privacy transformations for an information file depends heavily on the amount of file activity in response to a typical information request, size and structure of the file (e.g., short records, many entry points), structure of the data within the file, and on the number of different users. Since each of these factors may differ from file to file, design of a privacy system must take into account the relevant parameters. For example, a continuous key for encoding an entire file may be impractical, as entry at intermediate points would be impossible. If a complex privacy transformation is desired additional parallel hardware may be required, since direct implementation by programming may unreasonably increase the processing time. In order to provide the necessary control and integrity of the transformation system, and to meet the processing time requirements, a simple, securely housed processor similar to a common input-output control unit might be used to implement the entire file control and privacy system.

THE PROCESSOR

The processor and its associated random-access storage units contain the basic monitor program, system programs

for various purposes, and programs and data of currently serviced users. The role of the monitor is to provide the main line of defense against infiltration attempts through the software system by maintaining absolute control over all basic system programs for input-output, file access, user scheduling, privacy protection, etc. It should also be able to do this under various contingencies such as system failures and recovery periods, debugging of system programs, during start-up or shut-down of parts of the system, etc. Clearly, the design of such a fail-safe monitor is a difficult problem. Peters[9] describes a number of principles for obtaining security through software.

Penetration attempts against the monitor or other system programs attempt to weaken its control, bypass application of various countermeasures, or deteriorate its fail-safe properties. Since it is unlikely that such penetrations could be successfully attempted from outside the processor facility, protection relies mainly on adequate physical and personnel integrity. A first step is to keep the monitor in a read-only store, which can be altered only physically, housed under lock and key. In fact, it would be desirable to embed the monitor into the

basic hardware logic of the processor, such that the processor can operate only under monitor control.

Software integrity could be maintained by frequent comparisons of the current systems programs with carefully checked masters, both periodically and after each modification. Personnel integrity must, of course, be maintained at a very high level, and could be buttressed by team operation (forming a conspiracy involving several persons should be harder than undermining the loyalty of a single operator or programmer).

Integrity management procedures must be augmented with measures for controlling the necessary accesses to the privacy-protection programs; or devices for insertion of passwords, keys, and authorization lists; or for maintenance. These may require the simultaneous identification and authentication of several of the information-center personnel (e.g., a system programmer and the center "privacy manager"), or the use of several combination locks for hardware-implemented privacy-protection devices.

Hierarchy of Privacy Protection

Privacy protection requires a hierarchy of system-operating and privacy-protection programs, with the primary system supervisor at the top. Under this structure, or

embedded in it, may exist a number of similar but independent hierarchies of individual users' privacy-protection programs. It is neither necessary nor desirable to permit someone who may be authorized to enter this hierarchy at a particular level to automatically enter any lower level. Access should be permitted only on the basis of an authenticated "need to know." For example, if privacy transformations are employed by a particular user, his privacy programs should be protected against access by any of the system management personnel.

Time-Sharing

Various modes of implementing time-sharing in the information system may affect the privacy of information in the processor. In particular, copying of residual information in the dynamic portions of the storage hierarchy during the following time-slice seems likely. Since erasing all affected storage areas after each time-slice could be excessively time consuming, a reasonable solution may be to "tag" pages of information and programs as "private," or to set aside certain areas of the core for private information and erase only those areas or pages after each time-slice. Only the private areas or pages would need to be erased as part of the swapping operation.[14]

Also important is the effect of privacy transformations on processing time. Sophisticated privacy transformations, for example, if applied by programmed algorithms, may require a significant fraction of each time-slice. It may be necessary, therefore, to use hardware implementation of privacy transformations by including these in the hardware versions of the monitor or through the use of a separate parallel processor for all access-control and privacy-transformation operations.

Hardware Failures

With respect to integrity management, there is one aspect of hardware integrity which is the responsibility of the original equipment manufacturer, viz., a hardware failure should not be catastrophic in the sense that it would permit uncontrolled or even limited access to any part of the system normally protected. Whether this entails making the critical parts of the hardware super-reliable and infallible, or whether the system can be designed for a fail-safe form of graceful degradation is an open question. It is important to assure the user that the hardware has this basic characteristic.

