From: Emanuel Knill
To: Bierhorst, Peter L. (Assoc); Glancy, Scott C. (Fed); Liu, Yi-Kai (Fed); Jordan, Stephen P (Fed)
Subject: an accessible followup to my notes on min-entropy estimation
Date: Monday, August 8, 2016 1:37:39 PM

Hi all,

Thanks to Peter, I now have a reasonably digestible instance of the results in my notes from last year, see below. With this, I attempted an independent check of the Peter's soundness bounds. The simplest closed form comparison I can come up with favors Peter's bound and agrees to lowest order in the relevant quantities. It is based on a test factor strategy that is adaptive and has some interesting optimization opportunities to explore, which complicates a general comparison. At a glance, the optimized gain rate may be better, but with worse behavior with respect to the significance parameter. It is basically the arguments from my notes of last year and the reduction to probability estimation I mentioned at the end of our last workshop and that we have discussed since then. Everything simplifies in the relevant case where E's classical state is fixed at the beginning of the protocol. If correct, these results would be a good basis of a follow-up paper that deals with the issue of adaptation and the general convex operational ("super-quantum") setting for randomness expansion.

Remark: The simplification is possible because the standard interpretation of "untrusted devices" is translated into a complete absence of physical coupling between E and the devices (and other related constraints on couplings). This is expressed operationally by describing the rounds in terms of quantum operations acting independently on specific subsystems. This is what I learned with Peter's help. I am sorry, I should have realized this before, it is pretty obvious in retrospect, just from various diagrams in Colbeck's thesis. I didn't check whether anyone weakened these operational assumptions, but as pointed out to me by Peter, one of the latest papers from Renner's group uses the same interpretation (and has similar diagrams). In my notes, besides the known complication of considering general, continuous stochastic processes, I didn't constrain communication operationally, instead I explicitly defined what information cannot leak by conditional independence constraints on stochastic processes. This can be thought of as coming from the ``other" side, favoring generality: Take the most general situation with or without couplings, and pin down exactly how they should be minimally constrained to ensure the soundness property we need. In my opinion, this is more realistic (there is no such thing as isolated labs in reality, no control by E is implied), and it can be used to target one's physical barriers to exactly what information should not leak. It also becomes relevant when composing multiple protocols while making use of protocol outputs. Much of the difficulty of proving the results in my notes came from having to keep track of these stochastic processes. For the single-use protocol (no composition) and the randomness expansion application using our techniques, it now looks like one can reduce the general situation to the operationally-constrained one, so for now, this appears to be w.o.l.g. (as far as our soundness results are concerned), maybe even in the convex operational formalism for physical systems.

Here is an analysis leading to soundness. it is based on probability estimation by test martingales:

We are considering the sequence of trials with settings $Z_i = X_i Y_i$ and outcomes $C_i = A_i B_i$. Write $U_i=X_i Y_i A_i B_i $. (Juxtaposition denotes concatenation.) The system E with respect to which we are to produce private randomness is assumed to be classical and fixed initially. We condition everything on a specific value of E, making sure that nothing in the proof depends on the particular value e of E. Conditional on this value, we can then just consider $(U_i)_i$ without explicit mention of e. I write $U_{\leq k}=(U_i)_{i=1}^k$ and similarly for the other stochastic sequences in play. The usual fundamental assumptions are that Z_i is uniform conditional on $U_{\leq i-1 }$ and $Z_i C_i$ satisfies the non-signaling conditions conditional on $U_{\le i-1}$. Let T be a real-valued function of U, the reference Bell function. We assume that $\lang T\rang_{LR} \leq 1$ and $\lang T\rang_{PR} \leq 1+m$ for all local realistic distribution LR and all PR boxes PR. Equivalently, $1+m$ is the maximum expectation of T over all non-signaling distributions. For this treatment, T does not need to be positive, only the functions to be derived from T must be positive. But assuming that it is positive seems to be w.o.l.g. in view of extra parameters we can choose.

