Limits of Random Oracles in Secure Computation

Mohammad Mahmoody
Cornell University
Hemanta K. Maji
University of California Los Angeles
Manoj Prabhakaran
University of Illinois Urbana-Champaign

Background

Impagliazzo-Rudich [2] showed that a Random Oracle is not sufficient to implement public-key encryption information-theoretically, thereby establishing a fundamental qualitative separation between public-key and private-key cryptography. This also had implications for Secure Function Evaluation or SFE (wherein Alice and Bob with inputs x and y, resp., compute $f(x,y)$ without revealing further information): Oblivious Transfer and other "complete" functions cannot be implemented using only a Random Oracle.

Our Results

We show that an RO, by itself (without computational assumptions), is useful for secure function evaluation exactly as much as an ideal commitment functionality is: f can be securely computed in the RO-model iff it can be computed in the "commitment-hybrid" model.

In particular, for security against semi-honest (passive) adversaries, an RO is not useful at all .

This holds for all 2-party deterministic SFE functions (even unsymmetric ones) with polynomial-domains.

What Does It Tell Us?

Informally, the computational hardness needed for secure evaluation of *any* function that does not have an unconditionally secure protocol, is *more complex than* what one-way functions (or any other "mini-crypt" primitive that can be implemented in the RO-model) provide. Tthis can be formalized as the impossibility of a "fully blackbox reduction" [5] of SFE to one-way functions.

These are the first results since [2], separating secure computation from mini-crypt primitives.

Future Work

- Our result is specific to *deterministic* SFE, as our analysis uses their combinatorial structure. No such structure is known for randomized SFE. But if we can "compile out" the RO in *any secure protocol*, our result can be extended to randomized SFE as well.
- In ongoing work, we consider oracles other than RO, that can lead to separations of SFE from public-key encryptions as well. More generally, we ask if we can uncover many worlds in "Impagliazzo's universe" for various (qualitatively different) SFE functionalities.

Preliminaries

Decomposable and Undecomposable Functions:

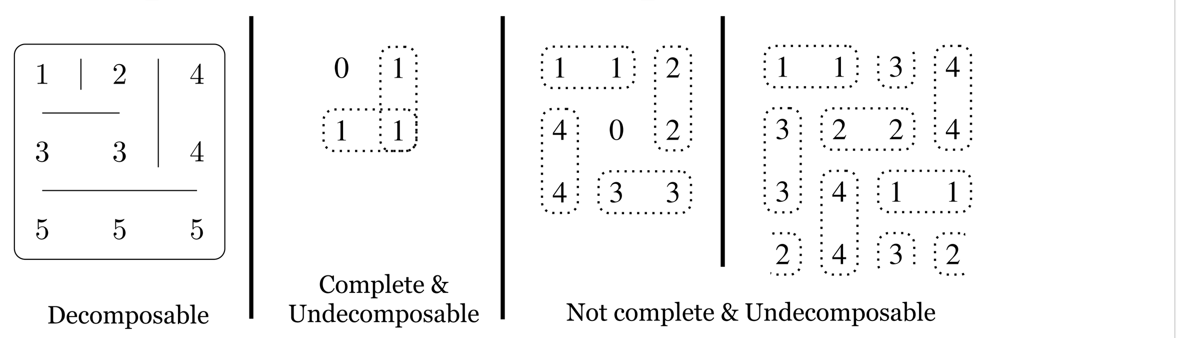

Decomposable functions are exactly those for which there are 2-party SFE protocols [3]. There are many undecomposable functions that are not complete. These are the ones for which our characterization newly rules out SFE protocols in the RO model.

Protocol Tree for a 2-party protocol: Has A (Alice) and B (Bob) nodes, that represent partial transcripts, with an edge from an A-node (resp. B-node) u to v if the next message from Alice (resp. Bob) given transcript u results in transcript v. The weight on an edge is the probability $P[v|u;x,y]$ where x,y are inputs.

Frontier Analysis of Protocols: We shall consider "frontiers" on the protocol tree where for the first time some property holds. For example [4] used frontiers F_X and F_Y where for the first time some significant amount of additional information about x (resp. y) is revealed by a single message.

