

A nondeterministic lattice of information

Carroll Morgan

University of New South Wales, NSW 2052 Australia
carrollm@cse.unsw.edu.au *

Abstract. In 1993 Landauer and Redmond [2] defined a “lattice of information,” where a partition over the type of secret’s possible values could express the security resilience of a sequential, deterministic program: values within the same cell of the partition are those that the programs does not allow an attacker to distinguish.

That simple, compelling model provided not only a refinement order for deterministic security (inverse refinement of set-partitions) but, since it is a lattice, allowed the construction of the “least-secure deterministic program more secure than these other deterministic programs”, and its obvious dual. But Landauer treated neither demonic nor probabilistic choice.

Later work of our own, and independently of others, suggested a probabilistic generalisation of Landauer’s lattice [1, 3] — although it turned out that the generalisation is only a partial order, not a lattice [5].

This talk looks between the two structures above: I will combine earlier qualitative ideas [6] with very recent quantitative results [4] in order to explore

- What an appropriate purely demonic lattice of information might be, the “meat in the sandwich” that lies between Landauer’s deterministic, qualitative lattice and our probabilistic partial order.
- The importance of compositionality in determining its structure.
- That it is indeed a lattice, that it generalises [2] and that it is generalised by [1, 3].
- Its operational significance and, of course,
- Thoughts on how it might help with constructing (secure) programs.

References

1. Mário S. Alvim, Kostas Chatzikokolakis, Catuscia Palamidessi, and Geoffrey Smith. Measuring information leakage using generalized gain functions. In *Proc. 25th IEEE Computer Security Foundations Symposium (CSF 2012)*, pages 265–279, June 2012.
2. Jaisook Landauer and Timothy Redmond. A lattice of information. In *Proc. 6th IEEE Computer Security Foundations Workshop (CSFW’93)*, pages 65–70, June 1993.

* I am grateful for the support of the Australian ARC via its grant DP120101413, and of NICTA, which is funded by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre-of-Excellence Program.

3. Annabelle McIver, Larissa Meinicke, and Carroll Morgan. Compositional closure for bayes risk in probabilistic noninterference. In *Automata, Languages and Programming, 37th International Colloquium, ICALP 2010, Bordeaux, France, July 6-10, 2010, Proceedings, Part II*, pages 223–235, 2010.
4. Annabelle McIver, Larissa Meinicke, and Carroll Morgan. Abstract Hidden Markov Models: a monadic account of quantitative information flow. In *Proc. LiCS 2015 (to appear)*, 2015.
5. Annabelle McIver, Carroll Morgan, Larissa Meinicke, Geoffrey Smith, and Barbara Espinoza. Abstract channels, gain functions and the information order. In *FCS 2013 Workshop on Foundations of Computer Security*, 2013. <http://prosecco.gforge.inria.fr/personal/bblanche/fcs13/fcs13proceedings.pdf>.
6. C.C. Morgan. *The Shadow Knows*: Refinement of ignorance in sequential programs. In T. Uustalu, editor, *Math Prog Construction*, volume 4014 of *Springer*, pages 359–78. Springer, 2006.