QQIF: Quantum Quantitative Information Flow

Arthur Américo Pasquale Malacaria

School of Electronic Engineering and Computer Science, Queen Mary University of London

Problem Statement

Quantitative Information Flow (QIF) is an area of research that aims to *quantify* how much confidential information systems leak, and to reduce this leakage. One of the most successful frameworks in the field is the g-leakage framework [1], which measures the information systems leak by assigning to it a quantity called *g-vulnerability*, which is predicated on the adversary's knowledge about the sensitive information and how much he expects to *gain* from this knowledge.

In this work, published in [2], we extend the g-vulnerability framework to a quantum setting. We consider quantum systems that have classical secrets, and we also adapt the quantum Blackwell-Sherman-Stein (BSS) theorem [3] — which, in its classical version, is a fundamental result for QIF — to our framework.

Classical QIF and g-vulnerabilities

Model: a *secret* is a random variable (r.v.) X taking values on a finite, nonempty set $\mathcal{X} = \{x_1, \dots, x_n\}$, according to some probability distribution p_X, which the adversary is aware of. A *system* takes as input the secret X, and produces an *observable* Y, a r.v. taking values on $\mathcal{Y} = \{y_1, \dots, y_m\}$. The system is modelled as a *channel* K , which is a matrix that, for each $x \in \mathcal{X}, y \in \mathcal{Y}$, gives the conditional probability $K(y|x)$ of $Y = y$ given that $X = x$. With the realisation of the observable $Y = y$, an adversary updates the knowledge he has about the secret from the initial distribution p_X to $p_{X|y}$ by Bayesian updating $p_{X|y}(x) = \frac{p(x)K(y|x)}{\sum_x p(x)K(y|x)}$.

K	y_1	y_2	y_3	y_4
x_1	$1/2$	$1/3$	$1/6$	0
x_2	$1/4$	$1/4$	$1/4$	$1/4$
x_3	$1/2$	0	$1/2$	0

$\longrightarrow$

	$p_{X\|y_1}$	$p_{X\|y_2}$	$p_{X\|y_1}$	$p_{X\|y_1}$
x_1	$3/5$	$2/3$	$1/3$	0
x_2	$1/5$	$1/3$	$1/3$	1
x_3	$1/5$	0	$1/3$	0

Figure 1. A channel K, and posterior distributions obtained from K and $p_X = (1/2, 1/3, 1/6)$

g-Vulnerabilities: A *gain function* is a function $g : \mathcal{W} \times \mathcal{X} \to \mathbb{R}$ such that $g(w, x)$ is the value of the gain the adversary when he chooses action $w \in \mathcal{W}$ and the secret value is $x \in \mathcal{X}$. The g-vulnerability of the secret before the execution of the system is given by the expected gain of the adversary if he chooses the optimal action

$$V_g(X) = \max_w \sum_x p_X(x) g(w, x). \tag{1}$$

Similarly, the *posterior* g-vulnerability is given by the expected value of the g-vulnerability after the execution of the system

$$V_g(p_X, K) = \sum_y p_Y(y) V_g(p_{X|y}) = \sum_y \max_w \sum_x p_X(x) K(y|x) g(w, x). \tag{2}$$

The quantity of information leakage can then be defined as the increase in g-vulnerability by the execution of the system.

The Blackwell-Sherman-Stein Theorem

The choice of gain function g often reflects the abilities and interests of the adversary. For example, $g(w, x) = \delta_{w,x}$ models an adversary interested in guessing the secret exactly in one try, whereas $g(w, x) = d(w, x)$ for some suitable distance function might represents an adversary aiming to obtain an approximation of the secret.

- This raises the question: when can we guarantee a system is more secure than another for all adversaries?

In [4], McIver et al answered this question by proving an important theorem for QIF, which was later discovered to be equivalent to the BSS Theorem [5].

The Blackwell-Sherman-Stein Theorem: Let $K_1 : \mathcal{X} \to \mathcal{Y}$ and $K_2 : \mathcal{X} \to \mathcal{Z}$ be channels. We have that $\forall p_X, \forall g \; V_g(p_X, K_1) \geq V_g(p_X, K_2)$ if, and only if, there is a channel $R : \mathcal{Y} \to \mathcal{Z}$ such that

$$\forall x, z \quad K_2(z|x) = \sum_y K_1(y|x) R(z|y).$$

That is, a channel K_2 leaks at most as much information as channel K_1 for all gain functions g if, and only if, K_2 can be obtained by *postprocessing* the outputs of K_1 by another channel R.

