Poster: Channels and the Information Order

Barbara Espinoza *(student)*, Geoffrey Smith *(faculty)*
School of Computing and Information Sciences
Florida International University, Miami, Florida 33199
{bespi009, smithg}@cis.fiu.edu

I. Introduction

Protecting confidential information from improper disclosure is a fundamental security goal, made more challenging due to the practical difficulty of preventing all leakage of secret information. For instance, a login program that rejects an incorrect password unavoidably reveals that the secret password differs from the one that was entered. One promising way to address information leakage is to consider it *quantitatively*, based on the intuition that a login program is acceptable in practice because it leaks only a "small" amount of information about the secret password. This viewpoint has led to the area of *quantitative information flow*, which has seen growing interest in the past decade. (See, for example [1], [2], [3].)

Measures of information flow are based on the information-theoretic notion of *channel* [4]. Channels capture the relationship between the inputs and the outputs of a system through a channel matrix which specifies, for each input, the conditional probability of observing each output of the system. The leakage of a channel is then calculated based on the extent to which observing the channel's output helps an adversary determine the value of the secret input.

Note that the amount of information that a channel leaks depends on the adversary's prior knowledge about the secret input. For instance, if the adversary already knows what the secret input is, then the channel cannot leak any additional information. Moreover, given two channels with the same set of secret inputs, which one is more secure in terms of information flow also varies with the adversary's prior knowledge. Fortunately, it has been shown [3] that channels under the *composition refinement* relation preserve their leakage ordering for *all* contexts. In light of these observations, in this poster we present some current work on the mathematical structure of channels under the composition refinement relation, showing that composition refinement is a partial order up to *semantic equivalence*. As we explain in Section V, channels are semantically equivalent if they are equivalent from the adversary's perspective.

II. Modeling the system as a channel

A *channel* is a triple $(\mathcal{X}, \mathcal{Y}, C)$, where $\mathcal{X}$ is a finite set of secret input values, $\mathcal{Y}$ is a finite set of observable output values, and C is a $|\mathcal{X}| \times |\mathcal{Y}|$ matrix, called the *channel matrix*. The intent is that $C[x,y]$ is the probability of obtaining output y when the input is x. Note that each entry of C is between 0 and 1, and each row sums to 1. An important special case is a *deterministic channel*, where each input yields a unique output.

Following [5], given a prior distribution π on $\mathcal{X}$ we can define the joint distribution p on $\mathcal{X} \times \mathcal{Y}$ as $p_{XY}(x,y) = \pi[x]C[x,y]$. Then, by marginalization we get jointly distributed random variables X and Y with marginal probabilities $p(x) = \sum_y p(x,y)$ and $p(y) = \sum_x p(x,y)$ respectively, and conditional probabilities $p(y|x) = \frac{p(x,y)}{p(x)}$ as well as $p(x|y) = \frac{p(x,y)}{p(y)}$ (provided that the denominators are nonzero). The conditional probabilities $p(x|y)$ can then be grouped into posterior distributions $p_{X|y}$ of the secret for each output y. Hence, assuming an adversary that knows C and π, after observing output y, the knowledge of the adversary about X is updated from π to $p_{X|y}$.

III. Composition refinement

Back in 1993, Landauer and Redmond [6] noticed that we can describe the information leakage of deterministic channels by considering the partition that the channel induces on the set of secret inputs. Each block in the partition contains all the inputs that map to a particular output. For example, if $C_{country}$ is a channel that outputs the country of birth of an individual, then it partitions the set of individuals according to their country of birth. Moreover, partitions are partially ordered by the refinement relation. It is said that partition $\approx$ is refined by partition $\sim$ if each block of $\sim$ is contained within some block of $\approx$. To illustrate this, let C_{city} be a channel that outputs both the country and city of birth of an individual. Hence, since the information provided by the partition induced by C_{city} is finer grained than the information provided by that of $C_{country}$ we say that the partition of C_{city} refines the partition of $C_{country}$.

