8 CLOSED SYSTEM SECURITY MODELLING

Before a group can enter open society it must first close ranks.
Charles Hamilton,
Black Power!,
1967.

En boca cerrada no entran moscas.
Spanish Proverb.

8.1. INTRODUCTION

In this chapter, the problem of generating composable security properties is considered to result from the attempt to define security for isolated components, and subsequently retrofit composition into the model. Such models are here called ''open'' models, in that they are based on the analysis of an individual component operating in isolation, without regard to what might actually be producing or consuming the signals that enter or exit it. This chapter suggests the use of ''closed'' modelling as a basis for the definition of security properties, where a closed model is one in which the component under scrutiny is assumed from the start to be interconnected with an assortment of information sources, sinks, and processors.

Rethinking security modelling from this perspective presents an opportunity to review some fundamental ideas about information that may profitably be changed from those conventionally used in open-system models.

8.2. INFORMATION, FLOW, AND KNOWLEDGE

Our primary goal is to express certain restrictions on the flow of information within the composite system, and to detect whether a given system design will obey those restrictions. This brings up the question of how we will define information and how we will determine where and when it flows.

The definition of information and detection of its transmission in a model is a difficult problem, and one that cannot be said to have a single, universally applicable solution. In particular, definitions that may be usefully

applied in an open security model may no longer be reasonable in a closed one. Several fields of research have struggled with this question, and the many different answers that have been proposed reflect some very different assumptions. We present here a survey of some of these methods.

8.2.1. Noninterference and Nondeducibility

For completeness, we briefly mention these security definitions which have become the conventional formal definitions of confidentiality. Noninterference, which is, strictly speaking, applicable only to deterministic systems, states that a system is secure if and only if the sequence of inputs from and outputs to low-level users is unaffected by changes in high-level users's inputs.

Nondeducibility is, in practice, a similar concept, though derived from slightly different principles. It requires that the hidden information (generally, high-level inputs) not be deducible from low-level observations. Unlike noninterference, nondeducibility may be applied to non-deterministic (more correctly, probabilistic) systems, in which it takes the form of a requirement that there must exist at least one system execution that combines any given high-level input string with any given low-level behaviour.

One other difference arises if the system has any inputs that are considered neither high nor low. Noninterference requires that the low-level observations remain unaltered when high inputs change, and so these ''other'' inputs must also be held fixed. Nondeducibility, however, only requires that the high and low coexist in some execution, and so permits the other inputs to be varied to meet this condition.

Both definitions have problems dealing with probabilistic systems. Noninterference simply isn't defined so as to be applicable to them. Nondeducibility can be applied, but may give unreasonably optimistic results if some desired system execution exists only at a very low probability.

Both definitions assume that any illicit channel may be exploited no matter how much computation may be required to recover the sensitive data. That is, they assume infinite computing resources on the part of the enemy. They also assume that the enemy has full knowledge of the principles of operation of the system, and can effectively compute all possible system behaviours.

8.2.2. Information Theory

The first quantitative treatments of information are connected with the study of communication. In connection with telegraphy, Nyquist [Nyquist24] used a logarithmic measure of the information content of a signal. If the signal can take n possible values, it can store $\log_b n$ units of information, in the

sense that it would take that number of b-way signals to represent a single n-way one. The choice of b affects only the size of the units, and if the base 2 is chosen the resulting units are known as *bits* [1].

The subsequent development of communication theory led to the employment of several related measures. In considering the problem of efficiently encoding a (discrete) data stream, it is advantageous to employ a variable length code, such that high-frequency source symbols (or sequences of them) are encoded into short codes and rare source symbols into long ones. The *entropy* of a signal s^n of length n symbols is defined as

$$H_n(S) = -\sum_{s^n} p(s^n) \log_2 p(s^n)$$

where the sum is taken over all possible length n source strings. It is the average number of bits needed to encode each string s^n in the most efficient possible encoding. This code efficiency is sometimes improved, and never worsened, by encoding longer and longer strings of source symbols at a time. The theoretical lower bound on the average number of bits required per source symbol (from [Shannon49]) is

$$\begin{aligned} H(S) &= \lim_{n\to\infty} \frac{1}{n} H_n(S) \\ &= -\lim_{n\to\infty} \frac{1}{n} \sum_{s^n} p(s^n) \log_2 p(s^n) \end{aligned}$$

where the sum is taken over all s^n, the source symbol strings of length n, and where $p(s^n)$ is the probability of the string s^n. This $H(S)$, called the *entropy* (per symbol) of the source, represents its information content in the precise sense that the source may not be encoded using, on average, fewer than H bits per source symbol. If only one source symbol $s \in S$ is encoded at a time, it theoretically requires $-p(s) \log_2 p(s)$ bits in a minimum length code, from which it is sometimes concluded that the lower-probability source symbols ‘‘contain more information’’ than the higher probability ones.

