abstract

Computer Science/Discrete Mathematics Seminar II
Topic:

Speaker:

Affiliation:

Date:

Time/Room:

A basic fact of algebraic graph theory is that the number of connected components in an undirected graph is equal to the multiplicity of the eigenvalue zero in the Laplacian matrix of the graph. In particular, the graph is disconnected if and only if there are at least two eigenvalues equal to zero. Cheeger's inequality and its variants provide an approximate version of the latter fact; they state that a graph has a sparse (non-expanding) cut if and only if there are at least two eigenvalues that are close to zero.

It has been conjectured that an analogous characterization holds for higher multiplicities, i.e., there are k eigenvalues close to zero if and only if the vertex set can be partitioned into k subsets, each defining a cut of small expansion. We resolve this conjecture. Our result provides a theoretical justification for clustering algorithms that use the bottom k eigenvectors to embed the vertices into R^k, and then apply geometric considerations to the embedding.

We also show that these techniques yield a nearly optimal tradeoff between the expansion of sets of size approximately n/k, and the kth smallest eigenvalue of the normalized Laplacian matrix, denoted λ_k. In particular, we show that in every graph there are at least $k/2$ disjoint sets (one of which will have size at most $2n/k$) having each expansion at most $O(\sqrt{\lambda_k \log k})$. This result was also discovered independently by Louis, Raghavendra, Tetali, and Vempala. This bound is tight, up to constant factors, for the "noisy hypercube" graphs.