A Note on Synergistic Effects of Protection Mechanisms

by Dr. Frederick B. Cohen ‡

In this paper, we examine the synergistic effect as it applies to protection systems. We begin with a brief history of protection systems, attacks on protection systems, and synergistic effects in other fields. We then describe several examples of synergistic effects in information protection and discuss how and why these synergies exist. Next we philosophize about a possible explanation for synergistic effects in general and in protection systems in particular, and give an intuitively based mathematical model for protection synergy. Finally, Finally, we summarize results, draw conclusions, and describe further work.

search terms: Information Protection, Access Control, Privacy, Integrity, Trusted Systems, Password Protection, Authentication, Auditing, Computer Security, Synergism

‡ This research was funded by ASP, PO Box 81270, Pittsburgh, PA 15217, USA

1 Background

In the course of investigating protection system design, it has become increasingly apparent to us that protection mechanisms interact in non-trivial ways. On several occasions, we found instances where two protection mechanisms operating in conjunction with each other cover a larger set of attacks than the sum of the sets of attacks covered by the individual mechanisms. This appears to present a case for claiming a synergistic effect between protection mechanisms, and consequently, the presence of a deeper underlying phenomena.

1.1 Protection Systems

There are many models for protection in computer systems, a large number of different mechanisms for implementing these models and still more implementations of these mechanisms. We do not propose to be comprehensive or even even-handed in our coverage in this paper, but for the purpose of giving useful examples, we will refer to several models and mechanisms, and pause briefly here to cite them.

1.2 Attacks on Protection Systems

Known attacks against protection systems are far too numerous to list here and we are almost certainly are unaware of some attack mechanisms in actual use. Again, we will describe some well known attacks in order to provide examples.

1.3 Synergy as a Scientific Principle

Synergy as a scientific principle has existed for quite some time. In fact, to many non-scientists, synergy supports an argument that science cannot explain many of the things that happen, and is thus impotent in regards to understanding the human condition. It is fairly common to hear expressions like 'The whole is greater than the sum of the parts'. From a scientific viewpoint, synergistic effects are widely recognized as reflecting some unknown or misunderstood principle. We point out some examples of synergy in other areas of science to support its validity as a principle to be applied to information protection.

2 Examples of Protection Synergy

Information protection provides a very rich set of examples wherein synergistic effects may be noted. We describe several such examples here to demonstrate the existence of synergistic effects and to show how synergy impacts the protection provided by modern information systems.

2.1 Auditing and Passwords

With auditing alone, we cannot reliably detect attempts at unauthorized use of a system. For example, anyone who attempts to use the system being audited can simply claim to be any legitimate user and perform all of the tasks that the legitimate user would be able to perform. Similarly, with password based authentication, we cannot reliably detect unauthorized use of a system, because an attacker can try passwords over an arbitrary period of time, and will eventually succeed in guessing a password and gaining unauthorized access. In combination however, these two techniques can be quite effective, because auditing can detect failed attempts at guessing passwords before the probability of successful attack becomes significant, and thus provides the means for reliably detecting the attack.

2.2 Privacy and Integrity

The Bell-LaPadula [6] model is designed to provide privacy by classifying information at different 'security levels' and controlling the flow of information between those levels, but because of covert channels [8] and the lack of integrity protection, such systems can be easily made ineffective [5,7] by computer viruses. Similarly, integrity shells [7,9,10], which provide an effective defense against computer viruses through the use of cryptographic checksums, can be made impotent unless cryptographic keys are kept private. Again we have a synergistic effect wherein viruses are only made ineffective when both privacy and integrity protection are in place.

2.3 Logical and Physical Protection in a Network

'Peer networks' often have different protection measures at different sites, and as a result of synergistic effects, protection at all sites may be compromized. One such example is the case where a physical protection scheme at one site and a logical protection scheme at another site cause an overall breakdown of protection [7]. The transitive nature of virus propogation in conjunction with the peer equivalence of printers at different sites allows a virus entering a printer during maintenance at a logically secured site to pass to the physically secured site as a peer. Once in the physically secured site, it propagates throughout and returns to the logically secured site by again exploiting the peer relationship. In this case, the synergy works for the attacker because either site alone would be safe from this attack, but together, they are vulnerable.

