

THE CHINESE WALL SECURITY POLICY

Dr. David F.C. Brewer and Dr. Michael J. Nash

GAMMA SECURE SYSTEMS LIMITED
9 Glenhurst Close, Blackwater, Camberley, Surrey, GU17 9BQ, United Kingdom

ABSTRACT

Everyone who has seen the movie Wall Street will have seen a commercial security policy in action. The recent work of Clark and Wilson and the WIPICIS initiative (the Workshop on Integrity Policy for Computer Information Systems) has drawn attention to the existence of a wide range of commercial security policies which are both significantly different from each other and quite alien to current "military" thinking as implemented in products for the security market place.

This paper presents a basic mathematical theory which implements one such policy, the Chinese Wall, and shows that it cannot be correctly represented by a Bell-LaPadula model.

The Chinese Wall policy combines commercial discretion with legally enforceable mandatory controls. It is required in the operation of many financial services organizations and is, therefore, perhaps as significant to the financial world as Bell-LaPadula's policies are to the military.

INTRODUCTION

Until recently, military security policy thinking has dominated the direction of computer security research both in the US and the UK. Clark and Wilson's seminal paper [1] has, however, drawn attention to the fact that commercial security needs are just as important as those of the defence community and, through the WIPICIS initiative [2], that the problems of the commercial community are at least as diverse and relevant to the computer scientist.

There are many well defined commercial security policies covering all aspects of Clark and Wilson's model [3]. One of these, the *Chinese Wall* security policy is perhaps as significant to some parts of the commercial world as Bell and LaPadula's policies [4, 5] are to the military. It can be most easily visualized as the code of practice that must be followed by a market analyst working for a financial institution providing corporate business services. Such an analyst must uphold the confidentiality of information provided to him by his firm's clients; this means he cannot advise corporations where he has *insider knowledge* of the plans, status or standing of a competitor.

However, the analyst is free to advise corporations which are not in competition with each other, and also to draw on general market information. Many other instances of Chinese Walls are found in the financial world.

Unlike Bell and LaPadula, access to data is not constrained by attributes of the data in question but by what data the subject already holds access rights to. Essentially, datasets are grouped into "conflict of interest classes" and by mandatory ruling all subjects are allowed access to at most one dataset belonging to each such conflict of interest class; the actual choice of dataset is totally unrestrained provided that this mandatory rule is satisfied. We assert that such policies cannot be correctly modelled by Bell-LaPadula.

It should be noted that in the United Kingdom the Chinese Wall requirements of the UK Stock Exchange [6] have the authority of law [7] and thus represent a mandatory security policy whether implemented by manual or automated means.

Furthermore, correct implementation of this policy is important to English Financial Institutions since it provides a legitimate defence against certain penal classes of offence under their law.

CHINESE WALLS

A Model

It is useful to devise a model of the Chinese Wall policy, not only to facilitate sound reasoning of its properties but also to permit comparison with models of other policies.

Since, in particular, we wish to compare it with the Bell-LaPadula (BLP) model we will adopt the latter's concepts of *subjects*, *objects* and *security labels*. Our model is then developed by first defining what we mean by a *Chinese Wall* and then, by devising a set of rules such that no person (subject) can ever access data (objects) on the *wrong* side of that wall.

Database organization

We start by informally defining the concept of a Chinese Wall and, in particular, what is meant by being on the *wrong* side of it. A formal specification will be found in Annex A.

All corporate information is stored in a hierarchically arranged filing system such as that shown in figure 1. There are three levels of significance:

- a) at the lowest level, we consider individual items of information, each concerning a single corporation. In keeping with BLP, we will refer to the files in which such information is stored as *objects*;
- b) at the intermediate level, we group all objects which concern the same corporation together into what we will call a *company dataset*;
- c) at the highest level, we group together all company datasets whose corporations are in competition. We will refer to each such group as a *conflict of interest class*.

Associated with each object is the name of the company dataset to which it belongs and the name of the conflict of interest class to which that company dataset belongs. For example, the conflict of interest class names could be the business sector headings found in stock exchange listings (Petrochemical, Financial Services, etc.); the company dataset names could be the names of companies listed under those headings. For convenience let us call the r -th object o_r and refer to its company dataset as y_r and its conflict of interest class as x_r .

