Preliminary Thought on Degrees of Security in Multiprocessor Systems *

C. A. Ellis
G. J. Nutt

CU-CS-036-74

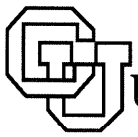

University of Colorado at Boulder
DEPARTMENT OF COMPUTER SCIENCE

* This work supported by National Science Foundation Grant #GJ-660

Preliminary Thoughts on Degrees of Security in Multiprocessor Systems *

by

C. A. Ellis and G. J. Nutt
Department of Computer Science
University of Colorado
Boulder, Colorado 80302

Report # CU-CS-036-74 Jan. 1974

* This work supported by National Science Foundation Grant #GJ-660.

INTRODUCTION

The existence of simultaneously running processes in a multiprocessor system causes certain issues in operating systems design to be critical: deadlock becomes more likely as the number of processes grows [2]; the scheduling of processes to processors becomes less intuitive (i.e. more complex) [5]; and sharing of resources under system authorization becomes more difficult to "guarantee" [4]. In this paper we are interested in quantifying the amount of authorized and unauthorized resource sharing in a class of multiple process, multiple resource systems which employ a particular protection mechanism. The exact mechanism defined here is patterned after the approach used in the Multi Associative Processor (MAP) design, described in detail elsewhere [1], [3], [7].

A security policy for an operating system is the particular set of rules under which all processes must exist. This policy is set by the system designer as he implements the software for his operating system. The implementation of the security policy can be substantially easier if a suitable protection mechanism is available in the hardware (or firmware) of the machine. An example of a mechanism used to implement a variety of policies can be seen in paging system design. A name space to memory space mechanism (such as an associative memory) is supplied, usually in the hardware of the machine. The existence of the virtual memory mapping mechanism allows a variety of paging policies to be implemented, e.g. first-in-first-out, random, etc. Examples of hardware protection mechanisms incorporated into machines for particular security policies are the key/lock memory protection on the SYSTEM/360, and the mechanism for ring structured security policy in MULTICS [8].

The overwhelming idea behind many security policies is that of allowing no unauthorized accesses of resources to take place in the system. It is sometimes the case that this "perfectionist's approach" is arbitrarily expensive to implement. Here we are willing to allow unauthorized accesses to take place, provided that a measure of the degree of unauthorized access is known a priori. A typical circumstance in which this situation might arise is given by the following example:

> Suppose that a certain system has a 2-bit conventional key/lock memory write protection mechanism. We shall assume, for simplicity, that there can exist only 4 key/lock pairs, 00-00, 01-01, 10-10, and 11-11 (i.e. we ignore the existence of a "skeleton key" for a supervisory process). Now, if the system is to support 5 processes, one of the following alternatives must be followed:
>
> a) Add another bit to the keys and locks.
>
> b) Assign the same key/lock pair to distinct processes.
>
> The first alternative is generally economically infeasible and the second places the burden of security on the processes using the shared key/lock combination. This latter approach amounts to the existence of potential unauthorized access.

In the remainder of this paper, we shall introduce a basic protection mechanism which is a generalization of the mechanism used in the MAP system [1, 3, 7]. Next, the idea of degrees of security is formally defined along with some measures of unauthorized access. It is then shown that a given class of security policies minimizes the number of unauthorized accesses that can take place in the system.

A BASIC PROTECTION MECHANISM

Suppose there exists a set of elements, U, in a computer system, the members of which participate in various forms of interaccess; (some of the members may be "processes in a domain", and others may be "resources"). An individual element can be classified as a subject if it instigates an access, or an object if it receives the activity associated with the access, e.g. a memory module, a file, another process, etc. The minimal form of access that we shall consider is simply message passing or process intercommunication; in a process-to-process conversation, the processes alternately take on the role of subject and object. Denote the set of all possible subjects by S and the set of all possible objects as θ; note that $U = S \cup \theta$ and that $S \cap \theta$ is not necessarily empty.

For all $s_i \in S$, associate an access key, $A(s_i)$, corresponding to a binary register. (This access key would be of fixed size n and be part of the process descriptor in an implementation). Similarly, with each $o_j \in \theta$, associate an access lock, $A(o_j)$, which corresponds to an n bit binary register. If $c \in S \cap \theta$, the access key and lock are identical. We shall refer to the k^{th} bit in A(c) as $A_k(c)$.

