

Finding large near-cliques in massive networks is a notoriously hard problem of great importance to many applications, including anomaly detection in security, community detection in social networks, and mining the Web graph. How can we exploit idiosyncrasies of real-world networks in order to solve this NP-hard problem efficiently? Can we find dense subgraphs in graph streams with a single pass over the stream? Can we design near real time algorithms for time-evolving networks? In this talk I will answer these questions in the affirmative. I will also present state-of-the-art exact and approximation algorithms for extraction of large near-cliques from large-scale networks, the k-clique densest subgraph problem, which run in a few seconds on a typical laptop. I will present graph mining applications, including anomaly detection in citation networks, and planning a successful cocktail party. I will conclude my talk with some interesting research directions.