Thus, if our system maintained information say on Bank-A, Oil Company-A and Oil Company-B:

- a) all objects would belong to one of three company datasets (i.e. y_r would be "Bank-A", "Oil Company-A" or "Oil Company-B") and
- b) there would be two conflict of interest classes, one for banks (containing Bank-A's dataset) and one for petroleum companies (containing Oil Company-A's and Oil Company-B's datasets), i.e. x_r would be "banks" or "petroleum companies"

This structure is of great importance to the model and is represented in the annex as an axiom (A1).

Simple security

The basis of the Chinese Wall policy is that people are only allowed access to information which is not held to conflict with any other information that they already possess. As far as the computer system is concerned the only information *already possessed* by a user must be information that:

- a) is held on the computer, and
- b) that user has previously accessed.

Thus, in consideration of the Bank-A, Oil Company-A and Oil Company-B datasets mentioned previously, a *new* user may freely choose to access whatever datasets he likes; as far as the computer is concerned a new user does not possess *any* information and therefore no conflict can exist. Sometime later, however, such a conflict may exist.

Suppose our user accesses the Oil Company-A dataset first; we say that our user now *possesses* information concerning the Oil Company-A dataset. Sometime later he requests access to the Bank-A dataset. This is quite permissible since the Bank-A and Oil Company-A datasets belong to different conflict of interest classes and therefore no conflict exists. However, if he requests access to the Oil Company-B dataset the request must be denied since a conflict *does* exist between the requested dataset (Oil Company-B) and one already possessed (Oil Company-A).

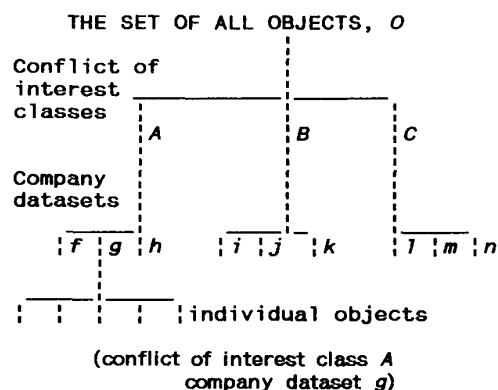


Figure 1 - The composition of objects

We note that it does not matter whether the Oil Company-A dataset was accessed before or after the Bank-A dataset. However, were Oil Company-B to be accessed before the request to access the Oil Company-A dataset, the restrictions would be quite different. In this case access to the Oil Company-A dataset would be denied and our user would possess {"Oil Company-B", "Bank-A"} (as opposed to the request to access the Oil Company-B dataset being denied and the user possessing {"Oil Company-A", "Bank-A"}).

What we have just described is a Chinese Wall. We note, in the first instance, that our user has complete freedom to access anything he cares to choose. Once that initial choice has been made, however, a *Chinese Wall* is created for that user around that dataset and we can think of "the *wrong* side of this Wall" as being any dataset within the same conflict of interest class as that dataset *within* the Wall. Nevertheless, the user still has freedom to access any other dataset which

is in a different conflict of interest class, but as soon as that choice is made, the Wall *changes shape* to include the new dataset. Thus we see that the Chinese Wall policy is a subtle combination of free choice and mandatory control.

More formally, we can introduce BLP's concept of a *subject* to represent a user, and indeed any program that might act on his behalf. Again, it is convenient to be able to identify any particular subject, let us call it s_u for the u -th subject. However, we conclude that it is necessary for the computer to remember who has accessed what. Annex A introduces a formalized device (D1), \mathbf{N} , to do this. Essentially \mathbf{N} is a matrix with a column for every object in the system and a row for every subject; any particular cell therefore corresponds to a particular (subject, object) pair. \mathbf{N} simply records who has accessed what, in other words who *possesses* what. For example if a subject, s_u , has accessed some object o_r then $\mathbf{N}(u, r)$ would be set true, otherwise it would be set false. Once some request by, say, s_u , to access some *new* object, say o_t , is granted then \mathbf{N} , of course, must be updated; in particular \mathbf{N} must be replaced by some new \mathbf{N}' in which $\mathbf{N}'(u, t)$ takes the value true.

Given this we can then define a security rule (axiom A2 in Annex A). This rule means that:

Access is only granted if the object requested:

- a) is in the *same company dataset* as an object already accessed by that subject, i.e. within the Wall, or
- b) belongs to an *entirely different conflict of interest class*.

(A2).

