Assaf Naor's talk on Sparsest cut

Bogdan Grechuk • 23 Feb 2026

This post continues describing our series of online seminars devoted to accessible presentations of some of the most significant mathematical theorems of the 21st century. I am delighted to report that on February 11 Prof. Avi Wigderson delivered a beautiful and highly accessible lecture on expander graphs. If you missed the talk, the recording is now available on YouTube.

The next lecture in the series will take place on

Wednesday, March 4, 2026, at 4:00 PM (UK time).

Prof. Assaf Naor (Princeton University, USA) will speak on the topic of *The Sparsest Cut Problem*.

To join the talk, simply click here at the scheduled time. No registration is required. For a full list of upcoming seminars, please visit the seminar webpage: https://th21.le.ac.uk/next-talks/.

For announcements of other talks, as well as descriptions of recent major mathematical breakthroughs, see the full list of posts on this blog.

The sparsest cut problem. The *sparsest cut problem* is a central combinatorial optimization problem with deep connections to geometry, analysis, and theoretical computer science. Given a finite graph $G=(V,E)$, the goal is to partition the vertex set V into two nonempty parts so as to minimize the ratio

$$r(S)=\frac{\#\{\text{edges crossing from } S \text{ to } V\setminus S\}}{\min\{|S|,|V\setminus S|\}}.$$

The minimum possible value,

$$\alpha(G)=\min_{\emptyset\neq S\subsetneq V} r(S),$$

is known as the *edge expansion* (or *Cheeger constant*) of G.

Computing $\alpha(G)$ exactly is NP-hard. Nevertheless, efficient approximation algorithms have been developed. In 1999, Leighton and Rao (Leighton and Rao 1999) obtained a polynomial-time $O(\log n)$-approximation algorithm: for every n-vertex graph G, one can find a cut with ratio at most $O(\alpha(G)\log n)$.

A major breakthrough came in 2009, when Arora, Rao and Vazirani (Arora et al. 2009) improved the approximation ratio to $O(\sqrt{\log n})$. Thus there exists a polynomial-time algorithm producing a cut satisfying

$$r(S) = O\big(\alpha(G)\sqrt{\log n}\big).$$

The weighted problem and the recent breakthrough.

A more general and substantially more delicate version is the *sparsest cut problem with general capacities and demands*. Here we are given two symmetric $n \times n$ matrices $C = (c_{ij})$ and $D = (d_{ij})$ with nonnegative entries. The matrix C encodes edge capacities, while D encodes demands between pairs of vertices.

For any nonempty proper subset $S \subsetneq [n] = \{1, \dots, n\}$, define

$$\mathrm{cut}(S) = \frac{\displaystyle\sum_{i \in S}\sum_{j \notin S} c_{ij}}{\displaystyle\sum_{i \in S}\sum_{j \notin S} d_{ij}}.$$

The objective is to minimize $\mathrm{cut}(S)$.

In a recent paper (Chang et al. 2024), accepted for publication in *Acta Mathematica*, Chang, Naor and Ren proved a polynomial-time

$$O(\sqrt{\log k})$$

approximation algorithm when the number of nonzero entries of D is at most k. More precisely, the algorithm outputs a subset S satisfying

$$\mathrm{cut}(S) \le O(\sqrt{\log k}) \min_{\emptyset \neq T \subsetneq [n]} \mathrm{cut}(T).$$

This matches the best possible asymptotic bound and settles a long-standing open problem.

Geometry behind the algorithms.

The proof draws on deep ideas from the geometry of metric spaces and the phenomenon of measure concentration. The guiding principle is to represent the vertices of a graph as points in a suitable metric space so that distances reflect combinatorial structure. Once such an embedding is constructed, one can partition the graph according to geometric clusters in the target space.

Semidefinite programming provides an efficient way to construct these embeddings. The quality of the resulting cut depends on the *distortion* of the embedding: better distortion yields better approximation ratios.

A metric space (X, ρ_X) admits a *bi-Lipschitz embedding* into a metric space (Y, ρ_Y) if there exist constants $m, M > 0$ and a mapping $f : X \to Y$ such that

$$m\,\rho_X(x,y) \le \rho_Y\big(f(x), f(y)\big) \le M\,\rho_X(x,y) \qquad \forall x, y \in X. \tag{1}$$

The ratio M/m is called the *distortion* of the embedding. Minimizing distortion is a classical theme in geometric functional analysis, with powerful algorithmic consequences.

In 1985, Bourgain proved that every n-point metric space embeds into a Euclidean space with distortion $O(\log n)$. For general metrics this bound is essentially optimal. However, improved bounds are possible for special classes.

One such class is that of *metrics of negative type*. A finite metric space (X, ρ) with points $x_1, \dots, x_n$ is of negative type if for all real numbers $c_1, \dots, c_n$ satisfying $\sum_{i=1}^{n} c_i = 0$,

$$\sum_{i,j=1}^{n} c_i c_j \rho(x_i, x_j) \le 0.$$

In 2008, Chawla, Gupta and Räcke (Chawla et al. 2008) proved that every n-point metric space of negative type embeds into Euclidean space with distortion $O((\log n)^{3/4})$. Shortly thereafter, Arora, Lee and Naor (Arora et al. 2008) improved this to $O(\sqrt{\log n}\log\log n)$.

The recent breakthrough of Chang, Naor and Ren (Chang et al. 2024) sharpens this to the optimal bound

$$O(\sqrt{\log n}),$$

matching a lower bound of Enflo (1969). This optimal embedding theorem directly yields the optimal $O(\sqrt{\log k})$ approximation algorithm for the sparsest cut problem with general capacities and demands.

Prof. Naor's lecture will present these ideas in an accessible way, explaining how geometric insights about metric embeddings ultimately lead to optimal algorithms for fundamental combinatorial optimization problems.

References

Arora, Sanjeev, James Lee, and Assaf Naor. 2008. "Euclidean Distortion and the Sparsest Cut." *J. Amer. Math. Soc.* 21 (1): 1–21.

Arora, Sanjeev, Satish Rao, and Umesh Vazirani. 2009. "Expander Flows, Geometric Embeddings and Graph Partitioning." *J. ACM* 56 (2): 1–37.

Chang, Alan, Assaf Naor, and Kevin Ren. 2024. "Random Zero Sets with Local Growth Guarantees." *arXiv Preprint arXiv:2410.21931*.

Chawla, Shuchi, Anupam Gupta, and Harald Räcke. 2008. "Embeddings of Negative-Type Metrics and an Improved Approximation to Generalized Sparsest Cut." *ACM Transactions on Algorithms (TALG)* 4 (2): 1–18.

Leighton, Tom, and Satish Rao. 1999. "Multicommodity Max-Flow Min-Cut Theorems and Their Use in Designing Approximation Algorithms." *J. ACM* 46 (6): 787–832.