Available online at www.sciencedirect.com
ScienceDirect
Electronic Notes in Theoretical Computer Science
ELSEVIER
Electronic Notes in Theoretical Computer Science 249 (2009) 75–91
www.elsevier.com/locate/entcs

Quantitative Notions of Leakage for One-try Attacks

Christelle Braun[a] Konstantinos Chatzikokolakis[b]
Catuscia Palamidessi[a]

[a] *INRIA and LIX, École Polytechnique*
Palaiseau, France
{braun,catuscia}@lix.polytechnique.fr

[b] *Technical University of Eindhoven*
Eindhoven, The Netherlands
{kostas}@chatzi.org

Abstract

Recent research in quantitative theories for information-hiding topics, such as Anonymity and Secure Information Flow, tend to converge towards the idea of modeling the system as a noisy channel in the information-theoretic sense. The notion of information leakage, or vulnerability of the system, has been related in some approaches to the concept of mutual information of the channel. A recent work of Smith has shown, however, that if the attack consists in one single try, then the mutual information and other concepts based on Shannon entropy are not suitable, and he has proposed to use Rényi's min-entropy instead. In this paper, we consider and compare two different possibilities of defining the leakage, based on the Bayes risk, a concept related to Rényi min-entropy.

Keywords: Information-hiding, hypothesis testing, probability of error, Rényi min-entropy.

1 Introduction

Information-hiding refers to a large class of problems including Secure Information Flow and Anonymity. There has been a growing interest in developing *quantitative* theories for this class of problems, because it has been recognized that non quantitative (i.e. possibilistic) approaches are in general too coarse, in the sense that they tend to consider as equivalent systems that have very different degrees of protection.

Concepts from Information Theory have revealed quite convenient in this domain. In particular, the notion of noisy channel has been used to model protocols for information-hiding, and the flow of information in programs. The idea is that the input of the channel represents the information to be kept secret, and the output represents the observable. The noise of the channel is generated by the efforts of the protocol to hide the link between the secrets and the observable, often achieved by using randomized mechanisms.

1571-0661
doi:10.1016/j.entcs.2009.07.085

Correspondingly, there have been various attempts to define the degree of leakage by using concepts based on Shannon entropy, notably the mutual information [14,4,7,8] and the related notion of capacity [10,9,2].

In a recent work, however, Smith has shown that the concept of mutual information is not very suitable for modeling the information leakage in the situation in which the adversary attempts to guess the value of the secret in one single try [12]. He shows an example of two programs in which the mutual information is about the same, but the probability of making the right guess, after having observed the output, is much higher in one program than in the other. In a subsequent paper [13], Smith proposes to use a notion based on Rényi *min-entropy*.

We look at the problem from the point of view of the *probability of error*: the probability that an adversary makes the wrong guess. We propose to formalize the notion of leakage as the "difference" between the probability of error *a priori* (before observing the output) and *a posteriori* (using the output to infer the input via the so-called MAP rule). We argue that there are at least two natural ways of defining this difference: one, that we call *multiplicative*, corresponds to Smith's proposal. The other, which we call *additive*, is new. In both cases, we show that it is relatively easy to find the suprema, which is nice in that it allows us to consider the worst case of leakage. The worst case is also interesting because it abstracts from the input distribution, which is usually unknown, or (in the case of anonymity) may depend on the set of users.

2 Preliminaries

2.1 Noisy channels and Hypothesis Testing

In this section we briefly review some basic notions about noisy channels and hypothesis testing that will be used throughout the paper. We refer to [5] for more details.

A *channel* is a tuple $\langle X, Y, p(\cdot|\cdot)\rangle$ where X, Y are random variables representing, respectively, the input with possible values $\mathcal{X} = \{x_1, x_2, \ldots, x_n\}$ (the *secrets* or *hypotheses*) and the output with possible values $\mathcal{Y} = \{y_1, y_2, \ldots, y_m\}$ (the *observables*). The distribution on X, $\vec{\pi} = (\pi_1, \ldots, \pi_n)$ is called *a priori* input distribution. We will also use the notation $p(x_i)$ and $p(y_j)$ to indicate the probabilities of the input x_i (i.e. $p(x_i) = \pi_i$) and the output y_j, respectively. We will denote by $p(y_j|x_i)$ the conditional probability of observing the output y_j when the input is x_i. These conditional probabilities constitute what is called the *channel matrix*, where $p(y_j|x_i)$ is the element at the intersection of the i-th row and j-th column.

The *a posteriori* probability $p(x_i|y_j)$ is the probability that the input is x_i, given that we observe the output y_j. The a priori and the a posteriori probabilities of x_i are related by Bayes theorem:

$$p(x_i|y_j) = \frac{p(y_j|x_i)\, p(x_i)}{p(y_j)}$$