NSPW Panel: 'Diversity as a Computer Defense Mechanism' - Position statement

Bev Littlewood
Centre for Software Reliability
City University
London EC1V 0HB, UK
+44 20 7040 8420
b.littlewood@csr.city.ac.uk

ABSTRACT

Redundancy and diversity are commonly applied principles for fault tolerance against accidental faults. Their use in security – to protect against intentional faults – is attracting increasing interest. I propose that there is a need for a formal probabilistic treatment, similar to the one that has brought successful insights in reliability.

Keywords

Diversity; probabilistic modeling; intrusion detection

1. INTRODUCTION

I am not a security expert! My background is in probabilistic models for the assessment of reliability and safety of software-based systems. Much of this work has been concerned with models of diversity: diversity of process (e.g. testing); diversity of system versions (e.g. multi-version software fault-tolerance); diversity of arguments to support dependability claims (e.g. multi-legged safety cases).

The formality of probabilistic modeling has brought greater rigour to the claims that are made about the efficacy of diversity: for example showing that claims for 'independence' must be treated with great scepticism. In my contribution to the Panel, I want to argue that probability modeling can bring a similarly useful rigour to the study of diversity in security.

2. THE NEED FOR PROBABILITY

There is not space here to preach at length about the necessity for a probabilistic approach (see [1] for a discussion on this). Suffice it to say that there is an inherent uncertainty here (e.g. in the process of attacks upon a system) and probability is the most appropriate calculus for uncertainty. Alternative formalisms to probability, such a fuzzy/possibility theory, do not have the power of probability, nor do they easily fit into a wider engineering framework, for example enabling quantitative risk assessment.

Rather informally, one might measure the security of a system in terms of its (probabilistic) ability to resist attacks. It is often argued by members of the security community that 'intentionality' in the attack process rules out the use of a probabilistic approach. I think this is mistaken, but intentionality *does* affect the details of any probabilistic modeling that would be carried out. For example, consider the stochastic point process of (security) failure events on the time axis. There seem to be two different kinds of uncertainty involved here. In the first place, there is a stochastic process consisting of 'novel' security breaches – these can be thought of as the first discoveries of security vulnerabilities. Secondly, there are 'local' stochastic processes comprising the cascades of breaches that will follow such discoveries, as their existence becomes known to possible attackers.

It seems to me that it is the first of these stochastic processes that should be of most interest, and it is this one that we have addressed in our work. In particular, we would like to know how this (non-stationary) stochastic process will develop in the future, as new vulnerabilities are discovered and removed. This problem is very similar to the classical problem of reliability growth in software, as design faults are identified in use and eliminated. The fact that security flaws are being sought *intentionally* by (possibly) malign humans, rather than (as in the case of reliability growth) by the perversity of nature, does not seem to me to detract from a probabilistic treatment. On the other hand, what happens after flaws have been discovered, and many attackers exploit them before they can be removed, is a very different stochastic process, and the role of intentionality may need to be modeled here (but it will still need to be a *stochastic* model because of inherent uncertainty).

3. DIVERSITY AND SECURITY

There has been a growing interest in the application of diversity in security in recent years. Many people think, for example, that the use of diverse intrusion detection systems may be useful, but others disagree. The issues here are, of course, empirical ones concerning efficacy: e.g. is diversity cost-effective in competition with other approaches? Little has been published on this, and it is often rather informal and anecdotal. For example, there are inappropriate appeals to 'independence' – just as there were in *safety* 20 years ago.

So what is diversity? Informally, it involves doing things in *deliberately different* ways (and comparing the results). The goal is always *diversity of failures*: e.g. in a 1-out-of-2 system we would ideally like *A* to succeed when *B* fails, and vice-versa. However, we would not generally believe deterministic claims – there is an inherent uncertainty about failure behaviour. The

NSPW 2005 Lake Arrowhead CA USA

efficacy of the diversity in a 1-out-of-2 system, then, would depend on the *conditional* probability of *A* failing *given that B has failed.*

How do we get diversity? The simplest way just involves *separation*, e.g. simply isolating software development teams to prevent communication. More constructively, diversity can be *forced* by requiring the use of different development processes, tools, etc. Sometimes it may even be possible to 'force' diversity in such a way as to make the versions different *with respect the particular faults that might be anticipated.*

Whilst there are plausible grounds for believing that, in some sense, diversity is a good thing, some claims are not believable: e.g. the *impossibility* of common failures; e.g. *statistical independence* between failures. So the important question is 'how good is it?' A key issue turns out to be dependence between available versions.