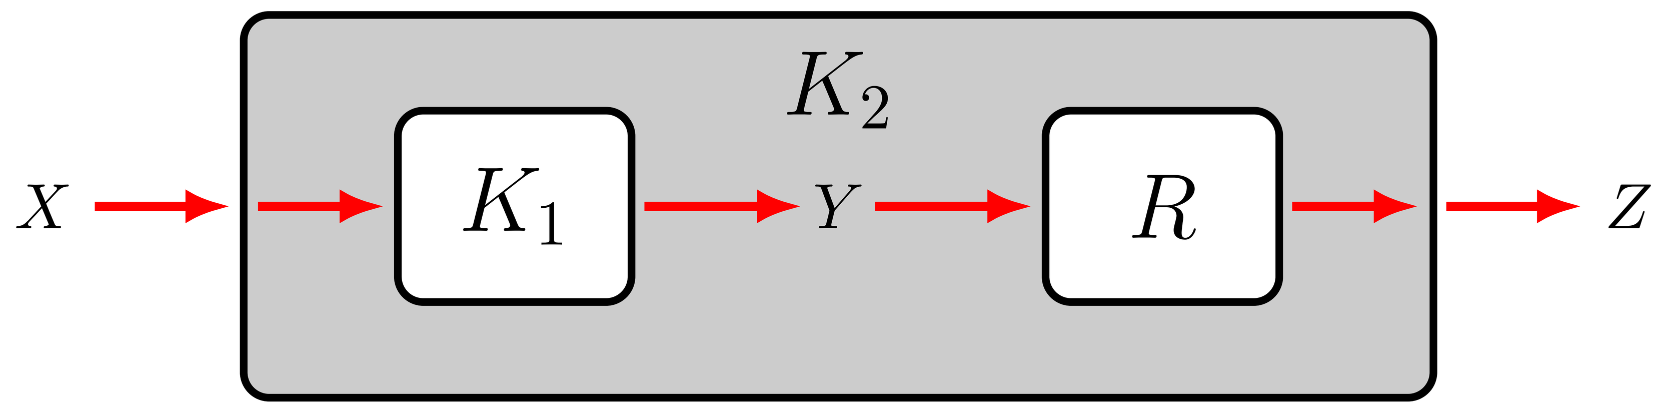

Quantum QIF

The g-leakage framework provides us with a natural way of extending QIF to a quantum setting. In this which the secret is still classical, modelled by a r.v. X taking values on a set $\mathcal{X} = \{x_1, \dots, x_n\}$. However here the system takes a secret value $x \in \mathcal{X}$ as input and performs a computation, producing a quantum state ρ^x. Thus, a system can be represented as a collection of states $\rho^{\mathcal{X}} = \{\rho^x\}_{x \in \mathcal{X}}$ indexed by $\mathcal{X}$, that are density operators on some Hilbert space $\mathscr{H}$.

An adversary then makes a measurement on ρ^x, selecting a POVM $E = \{E_y\}_{y \in \mathcal{Y}}$ from a set of "allowed" POVMs $\mathcal{P}$.. Notice that each POVM is indexed by a (finite, nonempty) set $\mathcal{Y} = \{y_1, \dots, y_m\}$, which is akin to the *output set* in classical QIF.

This construction is similar to *Quantum Statistical Models* in [3], but in this work we limit the set of feasible POVMs, as a way to modelling possible attackers.

Quantifying Information in QQIF

The quantification of information in QQIF is similar to the classical case. The adversary again has some prior knowledge about the secret, modelled by a probability distribution p_X, and a set of possible actions $\mathcal{W}$. The prior g-vulnerability in the quantum case is then the same as in the classical case, i.e. (1).

After the execution of the system, the attacker chooses a POVM $\{E_y\}_{y \in \mathcal{Y}}$ to perform a measurement on the resulting quantum state, and then chooses the action $w \in \mathcal{W}$ that maximises his gain. The *quantum posterior g-vulnerability* is thus

$$V_{g,\mathcal{P}}(p_X, \rho^{\mathcal{X}}) = \max_{E \in \mathcal{P}} \sum_{y \in \mathcal{Y}} \max_{w \in \mathcal{W}} \sum_{x \in \mathcal{X}} p(x) g(w, x) \mathbf{tr}(\rho^x E_y). \tag{3}$$

Notice that one can easily recover the classical case, in which a system is modelled by a channel $K : \mathcal{X} \to \mathcal{Y}$, from the quantum setting. This can be done by letting $\{|y\rangle\}_{y \in \mathcal{Y}}$ be an orthonormal basis of $\mathscr{H}$, defining the quantum states as $\rho_K^x = \sum_y K(y|x) |y\rangle \langle y|$, and letting the set of allowed POVMs to be the singleton $\mathcal{P} = \{E\}$, where $E_y = |y\rangle \langle y|$. In this case, (3) reduces to (2).

