CFEM workshop: Theory and Practice of Multiparty Computation
4 to 8 June, 2012
Aarhus University, Denmark

Speaker: Claudio Orlandi, Bar Ilan University

Title: Hiding the Input-Size in Secure Two-Party Computation

In the setting of secure multiparty computation, a set of parties wish to compute a joint function of their inputs, while preserving properties like privacy, correctness, and independence of inputs. One security property that has typically not been considered in the past relates to the length or size of the parties inputs. This is despite the fact that in many cases the size of a party's input can be confidential. The reason for this "omission" seems to have been the folklore belief that, as with encryption, it is impossible to carry out non-trivial secure computation while hiding the size of parties' inputs. In a surprising recent result, Ateniese et al. (PKC 2011) showed that it is possible to hide the input size of one of the parties when computing secure two-party set intersection. This suggests that the folklore belief may not be fully accurate.

In this work, we initiate a theoretical study of input-size hiding secure computation, and focus on the two-party case. We present definitions for this task, and deal with the subtleties that arise in the setting where there is no a priori polynomial bound on the parties' input sizes. Our definitional study yields a multitude of classes of input-size hiding computation, depending on whether a single party's input size remains hidden or both parties' input sizes remain hidden, and depending on who receives output and if the output size is hidden from a party in the case that it does not receive output. We prove feasibility and impossibility results for input-size hiding secure two-party computation. Some of the highlights are as follows:

- Under the assumption that fully homomorphic encryption (FHE) exists, there exist non-trivial functions (e.g., the millionaire's problem) that can be securely computed while hiding the input size of both parties.

- Under the assumption that FHE exists, every function can be securely computed while hiding the input size of one party, when both parties receive output (or when the party not receiving output does learn the size of the output). In the case of functions with fixed output length, this implies that every function can be securely computed while hiding one party's input size.

- There exist functions that cannot be securely computed while hiding both parties' input sizes. This is the first formal proof that, in general, some information about the size of the parties' inputs must be revealed.

This is joint work with Yehuda Lindell and Kobbi Nissim.