Department of Computer Science Technical Reports

Department of Computer Science

1980

Research in Secure Computing (Final Report)

Dorothy E. Denning

Peter J. Denning

Report Number:
80-349

RESEARCH IN SECURE COMPUTING

Final Report

NSF Grant No. MCS77-04835

November 1977 - August 1980

Dorothy E. Denning
Peter J. Denning
Principal Investigators

Computer Science Department
Purdue University
West Lafayette, Indiana 47907

CSD TR 349

August 1980

RESEARCH IN SECURE COMPUTING

Final Report
on
NSF Grant MCS77-04835

Introduction

This report summarizes our research and other activities during the support period. The objective of the project was to study several aspects of secure computing. Our results fall into four areas:

1. Statistical database security
2. Application of cryptography
3. Fault tolerant computing
4. Miscellaneous

The following four sections describe our activities in each of these areas. This is followed by a section listing the personnel associated with the project and a section listing the publications and reports of the project.

Statistical Database Security

A statistical database contains confidential records about individuals. A set of query programs responds to requests for raw statistics about groups of records having common characteristics. The raw statistics typically include counts, sums, means, and medians of data elements. An example of a query is "What is the mean salary of all female doctors?".

The problem is that a enquirer can correlate the responses to a series of queries to deduce confidential data about an individual, thereby compromising the individual's privacy. The enquirer thus can use inference to bypass the security policies of the system.

Most of the controls traditionally used (e.g., by Census bureaus) to protect confidential data from statistical disclosure apply to off-line, static, databases. They are not applicable to on-line, dynamic database systems which provide access to information about arbitrary subgroups of individuals. The research problem is to find controls suitable for these systems.

Our early research in this area with Mayer Schwartz was directed primarily toward understanding the nature of the problem [1,2,3]. It was well-known that compromise was easy to achieve if the database responded to queries for statistics about small (or large) subgroups. However. even if the database refuses to answer such queries, we discovered that compromise was often easy using a simple but powerful snooping tool called the "tracker" [2].

Initially, we did not realize the extent to which trackers posed a serious threat, for we had believed that finding a tracker could be computationally infeasible. However, in collaboration with Jan Schlörer at the Universität Ulm, W. Germany, D. Denning showed that trackers were usually trivial to find using a binary search strategy [4],

D. Denning studied proposed controls for on-line database systems as well as controls traditionally used by Census bureaus and other agencies to protect published statistics. The results were discouraging: most controls were either easy to circumvent or impractical to implement in general purpose database systems, She presented these findings at the 12th Annual Symposium on the Interface of Computer Science and Statistics [5].

At this point D. Denning redirected her research, looking for a practical yet effective control. She proposed a new control called Random Sample Queries (RSQ), which meets this objective [6].

The common feature of all attacks is that the user can control which set of records is queried. RSQs deal directly with the basic principle of compromise by making it impossible for a questioner to precisely control the formation of query sets.

RSQs uses sampling theory to control disclosure. However, unlike traditional sampling controls, RSQs uses a different sample to compute each statistic. In preliminary experiments, the control was found to introduce enough uncertainty that users cannot get reliable estimates of confidential data in individual records but can get meaningful statistics for groups of records.

The scheme works as follows. As each record satisfying a given characteristic C is located, the system determines whether to keep this record in the sampled query set. This determination should ideally be pseudo random so that the same sampled query set is used each time C is presented. Each queried record is selected independently with a given, fixed probability. Statistics are then computed from the sampled query set.

A broad goal of D. Denning's research is to develop a theory of control powerful enough to describe a wide range of disclosure techniques applicable to database systems. This will provide a framework for comparing the advantage, risks, and costs of various control. She has applied to NSF for continued funding of her research in this area. Matt Bishop will be working with her on this project if it is funded.

