Non-interference, who needs it?

Peter Ryan, Carnegie Mellon (Moderator)
John McLean, NRL
Jon Millen, SRI International
Virgil Gligor, University of Maryland, College Park

1. Overview

The concept of non-interference seeks to characterize the absence of information flows through a computer system. The intuition is startlingly simple. Suppose that we want to assert that no information may flow from user A to user B via the system S. We characterize this by asserting that B's view of S is unchanged by any alteration in A's behaviour. It is thus asserting that A can have no causal influence on B's interactions with and observations of the system.

Non-interference is such a simple and obvious characterization of MLS confidentiality that the security community is understandably reluctant to give it up. However, it has well known problems. First, in real systems high-level input interferes with low-level output all the time. High-level files can be encrypted, sanitized, or simply downgraded and sent on their way over low-level networks. Second, after fifteen years of trying, we still don't have any consensus as to what is the "correct" nondeterministic formulation of it. Nondeterministic versions tend to be too weak (e.g., Nondeducibility), too strong (e.g., Noninference), too cumbersome (e.g., PNI and AFM), too limiting (e.g., the Roscoe, Woodcock, Wulf determinism approach) too Baroque (e.g., Restrictiveness), or some combination of the five. In [2] it is argued that, in a process algebraic setting, the characterization of non-interference reduces to characterizing the equivalence of certain processes. This in turn is a fundamental and difficult question of theoretical computer science and one to which there is no universally agreed answer. Thus it is not even clear whether a "correct", Platonic notion of secrecy actually exists.

Non-interference would seem to be a fundamental notion in information security. It could be argued that, if we cannot get the specification and verification of the absence of information flows right, we really don't understand the foundations of our subject. On the other hand, it is such an abstract formulation that it seems remote from real concerns of security managers, policy makers and the developers of secure systems. Most "real" security policies are concerned with specifying who has access to what resources under what circumstances. Non-interference is never mentioned. Furthermore, non-interference is in practice impossible to realise in any real system: contention for resources etc render it infeasible. Even the so-called One-Way- Regulators (e.g. the NRL Pump) allow some downward flow, albeit of low channel capacity.

The study of non-interference arose from the need to understand why covert channels were possible, at a time when the only theoretical security models were access-control models, which were unable to explain them. The first wave of responses consisted of information flow models, which used the syntactic structure of statements to recognize possible flows, such as "indirect flow" from the condition of an if-then statement to variables that might be modified in its body. These models were found to overestimate flows. The second wave of models were the deterministic non-interference models, which were based on the notion of functional dependency. These models explained some covert channels, and found flows only where they really existed. Subsequent varieties of models found more channels by allowing for nondeterminacy in the computer system model, either "possibilistic" or probabilistic, and still other models addressed desirable features like composability.

What's wrong with these models? This question could be addressed at several levels. At the policy level, it has been suggested that no one cares about covert channels anymore, therefore models that purport to explain them are uninteresting. This does not really seem to be a valid response. There may be a shift in application areas, however. There is less emphasis in the design of multilevel operating systems, but more interest in something like the Bleichenbacher attack on the PKCS #1 cryptographic protocol standard [1], where a channel that is due partly to the algorithm and partly to the protocol design leads to compromise of encrypted data. Attacks that might expose a stored key are of great concern. The basic principles of information compromise still apply.

There is also the practical question of how non-interference theory can be translated into efficient algorithms for detecting covert channels. Non-interference anal-

Proceedings of the 14th IEEE Computer Security Foundations Workshop (CSFW'01)
1063-6900/01 $10.00 © 2001 IEEE

ysis is often implemented as a laborious verification effort, if it is implemented at all. This is an obstacle, but not a reason for abandoning non-interference research.

At a more technical level, there is concern over the angels-on-a-pinhead syndrome. In most non-interference models, a single bit of compromised information is flagged as a security violation, even if one bit is all that is lost. To be taken seriously, a non-interference violation should imply a more significant loss. Even at a theoretical level where timings are not available, and a bit per millisecond is not distinguishable from a bit per fortnight or a bit per century, a channel that compromises an unbounded amount of information is substantially different from one that cannot. Characterization of unbounded channels is suggested as the kind of goal that would advance the study of this subject, and some creative thought could no doubt suggest others.

A possible approach is to use non-interference as a high level specification of confidentiality and then map this down to a model of the particular architecture and mechanisms of the system in question. This induces constraints on underlying model, for example, on which access modes are allowed. The intuitive and compelling nature of the notion of non-interference makes this attractive but problems arise in ensuring that the system model is faithful, that all channels and access modes have been identified and accurately characterized, and that the mapping is accurate.

These observations seem to suggest that non-interference is little more than a rather intriguing topic of arcane debate, at best the source of compelling theoretical challenges on which learned but largely irrelevant papers can be written. Maybe the time has come to recognize that non-interference is now just a dinosaur of the Multi-Level Security era when intelligence agencies roamed the information security world.

The purpose of this panel is to address the question of what use, if any, non-interference really serves in the design, development and verification of secure systems and architectures? If the conclusion is that it serves no useful purpose, then why has it held such a fascination for the information security community for so long?

References

[1] D. Bleichenbacher. Chosen ciphertext attacks against protocols based on the RSA encryption standard PKCS #1. In *Advances in Cryptology - CRYPTO '98*, volume 1462 of *LNCS*, pages 1–12. Springer, 1998.

[2] P. Y. A. Ryan and S. A. Schneider. Process algebra and non-interference. *JCS special issue on CSFW'99*.