Comments on Eight TCC'09 Talks

Oded Goldreich
Department of Computer Science
Weizmann Institute of Science
Rehovot, Israel.
oded.goldreich@weizmann.ac.il

March 27, 2009. Minor revision: April 20, 2009

Preface: My comments refer to eight of the talks that I heard at the last TCC (see the conference's proceedings published as LNCS 5444). I wish to stress that this choice does not indicate that I consider other works less interesting. It is just that I either failed to attend the talks or failed to understand them sufficiently well (or, actually, did find the chosen eight more interesting).

1 An Optimally Fair Coin Toss

By Tal Moran, Moni Naor, and Gil Segev

The great success of theoretical research in cryptography, in the 1980s, led to a feeling that everything is possible. This view has been changing during the 1990s when the emergence of various impossibility results got more attention. Most of these impossibility results are not absolute but rather refer to natural well-defined abstractions of the known (and even "perceivable") proof techniques.

For example, while 2-message zero-knowledge proofs were proven to exist only for BPP sets, restricting the simulation (which underlies the definition of ZK) to black-box use of the adversary rules out also 3-message protocols as well as any constant-round public-coin protocols. Indeed, Barak's ZK protocol (simulated via non-blackbox use of the adversary's program) is the most dramatic example of a construction that bypasses the known limitations of a natural and general formulation of the known (and "perceivable") proof techniques.

The current paper provides yet another striking example. It refers to the tradeoff between the round complexity of a coin tossing protocol and the influence an adversary may have on its outcome (in a model that mandates that the honest party outputs a coin value also in the case that the other party aborts the protocol).

Cleve showed in 1986 that in any r-round protocol one party can influence the outcome by at least $\Omega(1/r)$, whereas a simple r-round protocol guarantees influence of at most $O(1/\sqrt{r})$. Furthermore, Cleve and Impagliazzo (1993) proved that, within a natural well-defined construction paradigm, no protocol can outperform the aforementioned simple protocol. Yet, the current paper presents an r-round protocol (which deviates from the latter construction paradigm) that guarantees influence of at most $O(1/r)$.

2 Complete Fairness in MPC Without an Honest Majority

By Dov Gordon and Jonathan Katz

Fairness refers to functionalities in which both parties are supposed to get an output. Loosely speaking, fairness means that one party may obtain its output only if all other parties can obtain their corresponding outputs. Note that fairness comes on top of privacy and correctness, and that it is reflected by the definition of the ideal model in which the actual execution is simulated: it refers to the operation the trusted party (in the ideal model); specifically, to whether or not the trusted party hands outputs to all parties simultaneously (or whether it can be "shut up" after handing outputs to some of the parties).

For two decades, Cleve's result regarding the impossibility of obtaining (almost) perfectly fair coin tossing protocols was interpreted as a "license to kill"; that is, as suggesting that we should give up on fairness in cryptographic settings in which there is no honest majority.

This perception has been corrected by the recent work "Complete Fairness in Secure Two-Party Computation" (of D. Gordon, C. Hazay, J. Katz, and Y. Lindell), and the current work extends the study from the two-party case to the multi-party case (lacking an honest majority). Indeed, these works indicate that fairness may be achievable in various cases, and so giving up on it in all cases may be too hasty. Thus, the question raised by these works is which functions can be computed fairly (without an honest majority).

A simple example provided in the current work refers to the computation of the OR function. In the two party case, any secure protocol (allowing abort) can be adapted to be fair by instructing the honest party to output its own input in case the other party has aborted. That is, an aborted execution is considered as one in which the aborting party has provided the zero input, and nothing is lost if this assumption is wrong (because in that case the adversary learns nothing). In the multi-party case, an abort is handled by re-initiating a new execution of the secure protocol without the aborting party, and this goes on until no parties that have aborted remain.

3 Security amplification for interactive cryptographic primitives

By Yevgeniy Dodis, Russell Impagliazzo, Ragesh Jaiswal, and Valentine Kabanets

The basic issue is amplifying noticeable hardness (i.e., problems that are hard for a noticeable fraction of the instances) into overwhelming hardness (i.e., problems that are hard for all but a negligible fraction of the instances). The archetypical example is Yao's Direct Product Lemma, and indeed the focus in general is on hardness amplification via direct product (i.e., defining a new problem who's instances are sequences of instances of the original problem).