We remark that partitions under the refinement relation form a complete lattice [6] which, in quantitative information flow, is known as the *Lattice of Information*. Furthermore, given deterministic channels C_1 and C_2, the partition induced by C_2 is finer than the partition induced by C_1 iff C_1 never leaks more than C_2 for any given context [7], [3].

These results do not extend to probabilistic systems, for probabilistic channels do not partition the set of secret inputs. However, we can work around this issue by considering the concept of *cascading*. A cascade of two channels is a classic construction where the output of the first channel is used as input to the second. We write $C_1 = C_2C_3$ to indicate that C_1 is the cascade of channels C_2 and C_3. Then, given channels C_1 and C_2 with the same set of secret inputs, it may be that C_1 is equivalent to C_2 followed by some post-processing; that is, $C_1 = C_2C_3$ for some channel C_3. In this case we say that C_1 is *composition refined* by C_2, denoted by $C_1 \sqsubseteq_\circ C_2$ [3]. An important property of cascading is that, post-processing with the second channel can only destroy information [4], [5], [3], therefore, if $C_1 \sqsubseteq_\circ C_2$, then C_1 never leaks more information than C_2 for any given context. Moreover, in the case of g-leakage [3], such strong leakage ordering implies composition refinement. For the *coriaceous conjecture* of [3] follows from techniques presented in [8]. Hence, composition refinement is the *only* way for the strong g-leakage ordering to hold.

Composition refinement and partition refinement are strongly connected. Given deterministic channels C_1 and C_2, the partition induced by C_2 is finer than the partition induced by C_1 iff $C_1 = C_2C_3$ for some channel C_3. That is, partition refinement and composition refinement coincide for deterministic channels [3]. To illustrate this, consider now a deterministic channel C_{filter} that given the city and country of birth of an individual outputs only the country of birth. Then, channel $C_{country}$ is equivalent to channel C_{city} followed by post-processing with channel C_{filter}, that is, $C_{country} = C_{city}C_{filter}$ and $C_{country} \sqsubseteq_\circ C_{city}$.

Because composition refinement is associated to a strong leakage ordering and, for deterministic systems, coincides with the Lattice of Information, it has been proposed as a candidate for generalizing the Lattice of Information to probabilistic systems [3].

IV. Generalizing the Lattice of Information

We are interested in understanding the extent to which composition refinement generalizes the Lattice of Information to probabilistic systems. With respect to order-theoretic properties, the first thing to remark is that composition refinement is a preorder, that is, it is a reflexive and transitive relation. However, composition refinement is not antisymmetric since there exist channel matrices such that $C_1 \sqsubseteq_\circ C_2$ and $C_2 \sqsubseteq_\circ C_1$ and yet $C_1 \neq C_2$. Hence, composition refinement is not a partial order.

But, because any preorder gives rise to a partial order on the quotient space, composition refinement can still be seen as a partial order. The idea is that, instead of ordering channels directly, we order classes of *composition equivalent* channels. Channels C_1 and C_2 are composition equivalent, denoted $C_1 \equiv_\circ C_2$, if $C_1 \sqsubseteq_\circ C_2$ and $C_2 \sqsubseteq_\circ C_1$. We can then say that composition refinement is a partial order up to composition equivalence. Moreover, in Sections V and VI, we will establish that it is also a partial order up to semantic equivalence of channels.

We now wonder whether channels under composition refinement form a lattice. For deterministic channels, the least upper bound of C_1 and C_2 is the channel that on input x produces as output the pair $(C_1(x), C_2(x))$ [6]. Such channel induces the coarsest partition that is finer than the partitions of both C_1 and C_2. However, this does not extend to the case of probabilistic channels. For two runs of a probabilistic channel carry more information than a single run.

V. Semantic equivalence of channels

Since composition refinement orders classes of composition equivalent channels, it is important to understand what are the structural properties of such classes of channels.