In general, it is neither necessary nor desirable for all $s_i \in S$ to be able to access (or even communicate with) each $o_j \in \theta$. The protection mechanism, thus, will allow only limited access as determined by the security policy, and is based on the:

Rule of Access: Let $s_i \in S$ and $o_j \in \theta$.

If $\sum_{k=1}^{n} A_k(a_i) \wedge A_k(p_j) \geq 1$.

(for n = register size) then s_i has access to o_j.

If $o_j \in \theta \cap S$, then the relationship is symmetric and "access" can be thought of as "communicate with".

A simple application of this mechanism to obtain a well-known security policy is illustrated by the following example:

Suppose that a certain system supports five processes simultaneously in which one of these processes is considered to be a supervisor and the other four are subordinate (user) processes with no ability to directly intercommunicate. By assigning the values shown below to 4-bit access registers, one can implement a policy in which no unauthorized communication can take place between user processes.

A (supervisor) = 1111

A (user #1) = 0001

A (user #2) = 0010

A (user #3) = 0100

A (user #4) = 1000

Private resources of user #i will have an access lock value of 2^{i-1}. (Note that no user can have resources which the supervisor will not be able to access, and with a 4-bit register, the supervisor can have no resource which is inaccessible to all users.) If a resource is to be shared by user #i and user #j, then the resource will have an access lock value of $2^{i-1} + 2^{j-1}$.

DEGREES OF SECURITY

The previous example presupposed that there was an adequate number of bits available in the access registers to represent any security policy with no unauthorized accesses. In a practical situation, the length of the access lock and key registers is fixed and the number of active elements may be large enough to prevent maximum security. Thus we turn our attention to the situation in which the tradeoff consideration is between the number of subjects in the system and the integrity of the security policy.

Definitions

1. Let x_{ij} and y_{ij} be boolean variables such that $x_{ij} = 1$ ($y_{ij} = 1$) implies that s_i has authorized (unauthorized) access to o_j; $x_{ij} = 0 (y_{ij} = 0)$ otherwise. $x_{ij} \wedge y_{ij} = 0$ necessarily for all $i (1 \leq i \leq |S|)$ and for all $j (1 \leq j \leq |\theta|)$ where $|Z|$ denotes the cardinality of the set Z.

2. Let $x_j (y_j)$ be the number of $s_i \in S$ which have authorized (unauthorized) access to $o_j \in \theta$, i.e.

$$a)\ x_j = \sum_{s_i \in S} x_{ij}, \qquad b)\ y_j = \sum_{s_i \in S} y_{ij}$$

3. Let $\bar{x}(\bar{y})$ be the average, over θ, of the number of subjects which have authorized (unauthorized) access to any particular $o_j \in \theta$, i.e.

$$a)\ \bar{x} = \frac{\sum_{o_j \in \theta} x_j}{|\theta|}, \qquad b)\ \bar{y} = \frac{\sum_{o_j \in \theta} y_j}{|\theta|}.$$

4. Let $\overset{\vee}{x}(\overset{\vee}{y})$ be the minimum, over θ of the number of subjects which have authorized (unauthorized) access to any particular $o_j \in \theta$, i.e.

$$\overset{\vee}{x} = \min_{o_j \in \theta} (x_j), \quad \overset{\vee}{y} = \min_{o_j \in \theta} (y_j)$$

5. Let $\hat{x}(\hat{y})$ be the maximum, over P, of the number of subjects which have authorized (unauthorized) access to any particular $o_j \in \theta$, i.e.

$$a)\ \hat{x} = \max_{o_j \in \theta} (x_j), \qquad b)\ \hat{y} = \max_{o_j \in \theta} (y_j)$$

6. The degree of security of a system is then

$$\bar{\delta}_{abs} = (1 + \bar{y})^{-1}$$

This definition yields a value of unity when the system allows no unauthorized accesses, and less than unity otherwise. It is absolute in the sense that it does not vary according to the size of the system, but only according to the average number of unauthorized accesses

7. The relative degree of security of a system is

$$\bar{\delta}_{rel} = \frac{|S| - \bar{x} - \bar{y}}{|S| - \bar{x}}$$

This definition yields a value of unity when the system allows no unauthorized access, and a value of zero when the maximum possible number of unauthorized accesses are allowed. This value is therefore relative to the size and structure of the system being analyzed.

