Security by specification and refinement

ETAPS 2010, Paphos, Cyprus

Summary

Since 1995 there has been intensive research in source-level assertion-style reasoning for developing probabilistic systems via stepwise refinement that moves through layers of abstraction [1,2,3]; in 2005 that work was summarised in a comprehensive text [4]. In the five years since then, refinement has been further adapted to deal with non-interference-style security, at first without probability [5,6] — but very recently with probability as well [7].

This tutorial will build on a basic understanding of specification and refinement to explain all those recent developments, and to show how they fit in to research themes that overlap both with secure systems and with quantitative reasoning.

In particular the tutorial will focus on systematic source-level developments of a number of well known probabilistic and security-sensitive algorithms, as examples, as well as explaining and illustrating general principles for doing work of this kind. There are many opportunities of up-to-the-minute topics for further research.

The tutorial will be in two parts, morning and afternoon.

Security by specification and refinement

ETAPS 2010, Paphos, Cyprus

Part 1: A refinement-oriented approach to qualitative non-interference

In this part we will explain how a variation of the standard notion of refinement is able to preserve both functional- and non-interference security properties. (The "classical" notion of refinement is concerned only with functional properties.) Crucial is that the resulting refinement relation be *monotonic* (analogous to *compositionality* for equivalence): that property is essential for hierarchical modelling and verification.

In contrast to other formal treatments of security analysis, our emphasis on refinement induces a focus on *comparative* security relative to the underlying attack model: an implementation which refines its specification cannot be said to be "secure" (or not). Rather we judge whether it is "at least as secure as its specification," i.e. is able to withstand the same attacks (or more).

Topics covered in Part 1 include:

- The underlying mathematical foundations of the secure-refinement order;
- A programming logic based on "ignorance preservation";
- A novel treatment of declassification; and
- An illustration of all of these on a number of well-known benchmark case studies, including *The Dining Cryptographers*, and *The Oblivious-Transfer Protocol,* and finally Multi-Party Computation.

Security by specification and refinement

ETAPS 2010, Paphos, Cyprus

Part 2: **Probabilistic models of secrecy and refinement**

The underlying attack model that determines the notion of compositionality we used in Part 1 gives only *qualitative guarantees*, and perforce assumes single protocol runs: more complicated scenarios where information may be accumulated by an attacker over *repeated runs* cannot be analysed exactly in that style. In Part 2 we address that problem by introducing a new notion of probabilistic secrecy and its refinement order.

Topics covered in Part 2 include:

- The underlying mathematical foundations of a probabilistic secrecy order;
- A probabilistic testing semantics and its relationship to other information-based orders including Shannon entropy;
- A treatment of quantitative declassification as a formalisation of information flow;
- An illustration of the technique on some topical case studies, including the so-called Dining Cryptographers' Café, which is a formal treatment of Chaum's Dining Cryptographers that includes (the formalisation of) repeated runs. We show for example how the method can characterise the relation between the coins' bias and information flow.

Security by specification and refinement

ETAPS 2010, Paphos, Cyprus

Summary of schedule

This is a whole-day tutorial (2 x 3 hours), with the morning session devoted to Part 1 topics, and the afternoon to Part 2 topics. The format of the tutorial will be based around theory together with practical demonstrations of the techniques.

The morning session will consist of lecture material covering the qualitative theory together with example specifications and proof. The first afternoon session will be spent on probabilistic models; for the last session, for those interested, we will present a complete derivation of a novel protocol, the Cryptographers' Café, emphasising the quantitative aspects of security refinement.

Bibliography

(1) Probabilistic predicate transformers. Morgan, McIver, Seidel. ToPLaS 18(3), 1996.

(2) Program development by stepwise refinement. Niklaus Wirth. CACM 14(4), 1971.

(3) *The B Book*. J-R Abrial. Cambridge University Press. 1996.

(4) *Abstraction, Refinement and Proof for Probabilistic Systems*. McIver and Morgan. Springer Monographs in Computer Science. 2005.

(5) The Shadow Knows: Refinement of ignorance in sequential programs. Carroll Morgan. In LNCS 4014 Proc. MPC 2006, Tarmo Uustalu (ed.) 2006.

(6) The Shadow Knows: Refinement of ignorance in sequential programs. Carroll Morgan. In Science of Computer Programming 74(8), 2009.

(7) Security, probability and nearly fair coins in the cryptographers' café. McIver, Meinicke and Morgan. Invited talk in LNCS 5850, Proc. FM09, Cavalcanti, Dams (eds.) 2009.