This rule is, however, insufficient in itself; it is necessary to ensure that initially, i.e. before any subject has requested access to any object, that \mathbf{N} has been correctly initialized (everywhere false). Moreover it is necessary to formally define the notion that, given such an initial state, then the first ever access will be granted. These two requirements are stated as axioms A3 and A4 (respectively) in Annex A.

This rule is analogous to the *simple* security rule defined in the BLP model and since these two rules will be compared with each other, we will call our rule (A2) the *simple* security rule as well. (Not surprisingly, we will define a $*$ -property rule as well, see later.)

Accessibility

The basic formulation of the Chinese Wall policy is expressed in Annex A as axioms A1-A4. Given these, we can prove some useful theorems about the policy. Annex A does this. The first two theorems merely confirm that our informal concepts of a Chinese Wall are indeed implied by these rules, in particular:

T1) **Once a subject has accessed an object the only other objects accessible by that subject lie within the same company dataset or within a different conflict of interest class.**

T2) **A subject can at most have access to one company dataset in each conflict of interest class.**

A third theorem concerns a most pragmatic consequence of a Chinese Wall policy, and that is *how many people do you need to operate one?*

T3) **If for some conflict of interest class X there are X_y company datasets then the minimum number of subjects which will allow every object to be accessed by at least one subject is X_x .**

This tells us that it is the conflict of interest class with the largest number of member company datasets (L , say) which will give us the answer. However, because of the free choice nature of the policy, should two or more users choose to access the same dataset then this number must be increased. Moreover, if we had an Automobile class with 5 motor companies ($L_y = 5$) and five users each with access to *different* motor companies, but *all* choose to access the *same* petroleum company, Oil Company-A, then who can access Oil Company-B?

Hence, in practice, it is not necessarily the largest conflict of interest class that gives us an accessibility problem. L_y is nevertheless the minimum number of subjects ever required and thus the minimum number of analysts the finance house must employ.

Sanitized information

It is usual in the application of the Chinese Wall policy for company datasets to contain sensitive data relating to individual corporations, the conflict of interest classes being in general business sectors. However it is considered important to be able to generally compare such information with that relating to other corporations. Clearly a problem arises if access to any corporation is prohibited by the simple security rule due to some previous access (a user can never compare Oil Company-A's data with Oil Company B's data, nor can he compare Bank-A's data with Oil Company-A's data if he has had previous access to Oil Company-B's data). This restriction can, however, be removed if we use a sanitized form of the information required.

Sanitization takes the form of disguising a corporation's information, in particular to prevent the discovery of that corporation's identity. Obviously, effective sanitization cannot work unless there is sufficient data to prevent backward inference of origin. Since it is not the objective of this paper to discuss the problem of inference we will assume

that sanitization is possible. In practice, this has been found to be the case [3].

It seems sensible, therefore, to regard all company datasets, bar one, as containing sensitive information belonging to some particular corporation. We reserve the remaining dataset, y_0 (say), for sanitized information relating to all corporations (D2 in Annex A).

It is unnecessary to restrict access to such sanitized information. This can be accomplished without need to rework any of the preceding theory by asserting that this distinguished company dataset, y_0 , is the sole member of some distinguished conflict of interest class, x_0 (A5). In other words, if an object bears the security label y_0 then it must also bear the label x_0 and vice versa. T2 tells us that all subjects can access this company dataset.

***-Property**

Suppose two subjects, User-A and User-B, have between them access to the three datasets, Oil Company-A, Oil Company-B and Bank-A, in particular User-A has access to Oil Company-A and Bank-A and User-B has access to Oil Company-B and Bank-A. If User-A reads information from Oil Company-A and writes it to Bank-A then User-B can read Oil Company-A information. This should not be permitted, however, because of the conflict of interest between Oil Company-A and Oil Company-B. Thus indirect violations of the Chinese Wall policy are possible.

We can prevent such violations by insisting that:

Write access is only permitted if

- a) **access is permitted by the simple security rule, and**
- b) **no object can be read which is in a different company dataset to the one for which write access is requested and contains unsanitized information.**

(A6).

Given this rule we can prove that:

- T4) **The flow of unsanitized information is confined to its own company dataset; sanitized information may however flow freely throughout the system.**

The rule is analogous to the **-property* security rule defined in the BLP model; we will therefore call our rule (A6) the **-property* security rule as well.

