

Panel: A Genealogy of Non-Interference

Peter Y. A. Ryan
Defence Research Agency
Room NX17
St Andrew's Rd
Malvern
Worcs WR14 3PS
ryan@rivers.dra.hmg.gb

1. Introduction

Non-interference is intended to capture the idea that one user's interaction with a system should not causally effect any observations of the system that can be made by another user. Thus the latter user should be unable to infer anything from observations of the system about the actual behaviour of the former. As such it is a central concept in security modelling.

Ever since Goguen and Meseguer first proposed a formalisation of the idea there has been a staggering proliferation of generalisations and alternative formulations. This, superficially very simple and intuitive idea, turns out to harbour remarkable subtleties, especially when extended to systems with non-determinism. It should be noted that the concept of non-determinism itself is very subtle and indeed at least two distinct flavours can be identified: underspecified and probabilistic, and it turns out that the distinction can be crucial from a security point of view, see for example [3]. Further elaborations of non-interference are possible if probability and/or time are introduced into the models.

The utility of NI has also been questioned: it is clearly an elegant abstraction but, at least in its pure form, seems rarely to be used for real. This has prompted a number of generalisations applicable to a wider class of applications:

- conditional
- partial
- intransitive

The current situation is very unsatisfactory: the literature is full of formulations that make varying claims of generality, superiority etc. New versions continue to be proposed (several were submitted to this year's IEEE Security and Privacy conference for example). Many different

frameworks and models of computation are employed making comparison difficult or impossible. The properties possessed by these formulations differ in significant respects: for example some satisfy "hook-up" or compositional properties and others don't. Some possess "input totality," others don't and in some formulations it is a non-issue.

It seems that much of the variation can be understood in terms of different models of computation, different notions of composition, different ways in which the "low-level" view is formalised and different notions of equivalence of processes. Indeed it seems that much of the discussion of ideas of non-interference parallels the discussions in the process algebra community as to definitions of equivalence of processes. For example "non-interference on strategies" seems to mirror the idea of testing equivalence (that in turn is equivalent to failures equivalence). "Forward correctability" appears to mirror Milner's idea of confluence and so on.

In short, the newcomer is faced with a bewildering array of formulations to choose from. He is puzzled as whether or not "non-interference is compositional," whether NI can be refined and so on. He is provided with little or no advice as to which formulation is most appropriate to his application.

2. Steps towards unification

Admirable efforts to bring order to this confusion are: Focardi and Gorrieri, [1], and McLean's selective interleaving approach, [2].

Roscoe's approach, [3], of formulating non-interference in terms of determinism of a suitable abstraction has the advantage of sidestepping the problems of choosing a suitable notion of equivalence, as most definitions of equivalence coincide for deterministic systems.

3. Purpose of the panel

The aim of this panel is to bring some order to the subject. Questions to be considered include:

Which versions of non-interference, including the various generalisations, are valid and useful and, of these, which are appropriate to which applications? For example Roscoe's formulation is clearly very effective (i.e refinements preserve it) but its range of applicability is not entirely clear.

Can we propose a framework in which the various generalisations of non-interference could be formulated in a unified way and in which the origin of the apparently conflicting results can be better understood?

4. Panelists

Roberto Gorrieri, Università di Bologna (Italy)

John McLean, Naval Research Laboratory (US)

A. W. Roscoe, Oxford University (UK)

Peter Y. A. Ryan, Defense Research Agency (UK),
panel organiser.

References

- [1] R. Focardi and R. Gorrieri. A taxonomy of trace-based security properties for CCS. In *Proceedings of the Seventh Computer Security Foundations Workshop*, pages 126–136, 1994.
- [2] J. McLean. A general theory of composition for trace sets closed under selective interleaving functions. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 79–93, 1994.
- [3] A. W. Roscoe. CSP and determinism in security modelling. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 114–127, 1995.