Imagine that you have several IDSs available. These are based on different principles, so might be expected to fail differently. It seems sensible to combine them. Suppose you can only afford to use two (e.g. because of costs; e.g. because of the problem of false positives). Do you necessarily choose the best two? Clearly not; in fact you would like their failures to be 'complementary' – ideally what is not caught by the first is caught by the second.

We can reason formally about all this. Imagine we could enumerate all possible attacks. For each attack, *x*, IDS *A* has probability $\theta_A(x)$ of missing it; for *B* the probability is $\theta_B(x)$; and for any *specific* attack the failures of the two IDSs are independent. Then for a *randomly selected* attack *X* we have

$$\text{P(A and B both fail to detect X)} = \sum_{x \in D} P(x)\, \theta_A(x)\ \theta_B(x) = Q_A\, Q_B + cov_x(\theta_A(x).\theta_B(x))$$

where $Q_A = \Sigma P(x)\ \theta_A(x)$ is the probability of *A* failing on a randomly selected attack (i.e. it is the 'unreliability' of *A*); *P*(*x*) is the probability that an attack is of type *x* (this distribution over all possible *x* characterises the nature of the threat environment, *cf* operational profile in reliability).

The key to all this lies in functions like $\theta_A(x)$, which we call 'difficulty functions': they represent how attacks differ in their difficulty, *and how such difficulty varies between A and B.*

What does this key equation tell us? Most importantly it shows how the difficulty functions determine the efficacy of a diverse system. In particular, it shows the trade-off between the efficacies of the individual IDSs (represented by Q_A, Q_B) and the diversity between them (represented by the covariance term).

Note that the first term on the right hand side of the equation represents the probability of simultaneous failure of both *A* and *B* under a naïve independence assumption. The equation thus makes clear that failure independence is a very special case: it can be regarded as just one point on a continuum of dependence. In effect this rules out claims for independence.

On the other hand, it shows that diversity is indeed 'a good thing', in the sense that the 1-out-of-2 system is always better than each component IDS – as, indeed, one would expect (diversity costs more!). In fact there are some precise theorems about efficacy which space constraints prevent being discussed here: for example, it can be shown that 'more diversity is always better' (and 'more' is defined precisely). See [2] for details in the context of reliability: similar results apply in the case of security.

4. DISCUSSION

Is any of this useful? How would it be used in practice?

Clearly, the difficulty functions, as defined here over all individual attacks, will be unknown in practice, and could not be estimated. However, in [1] we discuss how they could be defined over equivalence classes of attacks, and in this case they can be estimated statistically (at least in principle). We intend to look at historical data, and data from honey-traps, and attempt to estimate these for some sample IDSs.

Leaving aside issues of evaluation of particular systems, the modeling here brings some rigour and clarity into understanding the role of uncertainty in diversity. The notions of variation, and covariation, of 'difficulty' turn out, somewhat surprisingly, to be fundamental.

Whilst independence has been shown to be an unattainable goal, there is, in fact, the possibility of doing *better* than independence. This will happen when the covariance term in the equation is negative: roughly when what is 'difficult' for *A* is 'easy' for *B*, and vice-versa.

5. OTHER ISSUES, FURTHER WORK

So far, only diversity of protection against attack has been discussed. But the probability models are versatile, and can also be applied in wider contexts. Some areas where it may be fruitful to apply the models are:

- *Diversity of intruders.* Clearly diversity will be useful in a team of intruders: how do you pick the best red team of size *m* from a population of size *n*?
- *Diversity of intrusion procedures.* What is the best mix of attack procedures to use to find vulnerabilities? The probability modeling ideas have been applied to diverse fault-finding in software [3].
- *Diverse intruders, diverse sensors.* Can the models be extended to this case? It introduces a further element of diversity (or covariation) – between the attackers and the sensors. Interestingly, this has not been addressed in reliability: there seems to be a view that 'nature' does not mount diverse threats – is this true?

6. REFERENCES

[1] Littlewood, B. and L. Strigini. Redundancy and diversity in security. in *ESORICS (European Symposium on Research in Computer Security)*. 2004. Sophia Antipolis: Springer.

[2] Littlewood, B. and D.R. Miller, Conceptual Modelling of Coincident Failures in Multi-Version Software. *IEEE Trans on Software Engineering*, 1989. **15**(12): p. 1596-1614.

[3] Littlewood, B., et al., Modelling the effects of combining diverse software fault removal techniques. *IEEE Trans Software Engineering*, 2000. **26**(12): p. 1157-1167.