In coding and communication systems it is also sometimes useful to measure the mutual information content of two signals, considered as sequences of random variables. That is, how much information from one signal can be gained by observing the other. The *average mutual information* between X and Y is defined (in [Gray90]) as

$$I(X;Y) = \sum_{x^n, y^m} p_{X,Y}(x^n, y^m) \log_2 \frac{p_{X,Y}(x^n, y^m)}{p_X(x^n) p_Y(y^m)}$$

The consistency of these definitions is demonstrated by the fact that the information $I(X;X)$ communicated by X about itself equals $H_n(X)$, the

---

[1] Another popular choice for b is e, yielding the natural logarithm, in which case the units are known as *nats*.

information content of X.

In deriving expressions for the capacity of noisy bandlimited communication channels, Shannon used entropy as a measure of information content of a message. His use of this measure was narrowly restricted: the entropy of a set of signals, divided by the time required to communicate them over a channel, yields the signalling rate of the channel in units that are independent of the mechanism used. Equivalently, it is the signalling rate assuming the most efficient possible encoding of the source data.

These types of measures have recently become popular in the computer security field. If the system is viewed as a communications medium, as has been done in the traditional open security model, it is not surprising that communication theory should find application. This is particularly apparent in the case of covert channels, where it may be of interest to perform a quantitative analysis of a channel's transmission capacity [Millen87]. Related probabilistic techniques have also been proposed for modelling non-deterministic systems without suffering the overoptimism exhibited by straight nondeducibility [Wittbold90] [McLean90].

However, in the study of communication systems the ''information'' source and sink are left unspecified. In fact, it is common to model the data source as being random, with some known (and hopefully exploitable) statistical properties. Whether or not there is any meaning to these bits, and whether some bits are more important than others, is of no interest to the communications engineer, and is not reflected in the information theoretic definitions.

Indeed, this lack of interest is mandated by the nature of the communication problem itself. Since those systems do not create nor interpret the bits they convey, there is no ''handle'' that the communications theorist can use to assign significance to the data. For this reason, it is perhaps unfortunate that the term ''information'' has come into popular use as a synonym for signal entropy. The bandlimited signal with the highest entropy is white Gaussian noise, which we would normally consider to be information-free [2]. On the other hand, a signal which is almost always a "0" and only very rarely a "1" (eg., only when the missile fires) has very low entropy.

Open system security models inherit a disinterest in the nature of information when communication theoretic measurements are used, unaugmented, to characterize leaky information systems. In this case, however, the disinterest cannot reasonably be justified. Information systems do often interpret data, and occasionally even invent it. The opportunity exists to

---

[2] It is admitted that there exist cases in which noise may have a significance, derived from its application, e.g. if it is used to generate keys or passwords. We will tentatively accept this as ''information'', and worthy of protection, although with a truly complete security model it should be possible to derive the necessity of protecting key material rather than assuming *ab initio* that it is sensitive ''information''.

distinguish between a one bit per hour leak of an unspecified sensitive file and a similar leak of a password, because the system ''knows'' in some sense which bits are passwords and which are not.

Another criticism that may be levelled at information-theoretic definitions is that they may be making unreasonable assumptions about the capabilities of the receiver or, when these methods are applied to secure systems, of the adversary. If what an observer sees is affected by sensitive information, it is assumed that the observer will be able to decode those observations into some assertion about the sensitive information. In other words, whatever ''transfer function'' is represented by the transformation of high inputs to low perturbations is assumed to be invertible.

And yet, it is known that some functions whose inversion is theoretically possible (i.e. with unbounded computing power) are, in practical terms, not invertible. For example, assume that a sensitive input consists of a set of large prime numbers, and the output observed by the adversary is their product. Since the prime factors of any number are unique, the observation logically implies the input data and so information theory concludes that the adversary possesses all the sensitive data. But, factoring large numbers is computationally intractable (for sufficiently large numbers, that is), so the function may fail to be invertible in practice.

8.2.3. Knowledge Modelling

Hintikka [Hintikka62] is reported to be the first to attempt to model the concept of knowledge, using the notion of *possible worlds*. In this approach, an agent does not know the state of the real world, but is able to rule out some possibilities based on observations. A predicate ψ is then *known* precisely if it is true in each of the remaining possible worlds. This notion forms the basis of much of the formal reasoning about knowledge in the Artificial Intelligence research community.

