Understanding graphs through spectral densities

David Bindel

Department of Computer Science
Cornell University

SCAN seminar, 29 Feb 2016

Can One Hear the Shape of a Drum?

$$-\nabla^2 u = \lambda u \text{ on } \Omega, \qquad\qquad u = 0 \text{ on } \partial\Omega$$

Assume that for each n the eigenvalue λ_n for Ω_1 is equal to the eigenvalue μ_n for Ω_2. Question: Are the regions Ω_1 and Ω_2 congruent in the sense of Euclidean geometry?

I first heard the problem posed this way some ten years ago from Professor Bochner. Much more recently, when I mentioned it to Professor Bers, he said, almost at once: "You mean, if you had perfect pitch could you find the shape of a drum."

Mark Kac, American Math Monthly*, 1966*

Can One Hear the Shape of a Drum?

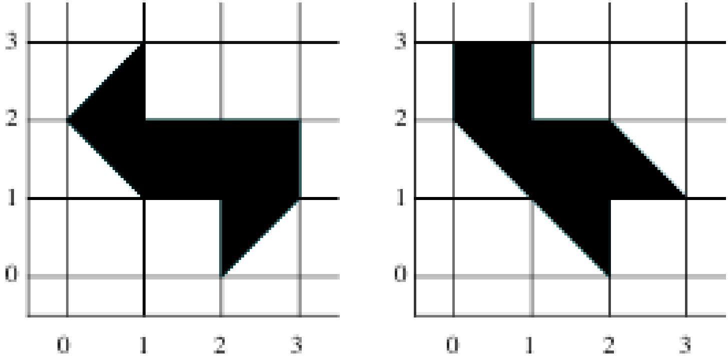

No in general (Gordon, Webb, Wolpert in 1992)
Yes with constraints (Zelditch in 2009)

What Do You Hear?

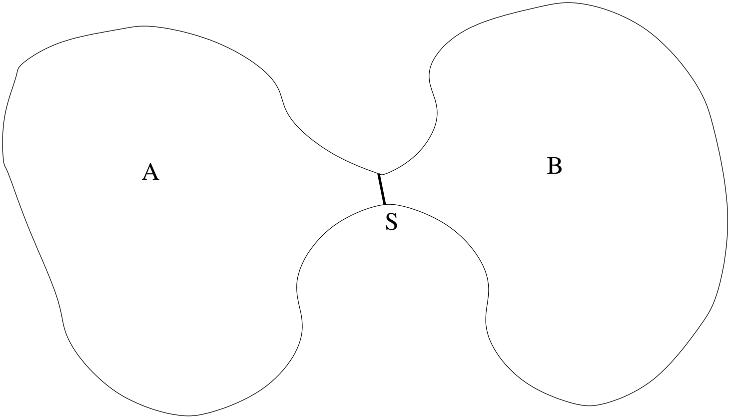

Size of bottlenecks (Cheeger inequality)

$$h \leq 2\sqrt{\lambda_2}$$

What Do You Hear?

Volume (Weyl law)

$$\lim_{x\to\infty} \frac{N(x)}{x^{d/2}} = (2\pi)^{-d} \omega_d \operatorname{vol}(\Omega), \quad N(x) = \{\text{\# eigenvalues } \le x\}$$

Can also tell lengths of geodesics for a closed Riemannian manifold.

Can One Hear the Shape of a Graph?

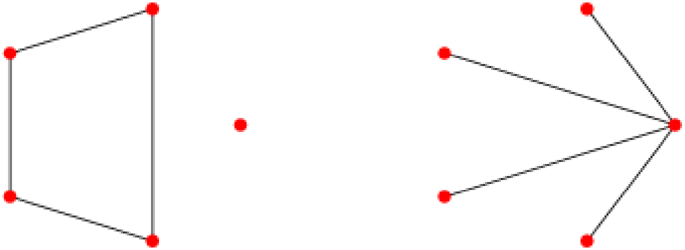

From eigenvalues of adjacency, Laplacian, normalized Laplacian?

A Bestiary of Matrices

- Adjacency matrix: A
- Laplacian matrix: $L = D - A$
- Unsigned Laplacian: $L = D + A$
- **Random walk matrix**: $P = D^{-1}A$
- **Normalized adjacency**: $\bar{A} = D^{-1/2}AD^{-1/2}$
- **Normalized Laplacian**: $\bar{L} = I - \bar{A} = D^{-1/2}LD^{-1/2}$
- Modularity matrix: $B = A - \frac{dd^T}{2n}$

All have examples of co-spectral graphs

... through spectrum uniquely identifies *quantum graphs*

What Do You Hear?