V. CONCLUDING REMARKS

It should be emphasized again that the threat to information privacy in time-shared systems is credible, and that only modest resources suffice to launch a low-level infiltration effort. It appears possible, however, to counter any threat without unreasonable expenditures, provided that the integrity and competence of key personnel of the information center is not compromised. Trustworthy and competent personnel will establish and maintain the integrity of the system hardware and software which in turn, permits use of other protective techniques. A concise assessment of the effectiveness of these techniques in countering a variety of threats is presented in Table II.

Privacy can be implemented in a number of ways ranging from system enforced "mandatory privacy," where available privacy techniques are automatically applied to all users with associated higher charges for processing time, to "optional privacy" where a user can specify what level of privacy he desires, when it should be applied, and perhaps implement it himself. The latter privacy regime appears more desirable, since the cost of privacy could be charged essentially to those users desiring protection. For

Table II.

SUMMARY OF COUNTERMEASURES TO THREAT TO INFORMATION PRIVACY

Threat \ Countermeasure	Access Control (passwords, authentication, authorization)	Processing Restrictions (storage, protect, privileged operations)	Privacy Transformations	Threat Monitoring (audits, logs)	Integrity Management (hardware, software, personnel)
Accidental: User error	Good protection, unless the error produces correct password	Reduce susceptibility	No protection if depend on password; otherwise, good protection	Identifies the "accident prone"; provides post facto knowledge of possible loss	Not applicable
System error	Good protection, unless bypassed due to error	Reduce susceptibility	Good protection in case of communication system switching errors	May help in diagnosis or provide post facto knowledge	Minimizes possibilities for accidents
Deliberate, passive: Electromagnetic pick-up	No protection	No protection	Reduces susceptibility; work factor determines the amount of protection	No protection	Reduces susceptibility
Wiretapping	No protection	No protection	Reduces susceptibility; work factor determines the amount of protection	No protection	If applied to communication circuits may reduce susceptibility
Waste Basket	Not applicable	Not applicable	Not applicable	Not applicable	Proper disposal procedures
Deliberate, active: "Browsing"	Good protection (may make masquerading necessary)	Reduces ease to obtain desired information	Good protection	Identifies unsuccessful attempts; may provide post facto knowledge or operate real-time alarms	Aides other countermeasures
"Masquerading"	Must know authenticating passwords (work factor to obtain these)	Reduces ease to obtain desired information	No protection if depends on password; otherwise, sufficient	Identifies unsuccessful attempts; may provide post facto knowledge or operate real-time alarms	Makes harder to obtain information for masquerading; since masquerading is deception, may inhibit browsers
"Between lines" entry	No protection unless used for every message	Limits the infiltrator to the same potential as the user whose line he shares	Good protection if privacy transformation changed in less time than required by work factor	Post facto analysis of activity may provide knowledge of possible loss	Communication network integrity helps
"Piggy-back" entry	No protection but reverse (processor-to-user) authentication may help	Limits the infiltrator to the same potential as the user whose line he shares	Good protection if privacy transformation changed in less time than required by work factor	Post facto analysis of activity may provide knowledge of possible loss	Communication network integrity helps
Entry by system personnel	May have to masquerade	Reduces ease of obtaining desired information	Work factor, unless depend on password and masquerading is successful	Post facto analysis of activity may provide knowledge of possible loss	Key to the entire privacy protection system
Entry via "trap doors"	No protection	Probably no protection	Work factor, unless access to keys obtained	Possible alarms, post facto analysis	Protection through initial verification and subsequent maintenance of hardware and software integrity
Core dumping to get residual information	No protection	Erase private core areas at swapping time	No protection unless encoded processing feasible	Possible alarms, post facto analysis	Not applicable
Physical acquisition of removable files	Not applicable	Not applicable	Work factor, unless access to keys obtained	Post facto knowledge form audits of personnel movements	Physical preventative measures and devices

example, if a user feels that simple passwords provide adequate protection, his privacy system can be implemented in that way, with probable savings in processing time and memory space.