The Quantum Blackwell-Sherman-Stein Theorem for QQIF

A QSM is a triple $\mathbf{R} = (\mathcal{X}, \mathscr{H}, \rho^{\mathcal{X}})$, where $\mathscr{H}$ is a Hilbert space and $\mathcal{X}, \rho^{\mathcal{X}}$ are a collection of states $\rho^{\mathcal{X}} = \{\rho^x\}_{x \in \mathcal{X}}$ indexed by $\mathcal{X}$ in $\mathscr{H}$. Given a QSM $\mathbf{R}$, an action set $\mathcal{W}$ and a gain function g, we define the *maximum expected payoff* as

$$\$_g(\mathbf{R}) = \max_E \frac{1}{|\mathcal{X}|} \sum_{x \in \mathcal{X}} \sum_{w \in \mathcal{W}} g(w, x) \mathbf{tr}(\rho^x E_w),$$

the maximum being taken over all possible POVMs indexed by elements in $\mathcal{W}$.

As in the classical case, there is a strict connection between the maximum expected payoff and the posterior g-vulnerability.

Proposition: Let $\mathbf{R} = (\mathcal{X}, \mathscr{H}, \rho^{\mathcal{X}})$ be a QSM, $\mathcal{W}$ an action set and g a gain function. Let p_u be the uniform distribution, and $\mathcal{P}$ be all POVMs in $\mathscr{H}$. Then,

$$\$_g(\mathbf{R}) = V_{g,\mathcal{P}}(p_u, \rho^{\mathcal{X}}).$$

In [3], Buscemi proved a quantum version of the BSS Theorem. The role that postprocessing plays in the classical version is performed by *statistical morphisms*, which are linear maps that include completely positive trace-preserving maps.

Definition: Let $\mathcal{G}(\mathscr{H})$ be the set of density operators in $\mathscr{H}$, and $\mathcal{L}(\mathscr{H})$ the set of linear operators in $\mathscr{H}$.. A family $\{F_w\}_{w \in \mathcal{W}}$ of operators over H is called a $\mathcal{W}$-test on a subset $\mathcal{G} \subset \mathcal{G}(\mathscr{H})$ if there is a POVM $E = \{E_w\}_{w \in \mathcal{W}}$ indexed by $\mathcal{W}$ such that for all $w \in \mathcal{W}, \rho \in \mathcal{G}$, we have $\mathbf{tr}(\rho F_w) = \mathbf{tr}(\rho E_w)$.

Definition: Let $\mathcal{G} \subset \mathcal{G}(\mathscr{H})$, $\mathcal{G}' \subset \mathcal{G}(\mathscr{H}')$. A linear map $L : \mathcal{L}(\mathscr{H}) \to \mathcal{L}(\mathscr{H}')$ induces a statistical morphism $L : \mathcal{G} \to \mathcal{G}'$ if 1) for all $\rho \in \mathcal{G}$, $L(\rho) \in \mathcal{G}'$, and 2) the dual transformation $L^* : \mathcal{L}(\mathscr{H}') \to \mathcal{L}(\mathscr{H})$ defined by trace duality maps $\mathcal{W}$-tests on $\mathcal{G}'$ to $\mathcal{W}$-tests in $\mathcal{G}$.

Given a collection of states $\rho^{\mathcal{X}}$, let $\mathcal{G}(\rho^{\mathcal{X}}) = \{\rho^x \mid x \in \mathcal{X}\}$. The proposition above allows us to give Buscemi's results in terms of the QQIF Framework:

The Quantum Blackwell-Sherman-Stein Theorem [3]: Let $\mathcal{P}$ be the set of all possible POVMs. Then, there is a statistical morphism $L : \mathcal{G}(\rho^{\mathcal{X}}) \to \mathcal{G}(\sigma^{\mathcal{X}})$ such that $\forall x \in \mathcal{X}$, $L(\rho^x) = \sigma^x$ if, and only if, for all gain functions g and all p_X,

$$V_{g,\mathcal{P}}(p_X, \rho^{\mathcal{X}}) \geq V_{g,\mathcal{P}}(p_X, \sigma^{\mathcal{X}}).$$

References

[1] M. S. Alvim, K. Chatzikokolakis, C. Palamidessi, and G. Smith, "Measuring information leakage using generalized gain functions," in *Proc. IEEE 25th Computer Security Foundations Symposium (CSF)*, pp. 265–279, 2012.

[2] A. Américo and P. Malacaria, "Qqif: Quantum quantitative information flow (invited paper)," in *2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, pp. 261–270, 2020.

[3] F. Buscemi, "Comparison of quantum statistical models: Equivalent conditions for sufficiency," *Communications in Mathematical Physics*, vol. 310, no. 3, pp. 625–647, 2012.

[4] A. McIver, C. Morgan, G. Smith, B. Espinoza, and L. Meinicke, "Abstract channels and their robust information-leakage ordering," in *Proc. of POST*, vol. 8414 of *LNCS*, pp. 83–102, Springer, 2014.

[5] D. Blackwell, "Equivalent comparisons of experiments," *The Annals of Mathematical Statistics*, vol. 24, no. 2, pp. 265–272, 1953.