The main idea is to give a bound $\Prob(c_{\le n}|z_{\leq n})\leq p_{max}$ at a significance level of ϵ. Here ϵ is the soundness parameter after n trials. This is a purely statistical estimation problem. It requires finding a statistic P_{max} such that $\Prob(\Prob(C_{\le n}|Z_{\leq n})\leq P_{max})\ge 1-\epsilon$. Note that the inner $\Prob$ defines a random variable (a conditional expectation), and the outer probability is determined by the joint distributions of all random variables in the argument.

To turns this into a standard sound protocol suitable for randomness extraction, let σ be the requested number of random bits at significance level ϵ. The protocol fails if $p_{max}>2^{-\sigma}$ and succeeds otherwise. This protocol is then sound according to our definitions. It is also true that a sound protocol for min-entropy outputs can be turned into a probability estimation protocol, but we do not need this here. (I have some notes explaining this that I sent out a while ago.)

The desired bound p_{max} is obtained by first relating the excess over 1 of the actual expectation $\lang T\rang$ to the maximum probability of measurement outcomes given the settings, on a given trial. (If indeces are omitted, it is a generic trial, with probability distributions conditioned on the past, which we generally define to be the previous trials' settings and outcomes, as well as, implicitly, e) Let q be the minimum probability of PR boxes in a representation of the distribution of U in terms of a mixture of local realistic distributions and PR boxes [See Peter's earlier paper on this decomposition]. Then the maximum probability of an outcome given a setting is bounded by $1-q/2$, where the maximum can be achieved if the LR part is deterministic and equal to one of the outputs of the (single) PR box in the distribution. The constraints on T imply that $q/2 \ge (\lang T\rang-1)/(2m)$. [This inequality is also explained in Peter's proof with a bit more detail. For current purposes, we don't need to know that only one, distribution dependent, PR box is needed in the decomposition, we just need that the non-signaling polytope is the convex closure of PR boxes and LR distributions.]

Define Q_i as the random variable which returns the value of q (as defined in the previous paragraph) for the i'th trial conditional on $U_{\le i-1}$. Before the i'th trial, the probability distribution for the next given $U_{\leq i-1}$ is determined, thus so is the value of q, making Q_i a random variable that is ``determined'' by $U_{\le i-1}$ in the technical sense of stochastic sequences. We use $\prod_{i=1}^n (1-Q_i/2)$ as an upper bound on the desired probability p_{max} (see below). Equivalently, we can estimate $\prod_{i=1}^n (1-Q_i/2)^{-\alpha}$ for any given $\alpha>0$. We do not know this product, but since it comes from a determined process we can estimate it via test-martingale methods. [One of the main differences to Peter's proof is that here we focus on inverse probabilities for estimation purposes and formulate trial-wise linear bounds for these. This bypasses the applications of the arithmetic-geometric mean inequality and enables direct use of the conventional test-martingale inequalities. But it introduces extra parameters to optimize.] The following inequality is useful for this purpose: For $0\leq q,q_0<1$, $\alpha>0$,

$$
(1-q/2)^{-\alpha} \ge (1-q_0/2)^{-\alpha}+\alpha(q-q_0)(1-q_0/2)^{-\alpha-1}/2.
$$

The expression on the right-hand side is the linear approximation to the function $(1-q/2)^{\alpha}$ of q at $q=q_0$. The inequality follows because the function is convex. Define $b_{\alpha,q_0}(q)$ to be the right-hand side of the inequality, a linear, monotone increasing, function of q. Let $T_{\alpha,q_0}=b_{\alpha,q_0}((T-1)/m)$. We now constrain α and q_0 such that T_{α,q_0} is a positive random variable. In general, we should optimize the choice of α and q_0 for the best bounds. It turns out that α has to be chosen initially and kept constant over the trials. (Not doing so requires using a randomness extractor that can use ``weighted'' min-entropy estimates.) But both T and q_0 can be chosen adaptively. Heuristically, q_0 should be chosen so that the inequality above is tight at the mean of T. That is, $q_0 = (\lang T\rang -1)/m$. For now, to simplify expressions, we set $q_0=0$, which is the default if the second subscript in expressions for b and T is omitted.

