GTACS @MTA, March 9th, 2016

When: Wednesday, March 9, 9:30 – 16:15, 2016

Where: Tel-Aviv Yaffo Academic College (MTA)

Weston Bld. 007 (וסטון 007)

Address: 10 Khever ha-Le'umim St, Tel Aviv-Yaffo
רחוב חבר הלאומים 10, תל-אביב יפו
(map)

Parking: Free parking at MTA lot, 10 Khever Ha-Le'umin St. (fills up early).
City parking in a short walking distance: Bloomfield Ahuzat Hahof parking lot, HaThiya St. 11, Tel-Aviv Yaffo**.**

Registration: Register here (participation is free but registration is required).

Description: Event announcement.

Program

9:30--10:00 Coffee + Gathering

10:00--11:00 **Zvika Brakerski (Weizmann)**
Hierarchical Functional Encryption (Even If Obfuscation Does Not Exist)

11:00--11:15 Coffee Break

11:15--12:15 **Benny Applebaum (TAU)**
Algebraic Attacks against Random Local Functions
and Their Countermeasures

12:15--12:30 Coffee Break

12:30--13:00 **Nimrod Aviram (TAU)**
DROWN: Breaking TLS using SSLv2

13:00--14:00 **Lunch break** (lunch provided by us)

14:00--15:00 **Niv Gilboa (BGU)**
Succinct Secure Computation of Branching Programs from DDH

15:00--15:15 Coffee Break

15:15--16:15 **Uri Stemmer (BGU)**
Algorithmic Stability for Adaptive Data Analysis

16:15 Sending you home

Abstracts

Zvika Brakerski (The Weizmann Institute of Science)
Hierarchical Functional Encryption (Even If Obfuscation Does Not Exist)

Functional encryption provides fine-grained access control for encrypted data, allowing each user to learn only specific functions of the encrypted data. We study the notion of hierarchical functional encryption, which augments functional encryption with delegation capabilities, offering significantly more expressive access control.

We present a *generic* transformation that converts any general-purpose public-key functional encryption scheme into a hierarchical one without relying on any additional assumptions. Our transformation applies across the board to any functional encryption scheme, including parameter regimes where constructions are known based on standard assumptions (such as LWE or even just the existence of PKE). Thus it implies that hierarchical functional encryption is likely to exist even if more elaborate primitives (such as program obfuscation) do not.

Joint work with Gil Segev.

Benny Applebaum (Tel Aviv University)
Algebraic Attacks against Random Local Functions and Their Countermeasures

Suppose that you have n truly random bits $x=(x_1,\ldots,x_n)$ and you wish to use them to generate $m >> n$ pseudorandom bits $y=(y_1,\ldots,y_m)$ using a local mapping, i.e., each y_i should depend on at most $d=O(1)$ bits of x. In the polynomial regime of $m=n^s$, $s>1$, the only known solution, originates from (Goldreich, ECCC 2000), is based on *Random Local Functions*: Compute y_i by applying some fixed (public) d-ary predicate P to a random (public) tuple of distinct inputs $(x_{i1},\ldots,x_{id})$. In this talk, we will try to understand, for any value of s, how the pseudorandomness of the resulting sequence depends on the choice of the underlying predicate.
Time permitting, I will describe a recent construction of Pseudorandom Functions with low complexity based on Goldreich's function.

Based on joint works with Lovett (STOC 2016, ECCC TR15-172) and Raykov (unpublished).

Nimrod Aviram (Tel-Aviv University)
DROWN: Breaking TLS using SSLv2

We present DROWN, a novel cross-protocol attack that can decrypt passively collected TLS sessions from up-to-date clients by using a server supporting SSLv2 as a Bleichenbacher RSA padding oracle. We present two versions of the attack. The more general form exploits a combination of thus-far unnoticed protocol flaws in SSLv2 to develop a new and stronger variant of the Bleichenbacher attack. A typical scenario requires the attacker to observe 1,000 TLS handshakes, then initiate 40,000 SSLv2 connections and perform 2^50 offline work to decrypt a 2048-bit RSA TLS ciphertext. (The victim client never initiates SSLv2 connections.) We implemented the attack and can decrypt a TLS 1.2 handshake using 2048-bit RSA in under 8 hours using Amazon EC2, at a cost of $440. Using Internet-wide scans, we find that 33% of

all HTTPS servers and 22% of those with browser-trusted certificates are vulnerable to this protocol-level attack, due to widespread key and certificate reuse. For an even cheaper attack, we apply our new techniques together with a newly discovered vulnerability in OpenSSL that was present in releases from 1998 to early 2015. Given an unpatched SSLv2 server to use as an oracle, we can decrypt a TLS ciphertext in one minute on a single CPU—fast enough to enable man-in-the-middle attacks against modern browsers. 26% of HTTPS servers are vulnerable to this attack. We further observe that the QUIC protocol is vulnerable to a variant of our attack that allows an attacker to impersonate a server indefinitely after performing as few as 2^25 SSLv2 connections and 2^65 offline work. We conclude that SSLv2 is not only weak, but actively harmful to the TLS ecosystem.

More details, as well as a version of the paper, can be found at drownattack.com.

Niv Gilboa (Ben-Gurion University)
Succinct Secure Computation of Branching Programs from DDH

Under the Decisional Diffie-Hellman assumption, we present a 2-out-of-2 secret sharing scheme that supports a compact evaluation of branching programs on the secrets by non-interactive computation on the shares.

We present several applications of this result to communication-efficient secure computation and related primitives. These include secure two-party protocols for evaluating NC1 circuits or log-space functions in which the communication complexity is linear in the input and output size, protocols for general circuits in which the communication complexity is slightly sublinear in the circuit size, efficient two-party function secret sharing for NC1 and as a consequence two-server PIR protocols supporting general NC1 queries.

Joint work with Elette Boyle and Yuval Ishai

Uri Stemmer (Ben-Gurion University)
Algorithmic Stability for Adaptive Data Analysis

Adaptivity is an important feature of data analysis - the choice of questions to ask about a dataset often depends on previous interactions with the same dataset. However, statistical validity is typically studied in a nonadaptive model, where all questions are specified before the dataset is drawn. Recent work by Dwork et al. (STOC, 2015) initiated the formal study of this problem, and gave the first upper bounds on the achievable generalization error for adaptive data analysis.

The results of Dwork et al. are based on a connection with algorithmic stability in the form of differential privacy. We extend their work by giving a quantitatively optimal, more general, and simpler proof of their main theorem that stable algorithms of the kind guaranteed by differential privacy imply low generalization error. We also show that weaker stability guarantees such as bounded KL divergence and total variation distance lead to correspondingly weaker generalization guarantees.

Joint work with Raef Bassily, Kobbi Nissim, Adam Smith, Thomas Steinke, Jonathan Ullman