3 A Possible Explanation for Protection Synergy

As we have seen, in information protection, synergy is commonplace, and in order to have effective protection, we must almost certainly consider synergistic effects at some level.

3.1 The Protected Subspace of the Computational Space

We can view the state space (S) of an information system as consisting of two mutually exclusive subspaces; the protected subspace (S_p); and the unprotected subspace (S_u). Hence:

$$S = S_p + S_u \text{ and } S_p - S_u = S_u - S_p = \emptyset$$

The protection policy defines the partition of S. For all nontrivial policies, and assuming a Turing machine model of computation [1], both S_p and S_u are infinite [2]. For all real-world systems, the state space is infinite, thus making most non-trivial problems relating to differentiating S_p from S_u too complex for practical resolution. We note that S encompasses worldly states (e.g. what human individuals know) as well as the computational states of computer systems (e.g. the accessibility of particular information by a particular authenticated identity).

As a practical matter, we normally model protection by seperating the problem into several non-orthagonal dimensions (e.g. accountability, integrity, privacy, availability), and apply techniques (e.g. audits, integrity shells, access controls, redundancy) to cover portions of those dimensions. By doing this, we cover a protection space intended to approximate S_p. Unfortunately, we may get very different coverage of S than S_p. To the extent that there is mismatch between S_p and the actual coverage, S_p may remain uncovered, resulting in a false sense of security, or some of S_u may be covered resulting in undesired protection.

3.2 A Multidimensional Space

We can mathematically model the protection space created by the multidimensional approach using a set of non-orthagonal coordinates $D = d_1, \ldots, d_n$ with the image of each of the techniques covering a subspace of each dimension. When we apply a technique (e.g. access control), we normally associate its impact on protection with covering a subspace of a single dimension (e.g. privacy), but this is only an approximation.

A more accurate depiction of the multidimensional model would be to describe the impact of a technique in terms of its image on each of the non-orthagonal dimensions of D. A still more accurate depiction would be to depict the shape of the coverage provided by the technique in D, but none of these is necessarily an accurate depiction of the coverage of S_p by the technique.

A further complexity results from the fact that the protection dimensions and techniques may not in fact cover S at all, exen though they may appear to do so. Since we are dealing with finite state automata with undecidable behavior, only a very small subset of the protection policies we may attempt to use provide actual coverage of the state space. Implied rights (e.g. implied flows under a POset [3]) may result in a completely uncovered S even though protection techniques (e.g. access control) are in place [4,5].

This problem arises from the fact that when we provide protection, we are trying to limit the achievable configuration of a system with worldly states. The solution to this particular problem comes from recognizing that information flow in a general purpose information system with sharing is transitive, and thus we can only limit the achievable states with a protection model which limits transitive flow. This is done mathematically by using a POset structure [4,5], but there is no guarantee that such a mechanism cannot be bypassed through some synergistic attack. Imperfect authentication, for example, could make a state in S_u achievable.

3.3 Probabilistic Spaces

It should now be clear that S_p is infeasible to accurately model or cover with current theory and practice, and therefore that any current protection system can fail. Even at the level of current physical theory, finite probabilities are associated with spontaneous transformations of matter, so to provide perfect protection, we would have to posit a different kind of physical universe.

As an alternative, we may try to model protection in terms of a probabilitic space wherein we associate probabilities with reaching each state in S from each other state in S. For each initial state $s_i \in S_p$, we can then assess the likelihood of reaching each of the states $s_j \in S_u$ through each of the paths from s_i to s_j, and hence derive the likelihood of reaching S_u from S_p based on the likelihood of starting in each $s_i \in S_p$. Hence we have a model as follows:

$$P = (S, \Pi : S \times S \rightarrow \Re_{0 \leq \pi \leq 1})$$

where P is the called the probability state space, S is the previously discussed state space of the system, and Π is the probability function mapping pairs of states from S into reals in the range of 0 to 1 inclusive.