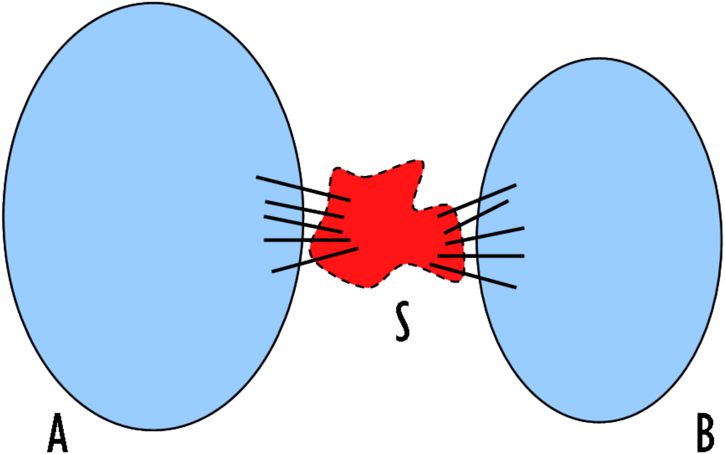

Size of separators (Cheeger inequality) – L

What Do You Hear?

What information hides in the eigenvalue distribution?

1. Discretizations of Laplacian: something like Weyl's law
2. Sparse random graphs: Wigner semicircular distribution
3. "Real" networks: less well understood

But computing all eigenvalues seems *expensive*!

Reminder: Spectral Mapping

Consider a matrix H, and let f be analytic on the spectrum. Then if $H = V\Lambda V^{-1}$,

$$f(H) = V f(\Lambda) V^{-1}.$$

(generalizes to non-diagonalizable case)

Another Perspective: Density of States

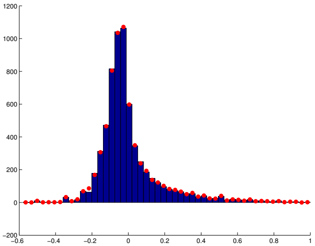

Spectra define a *generalized function* (a *density*):

$$\mathrm{tr}(f(H)) = \int f(\lambda)\mu(\lambda)\,dx = \sum_{j=k}^{N} f(\lambda_k)$$

where f is an analytic test function. Smooth out to get a picture: a *spectral histogram* or *kernel density estimate*.

Heat Kernels

DoS information equivalent to looking at the *heat kernel trace*:

$$h(s) = \mathrm{tr}(\exp(-sH)) = \mathcal{L}[\mu](s)$$

where H is a positive semi-definite operator.

$H = L \implies$
$h(s)/N$ = probability of self-return after time s from uniform start

Power Moments

DoS information equivalent to looking at the *power moments*:

$$\mathrm{tr}(H^j).$$

Has a natural interpretation for matrices associated with graphs:

- A: number of length k cycles.
- $\bar{A}$ or P: return probability for k-step random walk (times N).
- L: ??

Chebyshev Moments

Ordinary moments are not good for numerics – prefer Chebyshev:

$$d_j = T_j(A)$$

where $T_j(z) = \cos(j \cos^{-1}(z))$ is the jth Chebyshev polynomial.
Compute via three-term recurrence:

$$T_0(z) = 1$$
$$T_1(z) = z$$
$$T_{k+1}z = 2zT_k(z) - T_{k-1}(z)$$

Exploring Spectral Densities

Kernel polynomial method (see Weisse, Reviews of Modern Physics)

- Think of spectral distribution on $[-1, 1]$ as a generalized function

$$\int_{-1}^{1} \mu(x) f(x) \, dx = \frac{1}{N} \sum_{k=1}^{N} f(\lambda_k)$$

- Write $f(x) = \sum_{j=1}^{\infty} c_j T_j(x)$ and $\mu(x) = \sum_{j=1}^{\infty} d_j \phi_j(x)$, where $\int_{-1}^{1} \phi_j(x) T_k(x) \, dx = \delta_{jk}$
- Estimate $d_j = \mathrm{tr}(T_j(H))$ by stochastic methods
- Truncate series for $\mu(x)$ and filter (avoid Gibbs)

Much cheaper than computing all eigenvalues!