BELL-LAPADULA

It is instructive to compare the Chinese Wall policy with that of the Bell-LaPadula model [4, 5]. Both policies are described in similar terms: the composition and security attributes of the objects concerned and the rules for access, both in respect of simple-security and **-property*.

We take each of these topics in turn and identify a transformation which permits the Bell-LaPadula model (BLP) to describe that aspect of the Chinese Wall policy as faithfully as possible.

Object composition

The BLP model places no constraints upon the interrelationships between objects, in particular it does not require them to be hierarchically arranged into company datasets and conflict of interest classes. Instead it imposes a structure upon the security attributes themselves.

Each object within the BLP model [4] has a security label of the form (*class*, {*cat*}) where

- a) *class* is the classification of the information (e.g. unclassified, confidential, secret)
- b) *cat* is the formal category of the information (e.g. NOFOR).

Unlike the Chinese Wall policy, BLP attaches security attributes to subjects as well. They are complementary to object labels and have the form (*clear*, *NTK*) where:

- a) *clear* is the subject's clearance, i.e. the maximum classification to which he is permitted access
- b) *NTK* is the subject's need-to-know, i.e. the sum total of all categories to which he is permitted access.

Although this labelling scheme is seemingly quite different from that of figure 1, the two object compositions can be brought together quite simply by:

- a) reserving one BLP classification for sanitized information and another for unsanitized information
- b) assigning each Chinese Wall (x , y) label to a unique BLP category
- c) labelling each object containing unsanitized information with the BLP category corresponding to its (x , y) label and the BLP category of the sanitized dataset (x_0 , y_0).

The result of applying this transformation to the object composition of figure 1 is shown in figure 2.




Figure 2 Composition of objects into BLP-categories

Access rules

BLP also has a simple-security rule and a *-property rule:

- Simple security:** access is granted only if the subject's clearance is *greater* than the object's classification and the subject's need-to-know *includes* the object's category(ies)
- *-property:** write access is granted only if the output object's classification is *greater* than the classification of all input objects, and its category *includes* the category(ies) of all input objects.

We note that, in contrast to the Chinese Wall policy, these rules do not constrain access by the objects already possessed by a subject, but by the security attributes of the objects in question. Even so, a transformation does exist:

- give all subjects access to both sanitized and unsanitized objects
- define a need-to-know category for all possible combinations of object category which satisfy T2
- assign one such need-to-know category to each subject at the outset.

Referring to figure 2, for example, we would have 27 combinations:

{0, 1, 2, 3},
 {0, 1, 2, 6},
 {0, 1, 2, 9},
 {0, 1, 5, 3},
,{0, 7, 8, 9}

We then assign one of these to each subject, e.g. if we had three subjects, User-A, User-B and User-C, we could assign them:

User-A has {0, 1, 2, 3}
 User-B has {0, 4, 8, 6}
 User-C has {0, 7, 5, 9}

Now, the BLP simple security rule will work as we have granted subjects need-to-know combinations of company datasets for which no conflict of interest exists. BLP *-property also works. The category of the input objects must be included in the category of the output objects; input objects must therefore either contain sanitized information or be of the same company dataset as the output object. Moreover, unsanitized information dominates sanitized information.

Conclusion

These transformations give good account of BLP's ability to model to Chinese Wall security policy; they are not, however, entirely satisfactory.

For example, suppose in the real world management suddenly required User-A to have access to company dataset-8 because User-B was taken sick; what then? We cannot just change User-A's need-to-know to {0, 1, 8, 3} unless we know for certain that he has not yet accessed company dataset-2; otherwise we would violate the Chinese Wall. The BLP model does not provide the necessary information to answer this question.

Secondly, BLP only works if subjects are not given the freedom to choose which company datasets they wish to access. In other words, these transformations totally ignore the free choice nature of the Chinese Wall policy. This freedom of choice can be restored (e.g. by extending subject need-to-know to cover all company datasets) but only at the expense of failing to express the mandatory controls.

Thus, we can use BLP to model either the mandatory part or the free choice part of the Chinese Wall policy but not both at the same time. This is a somewhat fundamental restriction as the Chinese Wall policy requires both! The Chinese Wall model must therefore be regarded as distinct from BLP and important in its own right.

