

Privacy Amplification in the Isolated Qubits Model

Yi-Kai Liu


National Institute of Standards and Technology (NIST)
Gaithersburg, MD, USA

Joint Center for Quantum Information and Computer Science (QuICS)
NIST / University of Maryland




Tamper-Resistant Hardware

- One-time programs [GKR '08]



One-Time Programs

- Non-interactive
- Need trusted hardware
- Previous work: assume one-time memories
 - Abstract functionality, like oblivious transfer
 - [GKR '08] [Goyal et al] [Bellare et al]
- This work: assume isolated qubits
 - A special class of quantum-mechanical devices, with a natural restriction on the power of the adversary
 - [Liu '14]



Privacy Amplification

- Needed because real devices are never perfect
 - Some information always leaks (e.g., via side channels)
- Usual recipe: use a hash function h
 - Randomly chosen from a 2-universal family
- This doesn't work for us!
 - One-time programs are non-interactive
 - Need to announce h at the beginning of the protocol
 - Adversary knows h before he attacks the scheme



This Talk

- Deterministic privacy amplification
 - Secure in isolated qubits model
 - Non-interactive: uses a single fixed hash function, is secure against all adversaries simultaneously
- => One-time memories using isolated qubits
 - Only leak an exponentially small amount of information


Quantum Mechanics

- Limits the power of an adversary
 - No-cloning theorem
 - Measurement disturbs the quantum state
- However...
 - Adversary can do entangled measurements on many qubits at once
 - Quantum bit-commitment, oblivious transfer are impossible

Isolated Qubits Model

Real-world examples:
solid-state nuclear spins,
Si defects, NV centers

- Assume adversary cannot do entangling operations
 - Can only do adaptive single-qubit operations
 - “LOCC” = local operations and classical communication



Related Work

- “Nonlocality without entanglement” [Bennett et al, 1999]
 - There exist quantum operations that are “one-way,” in a world where everyone is restricted to LOCC operations
- Quantum bit-commitment secure against k -local adversaries [Salvail, 1998]
- Quantum bounded storage model [Damgaard et al, 2005]
- Quantum tokens [Pastawski et al, 2012]
- Password-based identification [Bouman et al, 2012]

Main Result

- Deterministic privacy amplification for one-time memories in the isolated qubits model
 - • Given a leaky string-OTM
 - Construct a bit-OTM with exponentially-small leakage
- Combine with construction of leaky string-OTM using isolated qubits [Liu, CRYPTO 2014]

Deterministic Privacy Amplification


- Use two r-wise independent hash functions

$$F, G: \{0,1\}^{\ell} \rightarrow \{0,1\}$$

$$r = \text{poly}(k)$$


Programming the OTM

Reading the OTM



Proof of Security


- **Key idea:**
- Don't analyze LOCC adversaries (decision trees)
- Instead look at **POVM elements** (individual paths through decision trees)
 - Tensor products of **single-qubit** operators
 - Simpler: can imagine that qubits are measured in any order
 - Fewer in number: $2^{O(n)}$ rather than $2^{2^O(n)}$



POVM element
 $M = M_1 \times M_2 \times \dots \times M_n$

Proof of Security

- **Key idea:**
- Show that **every POVM element** (every path through every decision tree) is bad for the adversary
 - Adversary's adaptive choices = clever ways of grouping together POVM elements
 - But this doesn't help!



Definition of Security: Leaky String-OTM

- Store two strings S and T , each ℓ bits long
 - Assume S, T are uniformly distributed
 - Ideal security goal: adversary can learn either S or T , but not both
- “Leaky” security:
 - For any LOCC adversary, have uncertainty about (S, T)
- $H_{\infty}^{\varepsilon}(S, T|Z) \geq (0.5 - \delta) \ell$
 - Z = adversary’s measurement outcome
 - $\varepsilon \leq \exp(-\Omega(k))$

Definition of Security: Single-Bit OTM

- Store two bits A_0 and A_1
- Every LOCC adversary learns at most one of A_0, A_1
 - There exists a binary random variable C , such that adversary doesn't learn A_C (even if he learns A_{1-C})
- $\Delta((A_C, A_{1-C}, C, Z), (U, A_{1-C}, C, Z)) \leq \epsilon$
 - Δ = statistical distance, $\epsilon \leq \exp(-\Omega(k))$
 - Z = adversary's measurement outcome
 - U = independent uniformly random bit

“Classical”
security
definition

Definition of Security: Single-Bit OTM

- **NB: our definition of security is mostly classical**
 - **Justification:** isolated qubits can't become entangled with anything else
 - **Caveat:** security claim only applies after the adversary measures the qubits
 - **Question about composability:** what if the adversary defers some measurements until later?

Proof of Security


- Recall: hash functions F, G
 - Input bits (a_0, a_1) expanded to strings (s, t) such that $F(s)=a_0, G(t)=a_1$
 - Let M be some measurement outcome (POVM element) that the adversary can observe
- First, prove security conditioned on a fixed M
 - **For any fixed M ,** with high probability over F and G , the scheme is secure
 - $E_{ST} [(-1)^{A_0+A_1} \mid \text{Adv. gets outcome } M]$
 $= \sum_{st} (-1)^{F(s)+G(t)} \Pr[S=s, T=t \mid \text{Adv. gets outcome } M]$

Proof of Security

- First, prove security conditioned on a fixed M
 - **For any fixed M** , with high probability over F and G , the scheme is secure
 - $E_{ST} [(-1)^{A^0+A^1} \mid \text{Adv. gets outcome } M]$
 $= \sum_{st} (-1)^{F(s)+G(t)} \Pr[S=s, T=t \mid \text{Adv. gets outcome } M]$
 - **Use large-deviation bounds** for sums and quadratic functions of (r -wise) independent rv's (Hoeffding, Hanson-Wright)
 - **Security property of leaky OTM** \Rightarrow distribution of (S, T) has high entropy \Rightarrow variance is small

Proof of Security

- Covering argument
 - ϵ -net for the set of all (tensor product) POVM elements
 - This has cardinality $\leq 2^{\text{poly}(k)}$
- **Union bound** over all points M in the net
- “Continuity argument”: perturbation of M does not affect security much
- So with high probability over F and G ,
for all M (simultaneously),
the scheme is secure



Outlook



- Experimental implementations?
 - Fault tolerance?
 - Adversaries who can perform noisy entangling gates?
- Composable security?
 - One-time programs? Other protocols?
 - Delayed measurements?
- Leakage resilience using quantum resources?