

Non-Interference for Weak Observers

David Clark Sebastian Hunt Pasquale Malacaria

July 20, 2004

We consider transformational programs, i.e. those which transform inputs into outputs, with two levels of confidentiality: high and low. Inspired by Giacobazzi and Mastroeni [GM04], we consider non-interference for the case that the low confidentiality user has only partial knowledge of low inputs and low outputs. We call such a user a *weak observer*. We first define a form of possibilistic non-interference for weak observers; then, after demonstrating that this is not strong enough, we define a probabilistic version. The basic idea behind our approach can be summed up thus: for a weak observer, even a deterministic program behaves non-deterministically.

Non-interference (NI) was first proposed by Goguen and Messeguer in 1982 [GM82] for deterministic systems. There have subsequently been a number of definitions for NI for non-deterministic systems beginning with Sutherland's definition of *Non-deducibility* (ND) in 1986 [Sut86]. This latter definition still admitted some influence of high level inputs on low ones, a flaw fixed by MucCullough's definition of Generalized NI (GNI) in the following year [McC87]. Our possibilistic definition of NI for weak observers is closely related to GNI, at least in spirit. Gray subsequently extended MucCullough's work to incorporate probabilistic considerations in 1990, defining the notion of P-restrictiveness in [WG90]. Our probabilistic definition of NI for weak observers is closely related to this.

1 NI for Strong Observers

Consider a program that transforms inputs to outputs. We assume a very simple denotational semantics in which the program is modelled by a function $g : H \times L \rightarrow H \times L$, where H, L are finite sets (we do not model non-termination). Given a pair (h, l) of inputs or outputs, we suppose that h is classified high-security and l is classified low-security. We consider observers who can observe l but not h : we are concerned with the question of whether such an observer can learn anything about the value of h by observing l and l' , where $(h', l') = g(h, l)$. When nothing can be learned, g is said to be non-interfering.

Let us say that an observer who can see the exact values of l and l' is *strong*. For a strong observer, NI can be characterised as follows. Let All_X be the relation on set X that creates a single equivalence class $(x \text{ All}_X x')$ for

all $x, x' \in X$) and let Id_X be the identity relation on X . We will drop the subscripts where the intended X is clear from the context. Then g satisfies NI for complete low-observation iff

$$g : \text{All} \times \text{Id} \Rightarrow \text{All} \times \text{Id}$$

where $g : R \Rightarrow S$ means $x R x'$ implies $g(x) S g(x')$.

For simplicity we ignore the high outputs by projecting g onto just the low outputs, giving $f : H \times L \rightarrow L$. NI for complete low-observation becomes:

$$f : \text{All} \times \text{Id} \Rightarrow \text{Id} \tag{1}$$

2 Possibilistic NI for Weak Observers

Now consider an observer who is able to observe l but only partially. For example, l may be an integer in twos-complement representation and the observer may be able to see only the lowest bit of l , thus learning its parity. Just as we can model the power of complete low-observation by Id , so we can model the power of this observer by the relation Par , where $x \text{ Par } x'$ iff x and x' have the same parity. In general, we may suppose that an observer can observe inputs with accuracy R and outputs with accuracy S , where R and S are equivalence relations on L . We will say that such an observer is *weak*. We wish to generalise the definition of NI to account for the weakness of the observer.

2.1 First Attempt

It is tempting to try to generalise (1) as follows:

$$f : \text{All} \times R \Rightarrow S \tag{2}$$

This corresponds to what is called *narrow abstract non-interference* in [GM04]. Unfortunately, as shown in [GM04], this condition can fail to hold, even when it seems clear that the observer is unable to learn anything about h . For example, consider the program

$L := H \text{ XOR } L$

where l and h are boolean. Suppose the observer can observe l precisely on output ($S = \text{Id}$) but is completely unable to observe l on input ($R = \text{All}$). Intuitively, it seems¹ that the observer can learn nothing about h : for any given observation on the output, both True and False are possible values for h as far as this observer can tell. However, this program fails to satisfy (2). As a counterexample, consider that $(\text{True}, \text{False}) \text{ All} \times \text{All} (\text{True}, \text{True})$ but $\text{True XOR False} = \text{True}$ $\not\equiv$ $\text{True XOR True} = \text{False}$. In [GM04], this is called a *deceptive flow*.

¹In fact, this depends on whether the observer has prior knowledge of the *statistics* of l . This is dealt with in section 3.

2.2 Second Attempt

In the above counterexample, note that there is actually *no* variation in the input value of h . The problem with (2) is that it requires *all* variation in inputs which is not initially observable, to remain so: “if it’s invisible to start with, it must be a secret”. Thus, by weakening the power of the observer to observe low-inputs (making R coarser), we actually make it *harder* for a program to satisfy (2). What is needed is a definition of NI which distinguishes between what is secret (and required to remain so) and what is merely unknown to the observer.