D. Denning has also colloborated with Tore Dalenius of Brown University and the University of Stockholm. They have proposed a hybrid method for releasing statistics derived from population surveys [7]. Traditionally, these statistics are released in one of two formats: microstatistics (individual data records with identifying information removed) or macrostatistics (summary statistics in the form of tables containing counts or aggregates). Microstatistics have the advantage of giving the statistician the freedom to compute any statistic desired from the raw data, but the disadvantage of sometimes inadequate protection. Macrostatics have the advantage of greater protection, but the disadvantage of only a limited subset of all possible statistics. Hybrid statistics (low-ordered moments of the data) offer some of the advantages of both formats.

Applications of Cryptography

D. Denning has investigated the application of cryptography to the design of secure computer systems. She proposed a method for implementing secure personal computing in a insecure network [8]. The method employs a public-key encryption device and hardware keys. A user can safely store confidential files in the central facility or transmit confidential data to other users on the network without relying on the security of the central facility on the communication links. Each user is responsible for his own security.

In order for two users to communicate securely in a computer network which uses single-key encryption (such as the DES), a method is needed whereby two users can exchange a secret communication key. D. Denning and Giovanni Maria Sacco, a Ph.D. student at Purdue, have studied the problem of distributing secret communication keys in such systems. They show that proposed key distribution protocols based on handshakes are inadequate, and propose a solution based on time-stamps [9].

In collaboration with Fred Schneider of Cornell University, D. Denning has studied methods for generating and distributing shared group encryption keys. They propose several methods which permit groups of users to securely broadcast and share confidential information in a computer network [10]. One method can be used to implement a master key, which is capable of unlocking several group keys. These results were presented at the 1980 Symposium in Security and Privacy.

Fault Tolerant Computing

In earlier research P. Denning has shown how small domains of protection can lead to an operating system which is highly reliable because it limits error propogation and hastens error detection. Working under P. Denning's supervision, T. Don Dennis completed his Ph.D. research, showing the feasibility of supporting small protection domains with a tagged capability architecutre [11]. His thesis examines the feasibility of using a tagged memory and a stack processor to implement a capability based computer. The hypothesis put forth here is that these two architectural features reduce the cost of the capability mechanism and result in a simpler implementation. The thesis begins with a review of memory protection and protection systems as a basic need of modern computer systems. A brief historical survey is presented which begins with Dennis and Van Horn's paper on the semantics of multi-programmed computations and ends with a review of the major capability machines which have since been built.

The thesis then introduces the memory and processor organization proposed for this design and compares this organization with that used in previous architectures. This discussion shows that a tagged-memory organization reduces the number of segments used by a process by allowing segments to contain both pointer and data information. This reduces memory management overhead and allows a simpler representation of objects. Moreover, the tagged memory simplifies the mechanisms used to change domains, pass parameters, and address information in primary memory. The stack processor is shown to further reduce the cost of the capability mechanism by providing an inexpensive way to allocate procedure activation records, handle domain changes, and address objects on the stack.

The thesis next presents the capability mechanism for the proposed design and discusses the process of mapping capability based virtual addresses into absolute primary memory addresses.

It shows how the abstract type concept is directly supported by the capability mechanism.

A discussion of the hardware facilities needed to support the design follows. This discussion presents a detailed view of the registers of the central processor, the organization of the process stack, the instruction set, and a possible firmware organization which could be used to implement the design.

The proposed architecture is then compared with two conventional machines and shown to be more efficient in representing programs. Moreover, evidence is presented which suggests that the performance of the proposed machine would be competitive with current state-of-art machines. Finally, two examples are presented which use the abstract type, process synchronization, and process communication facilities provided by the design.

The major contribution of this thesis is its thorough examination of the tagged memory approach to capability addressing. The research found that a tagged-memory capability machine is powerful enough to implement an operating system and the resulting operating system has a complexity about equal to that of the older generation of simple systems with little memory protection. The thesis also illustrates that advanced programming concepts can be implemented quite simply if the proper hardware support is provided.

Finally, the thesis shows that the process control mechanism can be designed to remove an important source of memory contention from the mutaul exclusion mechanism. This is seen as a partial solution to the memory contention problem found on many multi-processor systems. A paper has been submitted for publication [12].

