Robustness

Mário S. Alvim
(msalvim@dcc.ufmg.br)

Quantitative Information Flow

DCC-UFMG
(2021/1)

Introduction

- Given a channel with input X, we have seen that g-leakage provides a rich variety of ways to measure information leakage.

  More specifically:

  1. The prior models the adversary's prior knowledge about X.
  2. The gain function g models the operational scenario, which encompasses both
     a) the set of actions that the adversary can take, and
     b) the worth to the adversary of each such action, for each possible value of X.
  3. The choice of leakage allows us to measure either
     a) the relative increase in g-vulnerability (multiplicative leakage), or
     b) the absolute increase in g-vulnerability (additive leakage).

- Such an extensive vocabulary allows for precise operational significance.

  But it also brings worries about the robustness of leakage assessments.

  That's what we consider here.

The need for robustness

The need for robustness

- When measuring leakage we have consider many relevant questions.
- First: what do we know about the adversary's prior knowledge about X?
  - What actions might she take?
  - How valuable might they be to her?

It's often said that some adversarial behavior is "unlikely in practice"...

But in security this kind of thinking can be dangerous!

Anything that we think is "unlikely in practice" is arguably more likely: adversaries are thinking about what we are thinking, and will exploit it.

The need for robustness

- Another issue is that a channel might be developed with a certain operating context in mind, but it might later be migrated to a different operating context where the original assumptions do not hold.

  This is actually typical of software!

  Hence, if we calculate say $\mathcal{L}_g^{\times}(\pi, C)$ for a certain g and π and C, and decide that it's acceptably small, how sure can we be that C really is safe to deploy?

- Yet another questions is that wrt. the prior π, the adversary might have knowledge of which we are unaware.

  In the context of passwords, for example, large-scale studies have shown that user-selected passwords are not at all uniform.

  Hence analyzing channels with respect to (say) a uniform prior might well give a misleading view of the risks that they pose.

The need for robustness

- **Example 1** Suppose that $\mathcal{X} = \{0, 1, 2, \ldots, 9999\}$.

Let channel A be

$$Y := \texttt{ if } X\texttt{=0 then 1 else } X$$

A	1	2	3	...	9998	9999
0	1	0	0	...	0	0
1	1	0	0	...	0	0
2	0	1	0	...	0	0
3	0	0	1	...	0	0
.	.	.	.	.	.	.
.	.	.	.	.	.	.
.	.	.	.	.	.	.
9998	0	0	0	...	1	0
9999	0	0	0	...	0	1

and let channel B be

$$Y := \texttt{ if } X\texttt{=0 then 0 else 1.}$$

B	0	1
0	1	0
1	0	1
2	0	1
3	0	1
.	.	.
.	.	.
.	.	.
9998	0	1
9999	0	1

Under uniform prior $\boldsymbol{\vartheta}$, from Thm. 5.17 their multiplicative Bayes leakage is:

$$\mathcal{L}_1^{\times}(\boldsymbol{\vartheta}, \mathsf{A}) = 9999 \qquad \text{and} \qquad \mathcal{L}_1^{\times}(\boldsymbol{\vartheta}, \mathsf{B}) = 2 \quad .$$

As a result, we might decide that A leaks much more than B and that replacing A with B would be a good idea.

The need for robustness

- Example 1 (Continued)

But suppose that the adversary somehow knows that the value of X is actually either 0 or 1, each with probability $1/2$, so the prior is

$$\pi^n = (1/2, 1/2, 0, 0, \ldots, 0) \ .$$

Now observe that:

- A's output is always 1, which means that it leaks nothing about X, while
- B reveals the value of X exactly.

Hence we now have

$$\mathcal{L}_1^{\times}(\vartheta, \mathsf{A}) = 1 \qquad \text{and} \qquad \mathcal{L}_1^{\times}(\vartheta, \mathsf{B}) = 2 \quad .$$

Hence replacing A with B might be a decision that we would regret! ◁

The need for robustness

- The question of robustness arises also for the gain function g.

- Example 2 Suppose now that X is a 64-bit unsigned integer.

  Let C be the channel

  $$Y := \texttt{if } (X \texttt{ mod } 8) = 0 \texttt{ then } X \texttt{ else } 1$$

  and channel D be

  $$Y := X \mid \texttt{0x7} \quad ,$$

  where "$\mid$ 0x7" denotes **bitwise or** with 7, so D copies the first 61 bits of X into Y, but "masks out" the last 3 bits.