CLARK AND WILSON

It is interesting to consider the Chinese Wall policy in the light of Clark and Wilson's work. The Clark-Wilson model [1] defines a set of rules, based on commercial data processing practices, which together have the objective of maintaining data integrity. One such rule (E2), observed in many commercial systems [3], describes a finer degree of access control than that required by BLP. Compatibility of the Chinese Wall model with E2 is important because of the practical significance of that rule; it would not be useful if in order to satisfy E2 a system could not also satisfy the Chinese Wall policy.

Rule E2 requires access to be controlled on the basis of a set of relations of the form (user, program, object) which relates a user, a program, and the objects that program may reference on behalf of that user. Of greatest interest is the observation that if all users are permitted access to a given process, which in turn is allowed to access any object, then it does not follow that all users are permitted access to all objects. This may apply in the case of a Chinese Wall policy where, for example, some users may have access to a spreadsheet package but may not be permitted to use it on all company datasets to which they are permitted access.

A refined exposition of the Chinese Wall policy

The Chinese Wall policy as described so far refers to the users of the system. Thus we may interpret the use of the word *subject* in the various axioms, definitions and theorems as users. To be more explicit we define *user* as *a human being who causes or attempts to cause a transaction in a system*. In particular it is incorrect to associate *subject* with a (*user, process*) pair since, A2 for example, would imply that any particular user could gain access to two datasets within the same conflict of interest class simply by accessing them via different processes.

We may then accommodate Clark and Wilson's observation by simply requiring that:

- a) **users may only execute certain named processes (A7)**
- b) **those processes may only access certain objects (A8).**

Axiom A8 is defined in such a way as to maintain independence between what a user is allowed to do and what a process is allowed to do. This allows Clark and Wilson's observation to be incorporated into our model:

- a) without change to any of the foregoing theory (in particular we simply restate all axioms, definitions and theorems by replacing *subject* with *user*)
- b) redefining the overall *simple* and **-property* security rules to be the logical conjunction of all user-object, user-process and process-object permissions.

Thus a process is allowed to access an object if the user requesting it is allowed to execute it *and* that user, by the Chinese Wall policy, is allowed to access that object *and* that process is also allowed to access that object; write access is then permitted only if the **-property* rule (A6) is satisfied. Thus we conclude that the Chinese Wall model is not at variance with the observations of Clark and Wilson.

OVERALL CONCLUSIONS

In this paper we have explored a commercial security policy which represents the behaviour required of those persons who perform corporate analysis for financial institutions. It can be distinguished from Bell-LaPadula-like policies by the way that a user's permitted accesses are constrained by the history of his previous accesses.

We have shown that the formal representation of the policy correctly permits a market analyst to talk to any corporation which does not create a conflict of interest with previous assignments. Failure to obey this policy would be considered, at best, unprofessional and potentially criminally fraudulent. In some countries, showing that the policy has been enforced is sufficient to provide legal protection against charges of *Insider dealing*.

Thus this is a real commercial policy which can be formally modelled, but not using a Bell-LaPadula based approach. However, the Clark and Wilson's access control observations are not at variance with either the policy or its model as developed within this paper.

REFERENCES

- [1] CLARK, D.R., and WILSON, D.R.
"A Comparison of Commercial and Military Computer Security Policies" IEEE Symposium on Security and Privacy, Oakland, April 1987.
- [2] WIPCIS
"Report of the Invitational Workshop on Integrity Policy in Computer Information Systems (WIPCIS)" Published by the NIST (formerly the NBS), April 1988.
- [3] BREWER, D.F.C.
"The Corporate Implications of Commercial Security Policies", Proceedings of Corporate Computer Security 89, PLF Communications, London, February 1989.
- [4] BELL, D.E, and LaPADULA, L.J.
"Secure Computer Systems: Unified Exposition and Multics Interpretation" ESD-TR-75-306, MTR 2997 Rev. 1, The MITRE Corporation, March 1976.
- [5] McLEAN, J.
"The Algebra of Security" IEEE Symposium on Security and Privacy, Oakland, April 1988.
- [6] Her Majesty's Stationery Office, London, Securities and Investment Board Rules, Chapter III Part 5:08.
- [7] Her Majesty's Stationery Office, London, Financial Services Act 1986, S48(2)(h), 1986.

ANNEX A FORMAL MODEL

BASIC MODEL

Let S be a set of subjects, O be a set of objects and L a set of security labels (x, y). One such label is associated with each object. We introduce functions $X(o)$ and $Y(o)$ which determine respectively the x and y components of this security label for a given object o .