Note that, assuming that the adversary knows C and π, the posterior distributions $p_{X|y}$ and their probabilities $p(y)$ are what C reveals to the adversary about X. Hence, following McIver et al. [8], the leakage semantics of a channel $(\mathcal{X}, \mathcal{Y}, C)$ is a mapping $[\![C]\!] : \mathbb{D}\mathcal{X} \rightarrow \mathbb{D}\mathbb{D}\mathcal{X}$ from prior distributions on the set of secret inputs $\mathcal{X}$ to hyper-distributions, i.e. distributions on posterior distributions. We have found that the leakage of a channel C (under the leakage measures discussed in the literature: mutual information [4], min-entropy leakage [2] or g-leakage [3]) depends only on $[\![C]\!]$.

We say that two channels are semantically equivalent $C_1 \equiv_s C_2$ if they denote the same mapping, that is, if $[\![C_1]\!] = [\![C_2]\!]$. It turns out that when some columns of C are scalar multiples of one another, they can be merged and the result is a semantically equivalent channel. If we also sort the resulting columns lexicographically, we obtain a well-defined *reduced channel* $\{C\}_r$. Channels C and $\{C\}_r$ can be seen as being equivalent from the point of view of the adversary.

We have found that semantic equivalence is simple to check, in fact, $C_1 \equiv_s C_2$ iff $\{C_1\}_r = \{C_2\}_r$. Even more remarkable, based on this property, we have been able to prove that semantic equivalence and composition equivalence are the same relation.

VI. Concluding remarks

Because semantic equivalence and composition equivalence coincide, composition refinement can be seen as a partial order, up to semantic equivalence. That is, composition refinement partially orders semantic denotations of channels, or equivalently, partially orders reduced channels. Hence, reduced channels (or semantic denotations of channels) are to probabilistic channels, as partitions are to deterministic channels.

Acknowledgment

We are grateful to Annabelle McIver, Larissa Meinicke, and Carroll Morgan for their insightful comments. This work was partially supported by the National Science Foundation under grants CNS-0831114 and CNS-1116318.

References

[1] D. Clark, S. Hunt, and P. Malacaria, “Quantitative analysis of the leakage of confidential data,” in *Proc. Workshop on Quantitative Aspects of Programming Languages*, ser. Electr. Notes Theor. Comput. Sci, vol. 59 (3), 2001, pp. 238–251.

[2] G. Smith, “On the foundations of quantitative information flow,” in *Proc. 12th International Conference on Foundations of Software Science and Computational Structures (FoSSaCS ’09)*, ser. Lecture Notes in Computer Science, L. de Alfaro, Ed., vol. 5504, 2009, pp. 288–302.

[3] M. S. Alvim, K. Chatzikokolakis, C. Palamidessi, and G. Smith, “Measuring information leakage using generalized gain functions,” in *Proc. 25th IEEE Computer Security Foundations Symposium (CSF 2012)*, Jun. 2012.

[4] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. John Wiley & Sons, Inc., 2006.

[5] B. Espinoza and G. Smith, “Min-entropy leakage of channels in cascade,” in *Formal Aspects of Security and Trust*, ser. Lecture Notes in Computer Science, G. Barthe, A. Datta, and S. Etalle, Eds. Springer Berlin Heidelberg, 2012, vol. 7140, pp. 70–84.

[6] J. Landauer and T. Redmond, “A lattice of information,” in *Proc. Computer Security Foundations Workshop VI*, Jun. 1993, pp. 65 –70.

[7] P. Malacaria, “Algebraic foundations for information theoretical, probabilistic and guessability measures of information flow,” *CoRR*, vol. abs/1101.3453, 2011.

[8] A. McIver, L. Meinicke, and C. Morgan, “Compositional closure for Bayes risk in probabilistic noninterference,” in *Proc. ICALP’10*, 2010, pp. 223–235.