8. The minimum (maximum) degree of security of a system, $\check{\delta}_{abs}$ ($\hat{\delta}_{abs}$) is

a) $\check{\delta}_{abs} = (1 + \hat{y})^{-1}$

b) $\hat{\delta}_{abs} = (1 + \check{y})^{-1}$

Recalling the example from the previous section, we consider the frequently occurring case of m independent simultaneously active processes (not including any supervisory process) each of which needs authorized access to one (its own) memory module. Furthermore assume that k (the number of bits in an access register), is less than m. Security of memory access in this case can be formulated as a problem in which the processes form a set of m (non-object) subjects having keys of length k and their memories form a set of objects having lock registers of length k. By the Rule of Access, $s_i \in S$ can attain access to $o_j \in \theta$ if and only if $A_h(a_i) = A_h(p_j) > 0$ for some $1 \leq h \leq k$. This example defines a class of systems with protection mechanisms which we call M_{ind}. This class is characterized by $|S| = |\theta| = m > k$, $S \cap \theta = \emptyset$, and a unique one-to-one numbering of S and θ such that $x_{ij} = \delta_{ij}$, (Kronecker Delta means $x_{ij} = 1$ if $i = j$, $x_{ij} = 0$ otherwise) for $1 \leq i \leq k$ $1 \leq j \leq n$. A method of assigning locks to each $o_j \in \theta$ and keys to each $s_i \in S$ specifies the security policy. The fact that $m > k$ means that it is impossible to assign

a unique bit or non-empty set of bits to each process and the corresponding resource to which it has authorized access. Thus any security policy will allow some unauthorized accesses in this case. One class of security policies assigns $\left[\frac{m}{k}\right]$ processes and $\left[\frac{m}{k}\right]$ resources to each of the possible keys containing a single 1 except for $\left(\text{remainder } \left[\frac{m}{k}\right]\right)$ of those keys which are assigned $\left(\left[\frac{m}{k}\right] + 1\right)$ processes and resources where [Z] denotes the greatest integer less than or equal to Z. Thus all processes and resources are assigned a key of the form 00...010...00. We call this class P_{opt} because the following theorem can be proved.

Theorem:

For any protection mechanism in the Class M_{ind}, strategies in the class P_{opt} are optimal in the sense that they maximize the minimum degree of security $\overset{\vee}{\delta}_{abs}$.

Case of k = 2

a. $k|m$, then obviously half of the key and locks are 01, half are 10 one of those keys is authorized, so $\left[\frac{m}{k}\right] - 1$ unauthorized accesses are possible = $\hat{n}$. This is minimal $\hat{n}$ value because any strategy using 00 allows no access, any strategy using 11 allows subjects to access that o_1, so $\hat{n} = m - 1 > \frac{m}{2} - 1 = \left[\frac{m}{2}\right] - 1$. q.e.d.

b. $k|m$, then $\left[\frac{m}{k}\right] + 1$ of the objects and subjects 01 and $\left[\frac{m}{k}\right]$ use 01 (or vice-versa, so there exist $\frac{m}{k}$ and $\frac{m}{k} - 1$ unauthorized accessors in the two cases, i.e. $y_{ij} = \left[\frac{m}{k}\right]$ or else $y_{ij} = \left[\frac{m}{k}\right] - 1$ for all i,j; thus $\hat{y} = \left[\frac{m}{k}\right]$. That this $\hat{y}$ value is minimal can be seen using the same argument as a above.

Case of k > 2

Assume induction hypothesis for all k' < k:

if $m|k'$, then optimal is m/k -1

if $m|k'$, then " " [m/k]

a. case $k|m$

Suppose $\exists$ better schedule*, i.e. $\hat{y}^* < m/k - 1$ and this schedule is not solely of form 0...010...0, then $\exists$ o_i such that its key has at least 2 ones in it