Independence Learner: Following [2,1], there exists an Eve who queries the RO polynomially many times at each round, so that a locality property holds: conditioned on Eve's view so far, Alice's next message is almost independent of Bob's input (and vice-versa).

Augmented Protocol Tree: Contains Alice, Bob and Eve nodes. On an edge coming out of an Eve-node, Eve's interaction with RO is added to the transcript.

Proof Intuition

Suppose an undecomposable function f has a semi-honest SFE protocol in the RO model.

Plan: Define frontier F_X in the augmented protocol tree where a significant amount of new information about x is revealed by Alice, or is accumulated since last message from Alice. Similarly F_Y. Then:

- F_X and F_Y are almost "full": a transcript should pass through both, except with small probability.
- F_X occurs (on a random transcript path) "at or above" F_Y only with small probability; similarly for F_Y occurring "at or above" F_X.

Together we get a contradiction.

☑ *Fullness of frontiers*: because some information about both inputs must always be revealed (because of correctness and security, and undecomposability of f).

☑ *F_X is not strictly above F_Y* (and similarly, *F_Y is not strictly above F_X*) with significant probability: Else, Alice is revealing information about x independent of y; can be shown to be insecure if f is undecomposable.

▷ But could F_X and F_Y coincide?

☑ Intuitively, locality property ⇒ child of an A-node not on F_Y, and child of a B-node not on F_X.

▷ But F_X & F_Y could coincide at children of Eve-nodes.

i.e., information first revealed could be to Eve, it could depend on *both x and y*, and even be $f(x,y)$ itself.

To rule this out we give an attack to show that in case Eve's oracle queries reveal some information about x *and y*, then one of the two parties can extract (non-ideal) information using an imaginary execution (with a simulated RO) in which it alters its input.

Some Technical Details

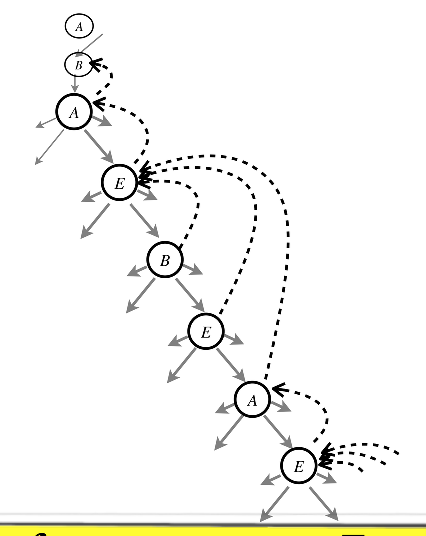

Apred relation: If a node v is the child of an A-node, then $Apred(v)$ = $Parent(v)$, else, $Apred(v)$ is the last node that is a child of an A-node on the path from root to v. (See fig.)

Similarly *Bpred* is defined.

F_X = {v| v is first node on a path from root s.t. $\exists y,x,x'$, $P[v|y]\geq\theta$ and $P[v|w;x,y]>(1+\delta).P[v|w;x',y]$ where w=$Apred(v)$}

(for suitably chosen δ and θ). The distributions are based on protocol execution with a random oracle and random inputs.

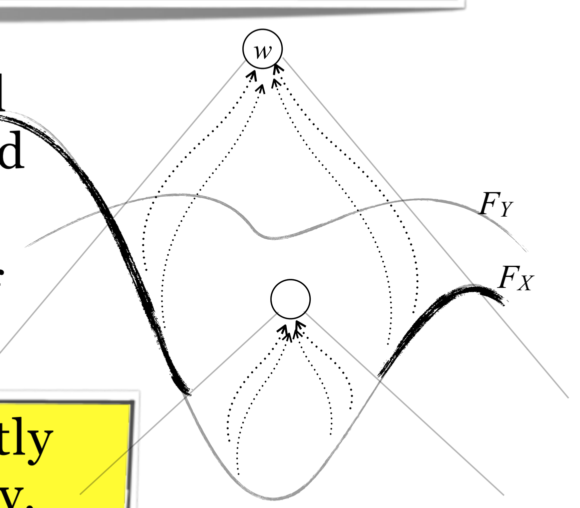

Similarly F_Y is defined in terms of *Bpred*.