We often use a simplification of this model by making assumptions about independence and grouping large numbers of states. As an example, we commonly use information theoretic approaches to approximate the difficulty of unauthorized access when passwords are in use.[11]

3.4 Where to Look For Synergism

If we take this model to be accurate, synergistic effects will appear at different places depending on how we choose to view the protection space. If we view P, we will almost certainly have an intractable problem detemining coverage, but our depiction will be accurate. If we view the shape of coverage in D, we will see synergism wherever the D inaccurately reflects P. If we view the image of coverage on each of the dimensions of D, we will also find synergism where the shape of coverage in D is not completely modeled by the image on each $d_i \in D$ (e.g. not an n-cuboid in D [1]). Finally, if we take the common view of coverage only being the image of coverage on a single dimension of D, we will also find synergism in the image of coverage in the remaining dimensions of D.

One interesting point here is that there is a hierarchy of views, each more accurate than the next, with the least accurate being the most common view wherein protection mechanisms are seen as an image of their coverage on one dimension of D, and the most accurate being the probability space view of P. As we use more accurate viewpoints, we get higher complexity and a more accurate depiction of the actual coverage.

3.5 How Attackers Exploit Protection Systems

Another way to look at the impact of synergy is from the attackers viewpoint. We

[1] Since the dimensions of D are non-orthagonal, the result of reconstruction from images on each dimension of D will not be a true n-cube

may think of an attacker as someone who 'steps out of' the system of controls, but from the hierarchy of views, we may see the successful attacker as someone who exploits a more accurate view of the protection system than the defender.

Suppose, for example, that the defender designs a system to cover a large portion of a multidimensional space D

3.6 Changes in the Multidimensional Space

4 Summary, Conclusions, and Further Work

5 References

1. **1)** A. Turing, "On Computable Numbers, with an Application to the Entscheidungsproblem", London Math Soc Ser 2, 1936.
2. **2)** M. Harrison, W. Ruzzo, and J. Ullman, "Protection in Operating Systems", CACM V19#8, Aug 1976, pp461-471.
3. **3)** F. Cohen, "Protection and Administration of Information Networks with Partial Orderings", IFIP-TC11, "Computers and Security", V6#2 (April 1987) pp 118-128.
4. **4)** F. Cohen, "Computer Viruses", Dissertation at the University of Southern California, 1986.
5. **5)** F. Cohen, "Computer Viruses - Theory and Experiments", DOD/NBS 7th Conference on Computer Security, originally appearing in IFIP-sec 84 (1984), also appearing as invited paper in IFIP-TC11, "Computers and Security", V6#1 (Jan. 1987), pp 22-35 and other publications in several languages.
6. **6)** D. E. Bell and L. J. LaPadula, "Secure Computer Systems: Mathematical Foundations and Model", The Mitre Corporation, 1973.
7. **7)** F. Cohen, "A Short Course on Computer Viruses", ASP Press, PO Box 81270, Pittsburgh, PA 15217, USA, 1990.
8. **8)** B. W. Lampson, "A note on the Confinement Problem", Communications of the ACM V16(10) pp613-615, Oct, 1973.

9. **9)** F. Cohen, “A Cryptographic Checksum for Integrity Protection in Untrusted Computer Systems”, IFIP-TC11 Computers and Security, V6(1987).

10. **10)** F. Cohen, “A Cost Analysis of Typical Computer Viruses and Defenses”, IFIP-TC11, “Computers and Security”, (accepted awaiting publication, 1991).

11. **11)** F. Cohen, “Algorithmic Authentication of Identification”, Information Age, V7#1 (Jan. 1985), pp 35-41.