First some notation. Let R be an equivalence relation on L and let $X \subseteq L$. We use L/R for the set of equivalence classes of L under R and we implicitly lift f to a map between power sets as follows:

$$f(h, X) \stackrel{\text{def}}{=} \{f(h, x) \mid x \in X\}$$

We also overload the usual notation for equivalence classes as follows:

$$[X]_R \stackrel{\text{def}}{=} \{[x]_R \mid x \in X\}$$

where $[x]_R$ means the R -equivalence class to which x belongs.

We express the NI condition as an invariant. We call this a possibilistic NI condition because it is the set of possible S -equivalence classes on low output that remains invariant under variation of the high inputs. Define $\Psi_X : H \rightarrow L/S$ by

$$\Psi_X(h) \stackrel{\text{def}}{=} [f(h, X)]_S$$

The NI condition is then: for each equivalence class on low inputs, Ψ is invariant on high inputs. That is:

$$\forall X \in L/R, \forall h, h' \in H \cdot \Psi_X(h) = \Psi_X(h') \tag{3}$$

This condition corresponds to a (less general) version of the abstract non-interference property of [GM04].

Checking the XOR program above, it is easily seen that it satisfies (3).

3 Probabilistic NI for Weak Observers

As is well known [WG90], possibilistic non-interference can allow the possibility of probabilistic covert channels. Consider again the XOR program and the same observer ($R = \text{All}, S = \text{Id}$). Suppose now, however, that the observer knows that 90% of the time l is False on input. Then, guessing that the high input is equal to the low output the observer will be right 90% of the time.

We can strengthen the possibilistic NI condition into a probabilistic one. First, some more notation. We lift f again; this time from a map between sets to $f^* : H \times \mathcal{D}(L) \rightarrow \mathcal{D}(L)$, a map between probability distributions (we write

$\mathcal{D}(X)$ to mean the set of all distributions on finite set X). We define f^* in the obvious way:

$$f^*(h, \mu)(l') \stackrel{\text{def}}{=} \sum_{\{l \mid f(h, l) = l'\}} \mu(l)$$

Let $\mu, \eta \in \mathcal{D}(L)$ and let R, S be equivalence relations on L . We overload the previous notations for equivalence classes as follows. Firstly:

$$\mu/R \stackrel{\text{def}}{=} \{\mu|_X \mid X \in L/R, \mu(X) > 0\}$$

where $\mu(X) \stackrel{\text{def}}{=} \sum_{x \in X} \mu(x)$ and $\mu|_X$ is μ specialised to X ($\mu|_X(x) \stackrel{\text{def}}{=} \mu(x)/\mu(X)$). Secondly, $[\eta]_S \in \mathcal{D}(L/S)$ is the induced distribution:

$$[\eta]_S(X) \stackrel{\text{def}}{=} \eta(X)$$

for all $X \in L/S$.

If μ is the distribution on low inputs, the invariant $\Phi_{\mu'} : H \rightarrow \mathcal{D}(L/S)$ is defined, for a given specialised distribution $\mu' \in \mu/R$, as:

$$\Phi_{\mu'}(h) \stackrel{\text{def}}{=} [f^*(h, \mu')]_S$$

The NI condition is then: for each specialised input distribution, Φ is invariant on high inputs. That is:

$$\forall \mu' \in \mu/R, \forall h, h' \in H \cdot \Phi_{\mu'}(h) = \Phi_{\mu'}(h') \quad (4)$$

When μ is the 90% skewed distribution considered earlier, the XOR program fails to satisfy (4). With $R = \text{All}$, there is only one equivalence class and hence only one μ' , which is just μ . But $\Phi_{\mu}(\text{True})$ is 90% skewed towards True whereas $\Phi_{\mu}(\text{False})$ is equally skewed towards False.

On the other hand, when μ is uniform, the same program, with the same observer, *does* satisfy (4): in this case the unknown low input is acting as a uniformly distributed key and the program implements perfect encryption.

References

- [GM82] J. Goguen and J. Meseguer. Security policies and security models. In *IEEE Symposium on Security and Privacy*, pages 11–20. IEEE Computer Society Press, 1982.
- [GM04] R. Giacobazzi and I. Mastroeni. Abstract non-interference: Parameterizing non-interference by abstract interpretation. In *The 31st Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'04)*, pages 186–197, Venice, Italy, January 2004. ACM Press.
- [McC87] D. McCullough. Specifications for multi-level security and a hook-up property. In *Proceedings of the 1987 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, April 1987.

- [Sut86] D. Sutherland. A model of information. In *Proceedings of the 9th National Computer Security Conference*, September 1986.
- [WG90] James W. Gray, III. Probabilistic interference. In *Proc. 1990 IEEE Symposium on Security and Privacy*, pages 170–179, Oakland, CA, 1990.