

Many modern computing systems are distributed and built in a decentralized way on top of large networks. As one of the key challenges when running a distributed algorithm in a large network, typically no node can know the state of the whole network and each node thus each node has to base its decisions on only partial, mostly local information about the network. The major goal of the research on local distributed algorithms is to understand to what extent nodes in a network can achieve a global goal such as solving some graph problem, based on only local information.

In my talk, I will focus on the role of using randomization in local distributed graph algorithms. In many cases, there currently is a large gap between the best randomized and deterministic algorithms and understanding whether this gap is inherent, is considered to be one of the central open questions of the area. I will discuss recent results that shed some light on what randomness is really needed for and that also allow to identify simple and basic problems that are complete for the class of problems that have efficient randomized distributed algorithms and for which we do not know efficient deterministic algorithms.