Cryptography:

10:00-11:00 – Shan Muthukrishnan (Rutgers University):

Heavy Hitters on Data Streams and Recent Variants

Abstract:

The data stream model focused on processing data with sublinear storage, and one of the traditional tasks in this model is identifying the heavy hitters (items that appear with overwhelming frequency, HHs). In this talk, I will provide an overview of HH algorithms, and focus on some of the recent variants: HHs seen from modern software defined networking (SDNs), HHs with very high dimensional data motivated by web analytics, HHs with pan-private guarantees, and other notions of heavy hitters including H-Index variants and multigraph versions. This problem continues to represent what we can do efficiently under many computing, space, communication, and other constraints.

11:30-12:15 – Yehuda Lindell (Bar Ilan University):

High-Throughput Secure Three-Party Computation with an Honest Majority

Abstract:

Protocols for secure multiparty computation enable a set of parties to compute a joint function of their inputs, without revealing anything except for the output. In recent years, there has been huge progress in the design of efficient secure multiparty computation protocols, and it is now beginning to be deployed in practice. In this talk, we outline recent work for the setting of three parties and an honest majority (i.e., at most one corrupted party) achieving extraordinarily high throughput. We are able to achieve rates of over 7 billion AND gates per second for the case of semi-honest adversaries, and over 1 billion AND gates per second for the case of malicious adversaries (on a cluster of three 20-core machines). These rates make it possible to carry out very large computations in reasonable time.

12:15-13:00 – Tal Moran (IDC):

Topology Hiding Computation on All Graphs

Abstract:

A distributed computation in which nodes are connected by a partial communication graph is called topology-hiding if it does not reveal information about the graph beyond what is revealed by the output of the function. Previous results have shown that topology-hiding computation protocols exist for graphs of constant degree and logarithmic diameter in the number of nodes [Moran-Orlov-Richelson, TCC'15; Hirt et al., Crypto'16] as well as for specific graph families with larger diameter, such as cycles, trees, and low circumference graphs [Akavia-Moran, Eurocrypt'17], but the feasibility question for general graphs was open. In this work we positively resolve the above open problem: we construct a protocol for topology-hiding computation that works for any graph, and is secure under the Decisional Diffie-Hellman assumption.

Joint work with Adi Akavia and Rio LaVigne

14:30-15:15 – Yuval Ishai (UCLA & Technion):

Secure Arithmetic Computation with Constant Computational Overhead

Abstract:

Motivated by the goal of efficient secure computations on sensitive numerical data, we present a protocol for securely computing arithmetic circuits that requires only a constant (amortized) number of arithmetic operations per gate. Our protocol is based on new cryptographic assumptions that can be viewed as natural arithmetic analogues of well studied assumptions. Beyond the asymptotic result, a key building block in our protocol can yield concrete efficiency improvements for natural secure computation tasks.

Joint work with Benny Applebaum, Ivan Damgård, Michael Nielsen, and Lior Zichron

15:15-16:00 – Eylon Yogev (Weizmann):

Search Problems: A Cryptographic Perspective

Abstract:

The class TFNP is the search analog of NP with the additional guarantee that any instance has a solution. TFNP has attracted extensive attention due to its natural syntactic subclasses that capture the computational complexity of important search problems from algorithmic game theory, combinatorial optimization and computational topology. Thus, one of the main research objectives in the context of

TFNP is to search for efficient algorithms for its subclasses, and at the same time proving hardness results where efficient algorithms cannot exist.

Currently, no problem in TFNP is known to be hard under assumptions such as NP hardness, the existence of one-way functions, or even public-key cryptography. The only known hardness results are based on less general assumptions such as the existence of collision-resistant hash functions, one-way permutations less established cryptographic primitives (e.g. program obfuscation or functional encryption).

Several works explained this status by showing various barriers to proving hardness of TFNP. In particular, it has been shown that hardness of TFNP hardness cannot be based on worst-case NP hardness, unless NP=coNP. Therefore, we ask the following question: What is the weakest assumption sufficient for showing hardness in TFNP?

In this talk, I will answer this question and show that hard-on-average TFNP problems can be based on the weak assumption that there exists a hard-on-average language in NP. In particular, this includes the assumption of the existence of one-way functions. In terms of techniques, there is an interesting interplay between problems in TFNP and derandomization techniques.

Based on joint works with Pavel Hubáček, and Moni Naor.

16:30-17:30 – Bernhard Haeupler (Carnegie Mellon University):

Distributed Optimization Algorithms via Low-Congestion Shortcuts

Abstract:

How fast a non-local distributed optimization problem can be solved in a given network depends in a highly non-trivial manner on the topology of the network. This talk will introduce a simple graph structure, called low-congestion shortcuts, which often tightly characterize and capture this dependency. Low-congestion shortcuts furthermore make it easy to design near optimal distributed algorithms for a wide variety of problems. For example, this leads to MST and approximate min-cut and shortest-path algorithms which require only a near linear number of messages and have the optimal O~(sqrt{n} + D) running times on worst-case network topologies while also achieving a near instance-optimal O~(D) running times on planar, low-genus, low-treewidth and other non-pathological network topologies.