Comments on
A Logical Approach To Multilevel Security Of Probabilistic Systems
James W.Gray, Paul F. Syverson

Sabina Petride

March 13, 2001

The system consistes of $\mathcal{H}$ (the covert senders, or the high environment), $\mathcal{L}$ (the covert receivers, or the low environment), and Σ (the system that provides services to $\mathcal{H}$ and $\mathcal{L}$). The covert senders may attempt to send secret information via Σ to the receivers, and if this covert information flow is possible then Σ is insecure (it has a covert channel). This information flow is not considered as intentionally: the concern is that the covert senders are actually trojan horses. The worst case scenario is assumed: all covert senders and covert receivers are trojan horses. The information flow is not possible by bypassing Σ, which means that the senders cannot directly communicate with the receivers.

Σ's interface consists of a collection of channels with one input, one output per unit of time; Σ is considered an agent, and all the processing done in the high/low environment is modeled as one agent: $\mathcal{H}/\mathcal{L}$. The set of all channels is $\mathbf{C}$ and the set of low channels is $\mathbf{L}$. $\mathcal{H}$ can see inputs and outputs on all $\mathbf{C}$, but $\mathcal{L}$ can see inputs and outputs only on $\mathbf{L}$. All nonprobabilistic choices are factored out of the execution of the system by fixing an *adversary* at the start of the execution, where the *adversary* is modeled as a conditional probability of the next inputs on the histories of inputs and outputs up to the present.

It is assumed that at every point in time (here *time is discrete*) the input on high channels $\mathbf{C}$-$\mathbf{L}$ is independent of the input on the low channels $\mathbf{L}$

(*Secure Environment Assumption*):

$$\mathcal{A}(a|\alpha,\beta) = \mathcal{H}(a|\mathbf{C}-\mathbf{L}|\alpha,\beta) * \mathcal{L}(a|\mathbf{L}|\alpha|\mathbf{L},\beta|\mathbf{L}).$$

States of the system, computation trees and runs of the system are defined. Based on these, a formal language is given that involves terms, atomic predicates and general predicates, actions, formulas and temporal formulas. In particular, primed state variables are used to give values of variables at the next tick of the clock, subjective probabilities of formulas and knowledge operators are considered.

Axioms are given on temporal reasoning, epistemic reasoning, random variables, plus two additional axioms: if **S** knows something is always true, then **S** always knows it is true, and knowledge of objective probability implies knowledge of subjective probability. A semantic model is built such that soundness is guaranteed.

The reason for constructing this model is that **probabilistic noninterference** is satisfied by Σ iff a certain formula holds, and the presumption is that probabilistic noninterference captures the meaning of security in a system. The formula envolved basically states that it is always the case that $\mathcal{L}$ has knowledge of the probability distribution of its next output (*Formal Security Condition*, **FSC**).

No doubt is cast on the actual relevance of the probabilistic noninterference(**PNI**) to the problem of information flow. By definition, a system Σ satisfies **PNI** if a certain **L**-behavior makes impossible to distinguish between two different **H**-behaviors: the low environment gains no more information from a high environment then from another high environment. The fact that **PNI** captures the meaning of information flow problem can be argued: impossibility of gaining more information from a specific high environment doesn't rule out the possibility that *all* possible high environments allow information flow.

Since deriving **FSC** is difficult, we may verify an easier formula, called *Syntactic Verification Condition*(**SVC**); a system where **SVC** holds satisfies some properties (satisfies the *Applied Flow Model*), which encode a strictly stronger notion of security than **PNI**. Formally, **SVC** states that at all times $\mathcal{L}$ knows the *objective* probability distribution of its next output; **SVC** is easier to verify than **FSC**, and if the system allows, we may check only **SVC** and determine if Σ has covert channels or not.

Some flaws detected:

- **PNI** as enforcing lack of covert channels is questionable (one may argue that the definition of **PNI** is verified, but there is information flow) (!)
- **4** and **Dum** axioms for *temporal reasoning* are not correct
- the two examples are not concludent for the information flow problem.