With the above, the following inequality holds:

$$
\lang T_\alpha \rang \leq (1-Q/2)^{-\alpha},
$$

where the inequality holds conditionally on the past, as usual. The fact that Q is determined matters here. This implies that $T_{i,\alpha}/(1-Q_i/2)^{-\alpha}$ is a test factor, from which we deduce

\[
\Prob(\epsilon \prod_{i=1}^n T_{i,\alpha} \geq \prod_{i=1}^n (1-Q_i/2)^{-\alpha})\leq \epsilon,
\]

Let $V_\alpha =\prod_{i=1}^n T_{i,\alpha}$. Putting everything together, let $P_{max,\epsilon}=(V_\alpha \epsilon)^{-1/\alpha}$ and we find, after some manipulation of inequalities,

\[
\Prob(\Prob(C_{\le n}|Z_{\leq n})\geq P_{max,\epsilon})\leq \epsilon.
\]

The probability here is calculated with respect to the random variables shown, so $\Prob(.|.)$ has to be interpreted as a conditional expectation--a random variable determined by its arguments. In a bit more detail, $\Prob(C_{\le n}|Z_{\leq n})\leq\prod_{i=1}^n(1-Q_i/2)$ (by expansion of the conditional probabilities, see below), implying that the event in the argument of $\Prob$ is a subset of $\prod_{i=1}^n(1-Q_i/2)\geq P_{max,\epsilon}$. This event is the same as the event $P_{max,\epsilon}^{-\alpha}\geq \prod_{i=1}^n(1-Q_i/2)^{-\alpha}$, since α is positive, and by monotonicity of $(.)^{\alpha}$. (Expressions ϕ of random variables that result in a truth value for each combination of values of the random variables determine the event $\{x|\phi(x)\}$, that is the set of points in the event space for which ϕ is true.)

By ``expansion of the conditional probabilities'' we mean recursive application of the identities

\[
\Prob(C_{\le n}|Z_{\le n}) =
  \Prob(C_{n}|C_{\le n-1}Z_{\le n})\Prob(C_{\le n-1}|Z_{\le n})
  = \Prob(C_{n}|C_{\le n-1}Z_{\le n})\Prob(C_{\le n-1}|Z_{\le n-1}).
\]

The second identity follows by independence of Z_n from the past, namely $C_{\le n-1}$ and $Z_{\le n-1}$. Note that Q as defined is determined by $C_{\le n-1}$ and $Z_{\le n-1}$ and, by the above expressions, provides the desired bound on $\Prob(C_{n}|C_{\le n-1} Z_{\le n})$. [I think a version of this expansion is also in Peter's proof and is key to these arguments.]

To get a quick and direct comparison to Peter's bound at n, we can set $\alpha=2m$ Note that for T a test factor, m cannot be arbitrarily large. I seem to recall that $m\leq 1/2$ in general, in which case $\alpha\leq 1$. For $\alpha=2m$, $T=T_{\alpha}$, and we get $P_{max, \epsilon}=(V\epsilon)^{-1/2m}$, where $V=\prod_{i=1}^n T_i$. Peter's bound for soundness is given as

\[
V\epsilon = (1+2m(1-Q_{max,\epsilon}^{1/n}))^n,
\]

where $Q_{max,\epsilon}$ is his probability bound (written as δ in the bound; also, he writes ϵ_s for ϵ and we set $\epsilon'=1$ since the latter only matters when quantifying completeness). Inverting Peter's bound gives

\[
Q_{max,\epsilon} = (1+(1-(V\epsilon)^{1/n})/2m)^{n},
\]

to be compared to $P_{max,\epsilon} = (V\epsilon)^{-1/2m}$. Set $r=(V\epsilon)^{1/n}$ and take logarithms to reduce this to comparing $\log(Q_{max,\epsilon})= n\log(1+(1-r)/\alpha)$ and $\log(P_{max,\epsilon})= -n\log(r)/\alpha$. To be useful, we need $r\ge 1$. By concavity of $\log$, this comparison favors Peter's bound. For our situation $r=(V\epsilon)^{1/n}$ is close to 1, say $r=1+\gamma$ with γ small compared to α. Both expressions are equal to $-n\gamma/\alpha + O((\gamma/\alpha)^2)$.