Stochastic Trace and Diagonal Estimation

$Z \in \mathbb{R}^n$ with independent entries, mean 0 and variance 1.

$$E[(Z \odot HZ)_i] = \sum_j h_{ij} E[Z_i Z_j] = h_{ii}$$

$$\mathrm{Var}[(Z \odot HZ)_i] = \sum_j h_{ij}^2.$$

Serves as the basis for stochastic estimation of

- Trace (Hutchinson, others; review by Toledo and Avron)
- Diagonal (Bekas, Kokiopoulou, and Saad)

Power of diagonal estimation is under-appreciated...

Diagonal Estimation and LDoS

Diagonal estimation also useful for *local* DoS $\nu_k(x)$;
in the symmetric case with $H = Q\Lambda Q^T$, have

$$\int f(x)\nu_k(x)\,dx = f(H)_{kk} = e_k^T Q f(\Lambda) Q^T e_k$$

$$\nu_k(x) = \sum_{j=1}^{n} q_{kj}^2\,\delta(x - \lambda_j)$$

DoS is sum of local densities of states:

$$\mu(x) = \sum_{k=1}^{n} \nu_k(x)$$

LDoS Information

Can compute common *centrality measures* with LDoS

- Estrada centrality: $\exp(\gamma A)_{kk}$
- Resolvent centrality: $\left[(I - \gamma \bar{A})^{-1}\right]_{kk}$

Some motifs associated with localized eigenvectors at specific values:

- Chief example: Null vectors of $\bar{A}$ supported on leaves.
- Use LDoS + topology to find motifs?

What else?

Phase Retrieval in Graph Reconstruction

Reconstruct graph from *completely resolved* LDoS at all nodes?

- Assume $H = Q \Lambda Q^T$
- No multiple eigenvalues $\implies$ know $|Q|$ and Λ
- Can we recover signs in Q?

Feels a little like phase retrieval...

Of course, we usually have noisy LDoS estimates!

Other methods

Golub and Meurant: Gauss quadrature for $\nu_k(x)$ via Lanczos

- Good: No stochastic estimation error (vs KPM)
- Bad: Separate Lanczos *per node*

Roder and Silver: Max entropy estimation

- Good: Better resolution from Chebyshev moment estimates
- Bad: More expensive computation per node

Under investigation: Hybrid approach.

Exploring Spectral Densities (with David Gleich)

- Consider spectrum of normalized Laplacian (random walk matrix)
- Approximate via KPM and compare to full eigencomputation

Things we know

- Eigenvalues in $[-1, 1]$; nonsymmetric in general
- Stability: change d edges, have

$$\lambda_{j-d} \leq \hat{\lambda}_j \leq \lambda_{j+d}$$

- kth moment = probability of return after k-step random walk
- Eigenvalue cluster near 1 $\sim$ well-separated clusters
- Eigenvalue cluster near 0 $\sim$ triangles connected by one node

What else can we “hear”?

Experimental setup

- Global DoS
  - 1000 Chebyshev moments
  - 10 probe vectors (componentwise standard normal)
  - Histogram with 50 bins
- Local DoS
  - 100 Chebyshev moments
  - 10 probe vectors (componentwise standard normal)
  - Plot smoothed density on $[-1, 1]$
  - Spectrally order nodes by density plot

Suggestions for better pics are welcome!

Erdos

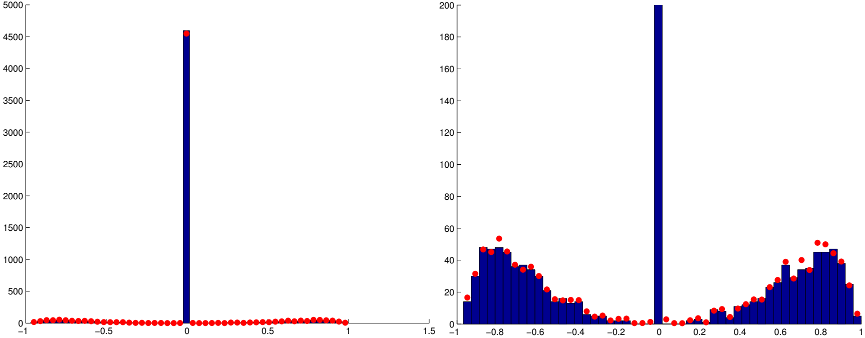
5000
4500
4000
3500
3000
2500
2000
1500
1000
500
0
−1
−0.5
0
0.5
1
1.5
200
180
160
140
120
100
80
60
40
20
0
−1
−0.8
−0.6
−0.4
−0.2
0
0.2
0.4
0.6
0.8
1