  Assuming a uniform prior ϑ, we can again use Theorem 5.17 to find that

  $$\mathcal{L}_1^{\times}(\vartheta, \mathsf{C}) \quad = \quad 2^{61}{+}1 \qquad \text{and} \qquad \mathcal{L}_1^{\times}(\vartheta, \mathsf{D}) \quad = \quad 2^{61} \quad .$$

  Hence C and D have almost exactly the same multiplicative Bayes leakage.

  But are they (almost) equally secure?

The need for robustness

- Example 2 (Continued)

It may appear that channel D is less secure, since:

- D always leaks the first 61 bits of X, whereas
- C usually leaks almost nothing: whenever X is not a multiple of 8, channel C outputs 1, which reveals only that X is not a multiple of 8 — meaning that $7/8 \cdot 2^{64}$ values remain possible.

Indeed if we evaluate leakage using the three-tries gain function V_3 (instead of the "one-try" V_1) we find that:

- the posterior vulnerability for D increases from $1/8$ to $3/8$, while
- the posterior vulnerability for C is hardly increased at all:
  - when $Y \neq 1$, the adversary already knows the value of X and does not benefit from three guesses, while
  - when $Y = 1$, she just gets "three stabs in the dark" among the $7/8 \cdot 2^{64}$ possible values, giving a negligible increase in posterior vulnerability.

Hence the three-tries gain function makes D leak more than C.

The need for robustness

- Example 2 (Continued)

But recall the gain function g_{tiger}, in which $\mathcal{W} = \mathcal{X} \cup \{\bot\}$, where the special action $\bot$ is used to opt not to input a guess, and

$$g_{\text{tiger}}(w, x) := \begin{cases} 1 & \text{if } w = x \\ 0 & \text{if } w = \bot \\ -1 & \text{otherwise} \end{cases} .$$

Under this gain function the situation is reversed.

- With channel D, the posterior vulnerability is 0, since a $1/8$ probability of guessing correctly is not high enough to overcome the risk of being eaten by tigers.
- With channel C, the posterior vulnerability is greater than 0: whenever $Y \neq 1$, the adversary knows the value of X, making safe for her to guess then.

Hence g_{tiger} makes channel C leak more than channel D. ◁

Approaches to robustness

Approaches to robustness

- Our examples show that analyzing a channel wrt. a particular prior/gain function is not enough to understand its information leakage in general.
- There has been considerable interest in achieving robustness in QIF.
- Several fruitful approaches have been developed, including:
  - capacity,
  - comparison of channels,
  - collateral effects of channels.

  We'll briefly introduce them now, and they'll be thoroughly cover ahead in this course.

Approaches to robustness: Capacity

- The first approach to robustness is capacity, which is the maximum leakage (of either kind) over all priors π and/or over all gain functions g.

  Capacity analysis enables us to conclude that a channel's leakage is no worse than some amount, over some range of operational scenarios.

  In this course we will consider a total of six capacity scenarios, based on:

  - whether we maximize over just the prior, over just the gain function, or over both the prior and the gain function; and on
  - whether we measure leakage multiplicatively or additively.

Approaches to robustness: Comparison of channels

- A second approach to robustness concerns the comparison of channels.

  This comparison aims at showing that one channel never leaks more than another, regardless of the operational scenario.

  We find that there is a structural refinement order ($\sqsubseteq$) that coincides with this strong leakage ordering.

  Moreover ($\sqsubseteq$) turns out to be a partial order on abstract channels.

  We'll study the theory of refinement after a preparatory discussion of channel compositions.

Approaches to robustness: Dalenius leakage

- A third perspective on robustness considers that a channel that leaks information about a secret X may have the surprising "collateral" effect of leaking information about a different secret Z.

  That can happen because X and Z may turn out to be correlated.

  Such correlations might be discovered at any time...

  Example 3 A medical research may determine a correlation between diet and susceptibility to disease.

  In this case, learning someone's favorite dish might be relatively harmless today, it might not be so in the future! ◁

  We refer to such leakage as Dalenius leakage.

  We find that it is possible to prove upper bounds on the Dalenius leakage of a channel, regardless of any correlations that might be discovered later.