While direct product is known to amplify hardness in various settings, it is know to fail in various interactive settings (see, e.g, the case of computationally-sound proof systems).

Actually, the current work does not address interactive protocols but rather "oracle type" functionalities such as MACs and PRFs. The general setting consists of a distribution over pairs of "keys", where a public (indexing) key is used to specify the oracle and private (verification) key is used to allow verification of the oracle answers. The task is to output a valid input-output pair (i.e., a pair accepted by the verification procedure) such that the input differs from all queries made to the oracle.

In the paper, the direct product problem is defined by interpreting each query (to the direct product problem) as a sequence of identical queries to the basic problems (which refer to different key-pairs), and this suffices for the cases of MACs and PRFs. However, it seems to me that the framework can be extended so to allow arbitrary queries to the direct product problem as long as one measures success by considering the number of coordinates in the final input-sequence that are different from all values that appeared in the (corresponding positions of the) query-sequences to the oracle; that is, denoting the final input sequence by $(i_1, ..., i_n)$ and denoting the $j^{\rm th}$ query by $(q_1^{(j)}, ..., q_n^{(j)})$, I refer to the set $\{k \in [n] : (\forall j)\, i_k \neq q_k^{(j)}\}$.

4 Adaptive ZK Proofs and Adaptively Secure OT

By Yehuda Lindell and Hila Zarosim

The problem of adaptive security has first arisen in the context of multi-party computation, but as pointed out by Beaver (STOC'96)[1] it is actually an issue even in the case of two-party computation. The current work is devoted to the study of adaptive security in the context of two-party computation.

Focusing on two central primitives, this work shows that one primitive (i.e., zero-knowledge) has adaptively secure protocols (based on any one-way functions), whereas constructing such protocols for the other primitive (i.e., oblivious transfer) based on enhanced trapdoor permutations requires non-blackbox techniques.

In addition to the intriguing nature of the two results, I wish to highlight the elegant ways in which they are derived. The positive result (i.e., adaptive zero-knowledge proofs for $\mathcal{NP}$) is derived by using a new type of instance-dependent commitment scheme, which (on the one hand) is perfectly tailored to the current application and (on the other hand) can be constructed based on any one-way function. The negative result (i.e., black-box separation of adaptively secure OT from standard OT) is obtained by an unprecedented reduction of the desired oracle separation to a known oracle separation.

5 On the (Im)possibility of Arthur-Merlin Witness Hiding

By Iftach Haitner, Alon Rosen, and Ronen Shaltiel

This work demonstrates that witness hiding protocols may have the same limitations as zero-knowledge protocols. Specifically, for certain natural distributions on the instances, constant-round public-coin protocols (of negligible soundness error) cannot be shown to be witness hiding by using black-box reductions (of the stand-alone search problem to the interactive witness-extraction problem).

I am particularly interested in this result because it *seems* to indicate that a strong notion of non-oblivious commitment scheme cannot be obtained by constant-round public-coin protocols. For details regarding what I mean here, see Appendix B of the final version of my paper with Boaz Barak on *universal arguments*. Indeed, I think it will be interesting to determine the possibility of such protocols. This also leads me to the somewhat related question regarding the possibility of constant-round public-coin proof systems (for $\mathcal{NP}$) that are *strongly witness-indistinguishable* (see Appendix C.3 in Vol II of *Foundation of Cryptography*).

[1]Seems that the issue was observed before by Beaver and Haber (Eurocrypt'92), in the context of two-party protcols for secure communication.

Some remotely related works. This conference has seen many black-box separation results. In addition to the two results that were mentioned above, I wish to mention the work "On the (Im)possibility of Key Dependent Encryption" (by Iftach Haitner and Thomas Holenstein). This result stands in contrast to a positive result by Hofheinz and Unruh (Eurocrypt'08), which refers only to encrypting a small amount of information regarding the key, and to a positive result of Boneh *et al.* (Crypto'08) that relies on a specific assumption.