Miscellaneous

We wrote a tutorial paper for Computing Surveys describing four types of security controls: access controls, flow control, inference controls, and cryptographic controls [13]. We described the general nature of controls of each type, the kinds of problems they can and cannot solve, and their inherent limitations and weakenesses. This paper was republished by Auerback for data processing management.

D. Denning wrote a short note on information flow into arrays, showing that flows caused by execution of array assignments are a special case of a more general pattern of implicit flow [14].

D. Denning presented a paper at COMPSAC 80, describing a method for maintaining routing records in automated record keeping systems [15]. The Privacy Act of 1974 mandates the use of routing data in all record keeping systems maintained by Federal Agencies. Routing data provides an account of the recipients of data stored in a record. The proposed technique uses time stamps to link records.

Personnel

Besides the principal investigators (D. Denning and P. Denning), these students contributed to the project:

1. T. Don Dennis, who completed his Ph.D. thesis in May 1980 on "A Capability Architecture" under P. Denning's supervision.

2. Kye Hedlund, Ph.D. student, who is working under D. Denning on computer architectures which exploit VLSI technology. Although the primary focus of his research is efficient computation, protection factors are also being considered.

3. Gerald Kreissig,,Ph.D. student, who is working under D. Denning. He is interested in protection models for access control, and presented a paper at the 1980 Symposium on Security and Privacy [16].

4. Matt Bishop, Ph.D. student, who is working under D. Denning on statistical database security as described earlier.

5. Giovanni Maria Sacco, Ph.D. student, who collaborated with D. Denning on key distribution methods as described earlier.

Publications and Reports

1. P. J. Denning and D. E. Denning, "Research in Secure Computing: A Program Report," CSD TR 251, Comp. Sci. Dept., Purdue Univ., Nov. 1977.

2. D. E. Denning, P. J. Denning, and M. D. Schwartz, "The Tracker: A Threat to Statistical Database Security," ACM Trans. Database Syst. 4, 1 (March 1979), 476-482.

3. M. D. Schwartz, D. E. Denning, and P. J. Denning, "Linear Queries in Databases," ACM Trans. on Database Syst. 4, 2 (June 1979) 476-482.

4. D. E. Denning and J. Schlörer, "A Fast Procedure for Finding a Tracker in a Statistical Database," ACM Trans. Database Syst. 5, 1 (Mar. 1980), 88-102.

5. D. E. Denning, "Complexity Results Relating to Statistical Confidentiality," Proc. Computer Science and Statistics: 12th Annual Symp. on the Interface, Waterloo, May 1979, 252-256.

6. D. E. Denning, "Secure Statistical Databases with Random Sample Queries," ACM Trans. Database Syst., (Sept. 1980).

7. T. Dalenius and D. E. Denning, "A Hybrid Scheme for Release of Statistics," Statistical Reporter (to appear).

8. D. E. Denning, "Secure Personal Computing in an Insecure Network," Comm. ACM 22, 8 (Aug. 1979), 476-482.

9. G. M. Sacco and D. E. Denning, "Handshakes are Shaky," Submitted for publication.

10. D. E. Denning and F. B. Schneider, "The Master Key Problem," Proc. 1980 Symp. on Security and Privacy, Oakland, April 1980.

11. T. D. Dennis, "A Capability Architecture," Ph.D. Thesis, Purdue Univ., May 1980.

12. P. J. Denning, T. Don Dennis, and J. Brumfield, "Low Contention Semaphores and ready lists," Technical Report TR-322 (Feb. 1980).

13. D. E. Denning and P. J. Denning, "Data Security," Computing Surveys 11, 3 (Sept. 1979) 227-249.

14. D. E. Denning, "Embellishments to the note on Information Flow into Arrays," ACM Sigsoft, Software Eng. Notes 5, 2 (April 1980) 15-16.

15. D. E. Denning, "A Method for Maintaining Routing Data in Automated record Keepting Systems," Proc. COMPSAC (1978), 215-219.

16. G. Kreissig, "A Model to Describe Protection Systems," Proc. 1980 Symp. on Security and Privacy, Oakland, April 1980.