

Dealing with betweenness in evolving graphs and imposed system workload imbalance (Invited talk)

Nicolas Kourtellis¹

¹ Telefonica Research, Barcelona, Spain

Abstract

Graph problems such as betweenness centrality can be notoriously hard to compute in a static setting. The problems of efficiency and scalability to handle million-node graphs are exacerbated in a dynamic setting, where the input is an evolving graph seen edge by edge, and the goal is to keep the graph measure under computation up to date. In order to solve such problems efficiently, and compute respective graph metrics online, some have proposed approximation or incremental methods, and others have proposed parallelized execution of properly defined subtasks on parallel and/or streaming platforms. In the first part of this talk, I will cover a framework that applies both strategies for computing betweenness centrality on evolving graphs while generating scores on pseudo-real time. In the second part, I will cover efforts to address system workload imbalance arising when computing such highly skewed graph metrics, which impose uneven load to the computation nodes of the streaming platform.