IMPLIES

Claim: *$Apred(F_X)$ occurs strictly above F_Y only with small probability.*

- Suppose not. Then:
- Relying on undecomposability, we identify a suitable 2×2 minor of inputs $(x,x')\times(y,y')$, so that $f(x,y)\neq f(x',y)$ but $f(x,y')=f(x',y')$, and $\exists$ $G_X\subseteq F_X$, s.t. $P[v|w;x,y]$ > $(1+\delta')P[v|w;x',y]$ where w=$Apred(v)$, and $Apred(G_X)$ occurs strictly above F_Y, and $P[G_X|x,y]$ is large. We contradict this:
- Let $G_X = S_X \cup R_X$ s.t. $Apred(S_X)$ are Alice nodes, and $Apred(R_X)$ are children of Alice nodes.
- $P[S_X|x,y]$ can be bounded using the locality property.
- We bound $P[R_X|x,y]$ by giving an attack at R_X.

Attack at the Frontier

R_X is the part of the frontier F_X such that for $v\in R_X$, w = $Apred(v)$ is the child of an Alice node, w occurs strictly above F_Y, and $P[v|w;x,y] > (1+\delta')P[v|w;x',y]$.

Claim: *$P[R_X|x,y]$ is small.*

- If not we show how a curious Bob with input y' can mentally switch to y and distinguish between x and x'. On reaching w Bob samples an alternate view $V_{B,y}(w)$ corresponding input y. He simulates a RO conditioned on this view and Alice's input x^* (which he does not know) using access to the actual RO (which is conditioned on x^* and $V_{B,y'}(w)$) : queries in blue and orange views are answered according to those views; queries in green region are freshly answered, and the other queries are answered using the actual RO.

This works because of the "safety property" of the independence learner: that (w.h.p.) the orange and green regions don't intersect the gray region.

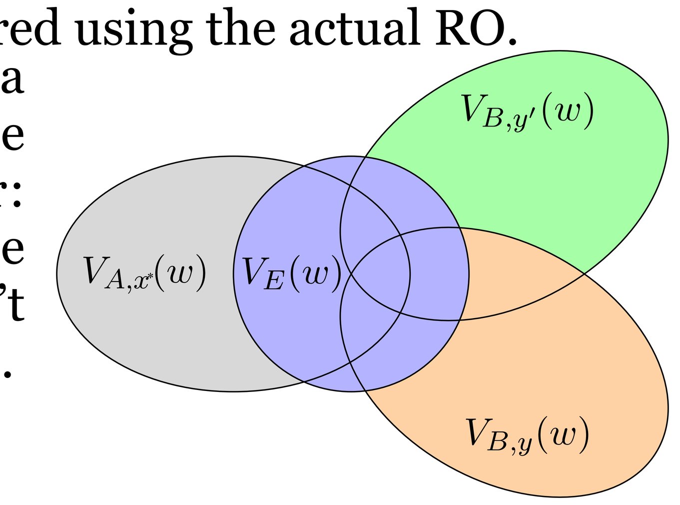

References

1. Barak, Mahmoody. *Merkle Puzzles are Optimal*. CRYPTO 2009.
2. Impagliazzo, Rudich. *Limits on the provable consequences of the one-way permutations*. STOC 1989.
3. Kushilevitz. *Privacy and communication complexity*. FOCS 1989.
4. Maji, Prabhakaran, Rosulek. *Complexity of multi-party computation problems: The case of 2-party symmetric secure function evaluation*. TCC 2009.
5. Reingold, Trevisan, Vadhan. *Notions of Reducibility between Cryptographic Primitives*. TCC 2004.

Research supported by NSF grants CNS 07-47027, CNS 07-16626, CCF 07-46990; AFOSR Award FA9550-10-1-0093; DARPA and AFRL contract FA8750-11-2-0211. The conclusions here do not necessarily represent the views of these agencies.