Logical statements about knowledge can be formally expressed and manipulated by the introduction of the *modal operator* K_i, enabling such assertions as $K_i \psi$, meaning ''player i knows ψ''. These modal operators, in turn, are often given a firm semantic interpretation through *Kripke structures*. A Kripke structure $M = (S, \pi, \mathcal{K}_1, \cdots, \mathcal{K}_n)$ is a set of *states* S, a π such that $\pi(s)$ yields, for each state s, a truth assignment to the primitive propositions, and some number of relations between states $\mathcal{K}_i$, one for each player. The modal operator semantics are then defined for K_i as follows: $M, s \models K_i \psi$ (read, '$K_i \psi$ is *satisfied* in state s of structure M'') if $M, t \models \psi$ for all t such that $(s, t) \in \mathcal{K}_i$. The $\mathcal{K}$ relations can be used to define the knowledge of each player, in the sense that $(s, t) \in \mathcal{K}_i$ iff states (i.e., worlds) s and t are indistinguishable as far as player i is concerned. One can also define a dual operator $P_i p$ as $\sim K_i \sim p$ which can be interpreted as saying that ''player i thinks that p is possible'' (i.e., does not know p to be false).

Using Kripke structures to define the modal operators means that they are, formally at least, simply new notation for conventional mathematical notions, and so inherit all the soundness and self-consistency of the underlying theory. It also means that several useful axioms may be derived for them, such as

$K_i \psi => \sim K_i \sim \psi$ ''Cannot simultaneously know ψ and $\sim\psi$''

$K_i \psi => \psi$ ''Can only know things that are true''

$K_i \psi => K_i K_i \psi$ ''If i knows ψ, then he knows that he knows it''

$\sim K_i \psi => K_i \sim K_i \psi$ ''If i does not know ψ, he knows that he does not''

Actually, the case where the $\mathcal{K}$ are equivalence relations (i.e. reflexive, transitive, and symmetric) is only one of the cases treated in [Kripke63], and is a model for a logic known as S5. In S5, all four of the above listed axioms are valid. Some work mentioned in this section are based on other logics whose models do not have equivalence relations, notably the logic S4 used in [Glasgow89], and KD45 used in [Moser89]. Figure 8.1 shows how different modal axioms are tied to properties of the model's $\mathcal{K}$ relation. (Note that axiom **T** implies axiom **D**). Various combinations of these have been given conventional names; a partial list is presented in Table 8.3. Note that some of the logics can be characterized in several equivalent ways.

Logic	$\mathcal{K}$ Properties	Valid Axioms	
		Basic	Implied
T	reflexive	**K T**	D
D	not isolated	**K D**	
B	reflexive, symmetric	**K T B**	D
S4	reflexive, transitive	**K T 4**	D
KD45	transitive, euclidean, not isolated	**K D 4 5**	
KB4	symmetric, transitive symmetric, euclidean	**K B 4**	
S5	reflexive, transitive, symmetric reflexive, euclidean not isolated, symmetric, transitive not isolated, symmetric, euclidean	**K T 5**	B D 4

Table 8.3: Some well-studied varieties of modal logics.

When the desired modality relates to knowledge, we can interpret the ''states'' as ''possible worlds'' and the relation between states as ''indistinguishability'', in which case the modal expression $K\,p$ is true in some state s iff p is true in all states indistinguishable from s, just as Hintikka proposed. Clearly, the indistinguishability relation is reflexive, transitive, and symmetric, and therefore the logic that naturally results from this formulation is the S5 logic defined above.

With these formalities completed, the resulting logic is a useful tool for representing and for reasoning about states of knowledge in a system. [Fagin86] contains an excellent discussion of the strengths and limits of this representation. It also describes a general model of knowledge in a multiplayer system, and explores the variations that result from varying some basic assumptions about the system, such as whether it is synchronous and whether knowledge, once acquired, is ever forgotten. From this starting point logical systems have been created that represent not only knowledge but knowledge about what is known (by oneself and others) [Fagin84].

Other modal logics can also be created based on a Kripke structure model. If a temporal logic is desired, one that permits statements of the form $\Box\,p$ (read, ''henceforth p'') with the meaning ''p is true now and will remain true''. In this case, the ''states'' of the Kripke structure can be interpreted as instants in time, and the relation between them as $(s_i, s_j) \in \mathcal{K}$ iff moment s_j is at or later than s_i. This relation is clearly reflexive, transitive, but not symmetric, and so the $\Box$ operator it defines has the properties of an S4 logic. Its dual operator $\Diamond$ (read ''eventually'') then means that p is either true now or will be true at some future time.

It is also possible to create a modal logic without reference to a defining structure. For example, we may have a notion of ''belief'' that is similar to knowledge, but differs in that something we believe may turn out to be false. A logic could be designed by starting with an S5-type knowledge logic, in which all of **K**, **D**, **T**, **4**, **B**, and **5** are valid, renaming the modal operator B for belief, and removing the axiom **T:** $B\,p \Rightarrow p$ so that belief no longer implies truth. However, without the Kripke structure as a guide it may not be apparent that the original set were not independent. In fact, **T** is a consequence of **B**, **4**, and **D**, or alternately of **B**, **4**, and