We will refer to the x as conflict of interest classes, the y as company datasets and introduce the notation x_j, y_j to mean $X(o_j)$ and $Y(o_j)$ respectively. Thus for some object o_j, x_j is its conflict of interest class and y_j is its company dataset.

Axiom 1:

$$y_1 = y_2 \implies x_1 = x_2$$

in other words, if any two objects o_1 and o_2 belong to the same company dataset then they also belong to the same conflict of interest class.

Corollary 1:

$$x_1 \neq x_2 \implies y_1 \neq y_2$$

in other words, if any two objects o_1 and o_2 belong to different conflict of interest classes then they must belong to different company datasets.

Proof:

$$y_1 = y_2 \implies x_1 = x_2 \quad (A1)$$

$$\text{not } (y_1 = y_2) \text{ or } (x_1 = x_2)$$

$$(y_1 \neq y_2) \text{ or not } (x_1 \neq x_2)$$

$$x_1 \neq x_2 \implies y_1 \neq y_2$$

Definition 1:

\mathbf{N} , a boolean matrix with elements $\mathbf{N}(v, c)$ corresponding to the members of $S \times O$ which take the value true if subject s_v has, or has had, access to object o_c or the value false if s_v has not had access to object o_c . Once some request $R(u, r)$ by subject s_u to access some new object o_r has been granted then $\mathbf{N}(u, r)$ must be set true to reflect the fact that access has now been granted. Thus, without loss of generality, any request $R(u, r)$ causes a state transition whereby \mathbf{N} is replaced by some new \mathbf{N}, \mathbf{N}' .

Axiom 2:

Access to any object o_r by any subject s_u is granted if and only if for all $\mathbf{N}(u, c) = \text{true}$ (i.e. by D1 s_u has had access to o_c)

$$((x_c \neq x_r) \text{ or } (y_c = y_r)).$$

Axiom 3:

$\mathbf{N}(v, c) = \text{false}$, for all (v, c) represents an initially secure state.

Axiom 4:

If $\mathbf{N}(u, c)$ is everywhere false for some s_u then any request $R(u, r)$ is granted.

Theorem 1:

Once a subject has accessed an object the only other objects accessible by that subject lie within the same company dataset or within a different conflict of interest class.

Proof:

If this proposition is untrue then it is possible for some subject s_u to have access to two objects, o_a and o_b , which belong to the same conflict of interest class but different company datasets, i.e.

$$\mathbf{N}(u, a) = \text{true and } \mathbf{N}(u, b) = \text{true and } x_a = x_b \text{ and } y_a \neq y_b$$

Let us assume without loss of generality that access to o_a was granted first. Then, when access to o_b was granted $\mathbf{N}(u, a)$ was already true and thus by A2 ($x_a \neq x_b$) or ($y_a = y_b$). Then for our two objects o_a and o_b to exist:

$$(x_a \neq x_b \text{ or } y_a = y_b) \text{ and } (x_a = x_b \text{ and } y_a \neq y_b)$$

$$(x_a \neq x_b \text{ and } x_a = x_b \text{ and } y_a \neq y_b) \text{ or } (y_a = y_b \text{ and } x_a = x_b \text{ and } y_a \neq y_b)$$

which is always false.

Theorem 2:

A subject can at most have access to one company dataset in each conflict of interest class.

Proof:

A subject need not have access to any dataset since by A3, \mathbf{N} is initially everywhere false. Suppose s_u then requests access to o_p . This request succeeds by A4 and our subject has access to one object within one company dataset.

By T1 the only objects then accessible to s_u lie within the same company dataset or within a different conflict of interest class, $x_q \neq x_p$. Hence at most only one company dataset within any particular conflict of interest class is accessible to a single subject.

Theorem 3:

If for some conflict of interest class X there are X_y company datasets then the minimum number of subjects which will allow every object to be accessed by at least one subject is X_y .

Proof:

Suppose there are N subjects and for some conflict of interest class, X, there are X_Y company datasets. Let all of these subjects have access to the same company dataset, i.e. N subjects have access to company dataset (X, 1); and, by T2, no subject has access to company datasets (X, 2)...(X, X_Y).