6 LEGO for Two Party Secure Computation

By Jesper Buus Nielsen and Claudio Orlandi

Recently, the generic compilation of two-party protocols from the semi-honest model to the malicious model (via zero-knowledge) was replaced by cut-and-choose techniques. These techniques were applied to Yao's garbled circuit as a whole, while the current work suggests to apply these techniques at the individual gate level. I find this suggestion extremely appealing, and only regret that the implementation is based on a specific problem (i.e., DLP). Note that a generic implementation will require an alternative way of coping with the problem of connecting the various gate gadgets.

Some remotely related works. I wish to mention two other works that refer to issues related to the efficient compilation of protocols. The first work, "Simple, Black-Box Constructions of Adaptively Secure Protocols" (by Seung Geol Choi, Dana Dachman-Soled, Tal Malkin, and Hoeteck Wee) referred to transformation of Oblivious Transfer from the (adaptive) semi-honest model to the (adaptive) malicious model. The second work "Black-Box Constructions of Two-Party Protocols from One-Way Functions" (by Rafael Pass and Hoeteck Wee) refers to black-box transformations of protocols (rather than a transformation of the protocol's code), which tends to be more efficient than non-blackbox transformations.

7 Chosen-Ciphertext Security via Correlated Products

By Alon Rosen and Gil Segev

Standard hardness amplification results refer to independently chosen inputs. Nevertheless, in the context of complexity theory (where the problems instances need not be easy for the party that generates them), hardness amplification results are also known for the case of correlated inputs (e.g., pairwise independent inputs, inputs generated by random walks on expander graphs, etc).

In this paper the (a priori known) correlation is among the solutions (e.g., the preimages of one-way function) rather than among the instances (e.g., the corresponding images). That is, for a generator G, we consider inverting f on the instances $f(x_1), ..., f(x_n)$, where $(x_1, ..., x_n) \leftarrow G(U_k)$, rather than inverting f on the instances $y_1, ..., y_n$, where $(y_1, ..., y_n) \leftarrow G(U_k)$

This paper shows a simple construction of CCA secure public-key encryption scheme based on a trapdoor permutation that is hard to invert for some correlated preimage patterns. Needless to say, it begs to study the plausibility of this as well as other "correlated input" problems. (Note that the Naor-Reingold pseudorandom synthesizers actually refer to such a correlated input problem.) It will be also interesting to explore other applications that correlated input problems may have (i.e., to the construction of various cryptographic primitives).

8 Towards a Theory of Extractable Functions

By Ran Canetti and Ronny Ramzi Dakdouk

A specific extractability assumption was proposed by Hada and Tanaka (in Crypto'98), and a general study of such assumptions has been begging since. Such a study has been implicitly initiated by Canetti and Dakdouk (ICALP'08) and is the focus of the current work.

Loosely speaking, an extractability assumption asserts that given any program that can generate a certain distribution, one can extract from this program a related distribution (while it is hard to transform instances of the first distribution to the second distribution). Specifically, for a function f, the extractability assumption asserts that given any program that generate a distribution Y on the image of f, one can derive a program that generates a distribution X over the domain of f such that $f(X) \approx Y$. Indeed, note that the assumption cannot possibly hold if f is a one-way permutation (nor if it is a regular one-way function with a dense range).

The current work distinguishes two types of extractability assumptions, called *noninteractive* and *interactive*. For reasons to be explained, I do not like these names and tentatively suggest to use *basic* and *correlated*. The foregoing example refers to the basic setting, whereas in the correlated setting the reference is to programs that can generate related Y and Y' in the image of f. Specifically, if $y = f(x, r)$ then y' should equal $f(x, r')$, where r' is given to the program after it has produces y. Thus, one may say that some interaction goes on here, but this interaction is very limited and this limitation is essential (because allowing unrestricted interaction reduces the current setting to the one of proofs of knowledge).

In any case, it is interesting to investigate these (and possible other) extractability assumptions (both from the angle of their plausibility and from the angle of their applicability).