Implementation of "optional" privacy could be based on a set of "privacy instructions" provided by the programming and query languages of the system. Their proper execution (and safeguarding of the associated password lists and keys) would be guaranteed by the system. The cost of their use would reflect itself in higher rates for computing time. For example, the AUTHENTICATE,K,M instruction requests the next password from list K which has been previously allocated to this user from the system's pool of password lists (or has been supplied by the user himself). The operating program now sends the corresponding message to the user who takes the correct password from his copy of the list K and sends it to the processor. If the comparison of the received password with that in the processor fails, the program transfers to location M where "WRITE LOG, A" instruction may be used to make a record of the failure to authenticate in audit log A. Further instructions could then be used to generate a real-time alarm at the processor site (or even

at the site of the user's terminal), terminate the program, etc. Other privacy instructions would permit application of privacy transformations, processing restrictions, etc.

The privacy protection provided by any system could, of course, be augmented by an individual user designing and implementing privacy protection schemes within his programs. For example, he may program his own authentication requests, augmenting the system-provided instructions for this purpose; and use his own schemes for applying privacy transformations to data in files or in the communication network. Since these additional protective schemes may be software implemented they would require considerable processing time, but would provide the desired extra level of security.

Further Work

This paper has explored a range of threats against information privacy and pointed out some of the systems implications of a set of feasible countermeasures. Considerable work is still needed to move from feasibility to practice, although several systems have already made concrete advances. Special attention must be devoted to establishing the economic and operational practicality of

privacy transformations: determining applicable classes of transformations and establishing their work factors; designing economical devices for encoding and decoding; considering the effects of query language structure on work factors of privacy transformation; and determining their effects on processing time and storage requirements.

Large information systems with files of sensitive information are already emerging. The computer community has a responsibility to its users to insure that systems not be designed without considering the possible threats to privacy and providing for adequate countermeasures. To insure a proper economic balance between possible solutions and requirements, users must become aware of these considerations and be able to assign values to information entrusted to a system. This has been done in the past (e.g., industrial plant guards, "company confidential" documents, etc.), but technology has subtly changed accessibility. The same technology can provide protection, but we must know what level of protection is required by the user.

REFERENCES

1. "The computer and invasion of privacy," Hearings Before a Subcommittee on Government Operations, House of Representatives, July 26-28, 1966, U.S. Government Printing Office, Washington, D.C., 1966.

2. Baran, P., "Communications, computers and people," AFIPS Conference Proceedings, Fall Joint Computer Conference 1965, Part 2, Vol. 27.

3. Rothman, S., "Centralized government information systems and privacy," Report for the President's Crime Commission, September 22, 1966.

4. David, E. E., and R. M. Fano, "Some thoughts about the social implications of accessible computing," AFIPS Conference Proceedings, Fall Joint Computer Conference 1965, Part 1, Vol. 27, pp. 243-251.

5. Pursglove, S. D., "The eavesdroppers: 'fallout' from R&D," Electronic Design, Vol. 14, No. 15, June 21, 1966, pp. 35-43.

6. Dash, S., R. F. Schwartz, and R. E. Knowlton, The Eavesdroppers (Rutgers University Press, New Brunswick, New Jersey, 1959).

7. Ware, W. H., "Security and privacy: similarities and differences"; presented at SJCC 67, Atlantic City, New Jersey.

8. Daley, R. C., and P. G. Neumann, "A general-purpose file system for secondary storage," AFIPS Conference Proceedings, Fall Joint Computer Conference 1965, Part 1, Vol. 27, p. 213.

9. Peters, B., "Security considerations in a multi-programmed computer system"; presented at SJCC 67, Atlantic City, New Jersey.

10. Baran, P., On Distributed Communications: IX. Security, Secrecy and Tamper-Free Considerations, RM-3765-PR (The RAND Corporation, Santa Monica, California, August 1964).

11. Shannon, C. E., "Communications theory of secrecy systems," Bell System Technical Journal, Vol. 28, No. 4, October 1949, pp. 656-715.

12. Dennis, R. L., "Security in computer environment," SP 2440/000/01, System Development Corporation, August 18, 1966.

13. Dantine, D. J., "Communications needs of the user for management information systems," AFIPS Conference Proceedings, Fall Joint Computer Conference 1966, Vol. 29, pp. 403-411.

14. Patrick, R. L., Private communication.