I-CORE ALGO Day

Cryptography and Security

9:15-10:00: Gathering

10:00-11:00: Rasmus Pagh (IT University of Copenhagen)

New algorithms for similarity search

The ability to handle noisy or imprecise data is becoming increasingly important in computing. One building block to achieve this are data structures for finding similar vectors (aka. “near neighbors”) in a collection of, possibly high-dimensional, vectors. This talk gives an overview of randomized techniques for high-dimensional similarity search, and discusses recent advances towards making these techniques more widely applicable by 1) reducing the space usage, and 2) eliminating the probability of error. Mathematical objects that are not typical tools in data structure design, more precisely covering designs and unbalanced expander graphs, turn out to play an important role. We conclude with some questions on explicit constructions of these objects whose resolution would lead to improved data structures.

11:00-11:30: Coffee Break

11:30-12:15: Alon Rosen

On the Cryptographic Hardness of Finding a Nash Equilibrium

In this talk we connect the complexity of finding a Nash equilibrium of a game with cryptography. In particular, we prove that finding a Nash equilibrium of a game is hard, assuming the existence of indistinguishability obfuscation and injective one-way functions with sub-exponential hardness. This is achieved by showing how these cryptographic primitives give rise to a hard computational problem that lies in the complexity class PPAD, for which finding Nash equilibrium is known to be complete.

Our result provides further evidence of the intractability of finding a Nash equilibrium, one that is extrinsic to the evidence presented so far.

This is joint work with Nir Bitansky and Omer Paneth.

12:15-13:00: Henry Corrigan-Gibbs (Stanford University)

Space-Hard Functions for Password Hashing

In this talk, I will introduce the problem of password hashing and will explain the reasons to prefer space-hard password hashing functions over conventional cryptographic hashes. Next, I will discuss the design and analysis of the Balloon hash functions, a new family of space-hard hash functions that:

 * Require a certain amount working space to compute efficiently
 * Exhibit password-independent memory access patterns
 * Are fast enough for real-world use.

The Balloon hash algorithms are surprisingly easy to describe but arguing formally about their security properties presents a number of technical challenges. Finally, I will conclude with discussion of recent results on space-hard hashing in the multi-instance setting
and with related open problems.

13:00-14:30: Lunch Break

14:30-14:50: Ido Shahaf (HUJI)

Searchable Symmetric Encryption: Optimal Locality in Linear Space via Two-Dimensional Balanced Allocations

Searchable symmetric encryption (SSE) enables a client to store a database on an untrusted server while supporting keyword search in a secure manner. Despite the rapidly increasing interest in SSE technology, experiments indicate that the performance of the known schemes scales badly to large databases. Somewhat surprisingly, this is not due to their usage of cryptographic tools, but rather due to their poor locality (where locality is defined as the number of non-contiguous memory locations the server accesses with each query). The only known schemes that do not suffer from poor locality suffer either from an impractical space overhead or from an impractical read efficiency (where read efficiency is defined as the ratio between the number of bits the server reads with each query and the actual size of the answer).

We construct the first SSE schemes that simultaneously enjoy optimal locality, optimal space overhead, and nearly-optimal read efficiency. Specifically, for a database of size N, under the modest assumption that no keyword appears in more than $N^{1 - 1/\log \log N}$ documents, we construct a scheme with read efficiency $\tilde{O}(\log \log N)$. This essentially matches the lower bound of Cash and Tessaro (EUROCRYPT '14) showing that any SSE scheme must be sub-optimal in either its locality, its space overhead, or its read efficiency. In addition, even without making any assumptions on the structure of the database, we construct a scheme with read efficiency $\tilde{O}(\log N)$.

Our schemes are obtained via a two-dimensional generalization of the classic balanced allocations ("balls and bins'') problem that we put forward. We construct nearly-optimal two-dimensional balanced allocation schemes, and then combine their algorithmic structure with subtle cryptographic techniques.

Joint work with Gilad Asharov, Moni Naor and Gil Segev

14:55-15:15: Marshall Ball (Columbia University)

Non-Malleable Codes for Bounded Depth Circuits

Informally, non-malleable codes provide a means of transforming an adversarial channel into a channel whose output is either identical to or unrelated to the input. We show how to construct efficient, unconditionally secure non-malleable codes for bounded output locality. In particular, our scheme is resilient against functions such that any output bit is dependent on at most n^{δ} bits, where n is the total number of bits in a codeword and $0\leq \delta < 1$ a constant. Notably, this tampering class includes $\textsf{NC}^0$.

Joint work with Dana Dachman-Soled, Mukul Kulkarni, and Tal Malkin.

15:15-16:00: Eli Ben-Sasson

Public-setup computational integrity from quasi-linear PCPs

Interactive proof (IP) and argument systems have many promising applications to secure decentralized payment systems like Bitcoin and Zerocash. The diversity of applications calls for IP systems that can be applied efficiently to "natural" languages in NTIME(T(n)) and which scale-well, asymptotically and concretely, with instance size.

Recently implementations of scalable IP systems rely on a trusted setup phase that involves a secret trapdoor, that, if compromised, ruins security.

In this talk I will discuss our recent implementation of a scalable IP system that has no trapdoors, and which is founded on efficient quasi-linear probabilistically checkable proofs (PCP)s.

Based on joint works with Iddo Ben-Tov, Alessandro Chiesa, Ariel Gabizon, Daniel Genkin, Matan Hamilis, Evgenya Pergament, Michael Riabzev, Mark Siberstein, Eran Tromer and Madars Virza.

16:00-16:30: Coffee Break

16:30-17:30: Tal Rabin (IBM)

Secure MPC with General Interaction Patterns

We present a unified framework for studying secure multiparty computation (MPC) with arbitrarily restricted interaction patterns such as a chain, a star, a directed tree, or a directed graph.

The talk will be self contained and does not require prior knowledge of MPC.

Joint work with Halevi, Ishai, Jain, Kushilevitz