(0...010...010...0). Call the set of all s_j with non-empty intersection with $A(o_i)$ the set C. Then by assumption, $|C| \leq \lceil m/k \rceil - 1$ (because $|C| \leq \hat{y}^* + 1$) so $y_i \leq m/k - 2$. The remaining o_j have bits disjoint with the key of o_i, so they must share less than or equal to k - 2 bits. We see that this yields a system of $|S| - |C| \geq \frac{k-1}{k} m + 1$ resources and k - 2 bits. By the induction hypothesis, the best we can do is $\hat{y}^* \geq \left\lceil \frac{\frac{k-1}{k} m + 1}{k-2} \right\rceil - 1 \geq$

$$\left\lceil \frac{m}{k} + \frac{1}{k-2} \right\rceil - 1 \geq \frac{m}{k} - 1. \Rightarrow\Leftarrow$$

b. case k | m

Suppose $\hat{y}^* < \lceil m/k \rceil$, where again $\exists o_i$ such that its key has two or more ones. The set C has $|C| \leq \lceil m/k \rceil$, so $y_i < \lceil m/k \rceil$. An analogous argument to a above shows that a system of $|S| - |C| \geq m - \left\lceil \frac{m}{k} \right\rceil \geq \left\lceil \frac{k-1}{k} m \right\rceil$ must share less than or equal to k - 2 bits. By the induction hypothesis, $\hat{y}^* \geq \left\lceil \frac{\frac{k-1}{k} m}{k-2} \right\rceil \geq \left\lceil \frac{m}{k} \right\rceil$ if the numerator is not divisible by k-2. $\Rightarrow\Leftarrow$

If $(k-2) \mid (m - \lceil m/k \rceil)$, then m - m/k approxiamtion yields the same quotient minus 1, so

$$\frac{m - \lceil m/k \rceil}{k-2} = \left\lceil \frac{m - m/k}{k-2} \right\rceil + 1 \text{ and application}$$

of the induction hypothesis for k' = k - 2 implies

$$\hat{y}^* \geq \frac{m - \lceil m/k \rceil}{k-2} - 1 = \left(\left\lceil \frac{m - m/k}{k-2} \right\rceil + 1 \right) - 1 = \left\lceil \frac{\frac{k-1}{k} m}{k-2} \right\rceil \geq \left\lceil \frac{m}{k} \right\rceil$$

q.e.d.

SUMMARY

In this paper we have discussed a protection mechanism and its application to implement a set of security policies. A measure of the degree of security for a given policy is defined, and it is shown that a class of security policies exist which maximizes the minimum degree of security. The ideas within this paper appear very promising, and there is a lot of work to be done here. Theorems must be obtained for $\overline{\delta}$ which is the most interesting security measure, and other definitions appear useful. For example, try the following:

Let σ_y be the variance of a security system where $\sigma_y = \sum \frac{(y_j - \bar{y})^2}{|M|}$

Let ρ_y be the deviance of a security system where $\rho_y = \hat{y} - \check{y}$. We then can say that the security range of a system is from $\hat{y}$ to $\check{y}$.

Thus all results reported here should be viewed as preliminary definitions and assertions prefacing an in depth study which the authors are interested in pursuing. Current research is concerned with further definitions to capture the essence of protection structures, and extension of the degrees of security results to all functions of the form $\sum_{k=1}^{n} A_k(a_i) \wedge A_k(p_j) > \ell$ for any $0 \leq \ell < n$. Indeed, it would be useful to consider and categorize the set of all functions of A_k which could be used as basic protection mechanisms.

REFERENCES

1. Arnold, R. D. "Multi Associative Processor System Architecture" University of Colorado, Department of Computer Science, MAP Memo No. 3, Dec. 1973.

2. Ellis, C. A. "On the Probability of Deadlock in Computer Systems" Proceedings of Fourth ACM Symposium on Operating Systems Principles, Oct. 1973.

3. Ellis, C. A. "Associative Operating Systems Study" University of Colorado, Department of Computer Science, MAP Memo No. 2, Dec. 1973.

4. Graham, G. S. and Denning, P. J. "Protection-Principles and Practice" AFIPS SJCC Conference Proceedings, vol. 40, (May 1972).

5. Graham, R. L. "Bounds on Multiprocessing Anomalies and Packing Algorithms" AFIPS SJCC Conference Proceedings, vol. 40 (May 1972).

6. Jones, A. K. "Protection in Programmed Systems" Ph.D. Thesis, Carnegie-Mellon University, Department of Computer Science, 1972.

7. Nutt, G. J. "Multi Associative Processor Evaluation Study" University of Colorado, Department of Computer Science, MAP Memo No. 1, Nov. 1973.

8. Schroeder, M. D. and Saltzer, J.H. "A Hardware Architecture for Implementing Protection Rings" Communications of the ACM, 15,3 (March, 1972).

9. Molho, L. M. "Hardware Aspects of Secure Computing" AFIPS SJCC Conference Proceedings, vol. 36, (1970).