We can access one of these, say (X, 2), by reallocating one of the subjects with access to (X, 1) to (X, 2), i.e. N - 1 subjects with access to (X, 1), one with access to (X, 2) and X_Y - 2 inaccessible datasets.

By induction, after n similar reallocations we have N - n subjects with access to (X, 1), one each for (X, 2)...(X, n + 1) and X_Y - (n + 1) inaccessible datasets. In order that all data sets are accessible we require X_Y = (n + 1) provided, of course, that the number of subjects with access to (X, 1) is at least one, i.e. N - n > 0. Hence we require the smallest value of N such that:

$$X_Y = (n + 1) \text{ and } N - n > 0$$

$$N - X_Y + 1 > 0$$

which has a minimum when

$$N - X_Y + 1 = 1$$

i.e. when N = X_Y.

SANITIZED INFORMATION

Definition 2:

For any object o_a,

y_a = y_o implies that o_a contains *sanitized* information

y_a <> y_o implies that o_a contains *unsanitized* information

Axiom 5:

$$y_o \leftrightarrow x_o$$

In other words, if an object bears the security label y_o then it must also bear the label x_o and vice versa. T2 tells us that all subjects can access this company dataset.

Axiom 6:

Write access to any object o_b by any subject s_u is permitted if and only if N^u(u, b) = true and there does not exist any object o_a (N^u(u, a) = true) which can be read by s_u for which:

$$y_a \leftrightarrow y_b \text{ and } y_a \leftrightarrow y_o.$$

Theorem 4:

The flow of unsanitized information is confined to its own company dataset; sanitized information may however flow freely throughout the system.

Proof:

Let T = {(a, b); N^u(u, a) = true and N^u(u, b) = true for some s_u in S}. We will interpret (a, b) as meaning that information may flow from o_a to o_b. The reflexive transitive closure T* then defines all possible information flows.

Let B = {(a, b); (a, b) in O x O and y_a <> y_b and y_a <> y_o}. This is the set of all object pairs excluded by A6.

Thus the only possible information flows remaining after the introduction of A6 are given by C = T* minus B:

$$\{(a, b); \text{not } (y_a \leftrightarrow y_b \text{ and } y_a \leftrightarrow y_o)\}$$

$$\{(a, b); \text{not } (y_a \leftrightarrow y_b) \text{ or not } (y_a \leftrightarrow y_o)\}$$

$$\{(a, b); (y_a = y_b) \text{ or } (y_a = y_o)\}$$

Hence information may only flow between objects which belong to the same company dataset or originate from the sanitized dataset.

EXTENSION FOR CLARK AND WILSON

We now formally introduce the set P, the set of processes which we may interpret as those programs or sub-programs which a user may use to access whatever objects that the Chinese Wall policy grants him access to. We let A be a relation of S x P, representing those processes which a user is allowed to execute. Thus:

Axiom 7:

A user, s_u, may execute a process, p_f, if and only if (u, f) is a member of A.

We augment L to include a third attribute, z (i.e. L = {(X, y, z)}, where z is a member of some set Z. Z(o_j) is the function which determines the z-component of the security label of a given object o_j (z_j for short) and introduce PZ to represent the power set of Z. We then associate with each and every process, p_f, a member of PZ, determined by the function PZ(p_f), or p_zf for short. We assert that processes can only access objects whose z-attribute is included in those of the process, i.e.

Axiom 8:

A process p_f may only access an object o_r if

$$z_r \text{ subset of } p_{z_f}.$$

The access rules are now governed by A2-A4 and A6-A8. In particular an initially secure state exists when no user has ever accessed any object (A3) and the first ever access by any user to any object is entirely unrestrained (A4).

Users may, however, only access that object (say s_u and o_r respectively) via some process p_f where (u, f) in A and for which z_r subset of $p z_f$ (A7, A8).

Users may then access some other object o_r via p_f if and only if for all $N(u, c) \equiv$ true:

$$\begin{aligned} & ((u, f) \text{ in } A) \text{ and} \\ & (z_r \text{ subset of } p z_f) \text{ and} \\ & ((x_c \langle x_r) \text{ or } (y_c \equiv y_r)) \end{aligned} \quad (A2, A7, A8)$$

and, finally users may only write information to an object o_b provided that access is not prohibited by any of the above rules and that there does not exist any object o_a which can be read for which:

$$y_a \langle y_b \text{ and } y_a \langle y_o \quad (A6).$$