Erdos (local)

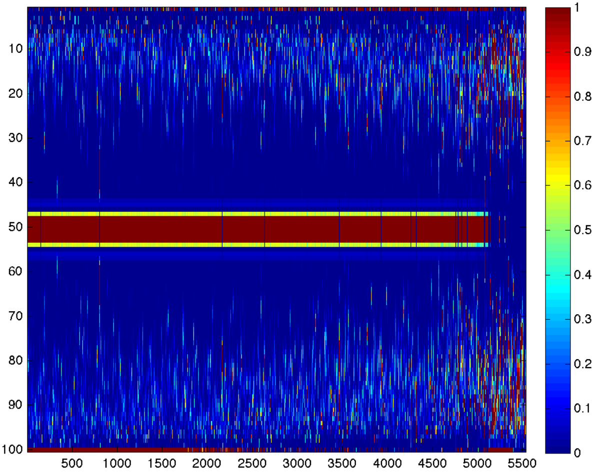

10
20
30
40
50
60
70
80
90
100
500 1000 1500 2000 2500 3000 3500 4000 4500 5000 5500
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0

Internet topology

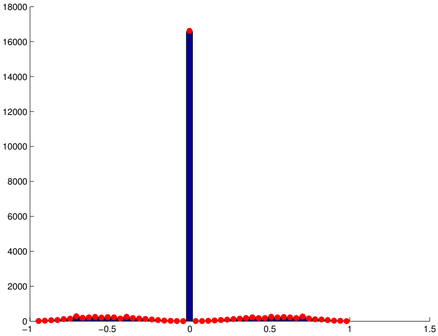

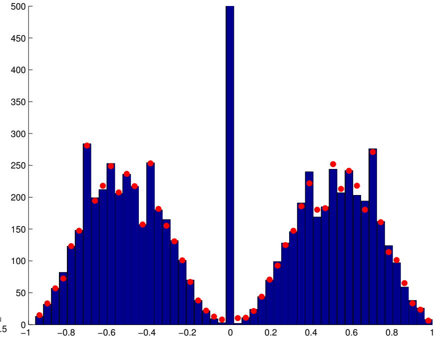

Internet topology (local)

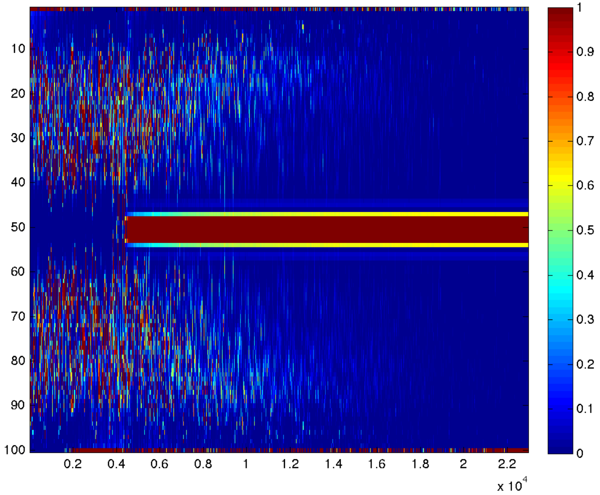

Marvel characters

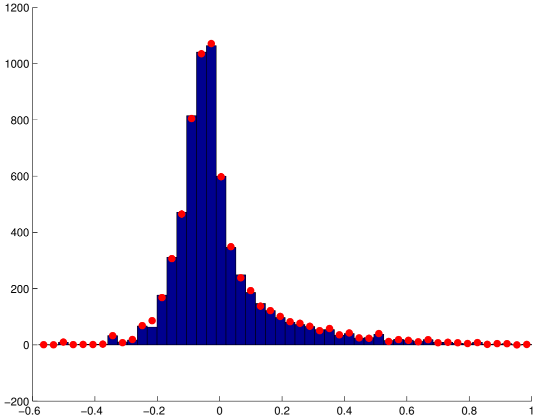

Marvel characters (local)

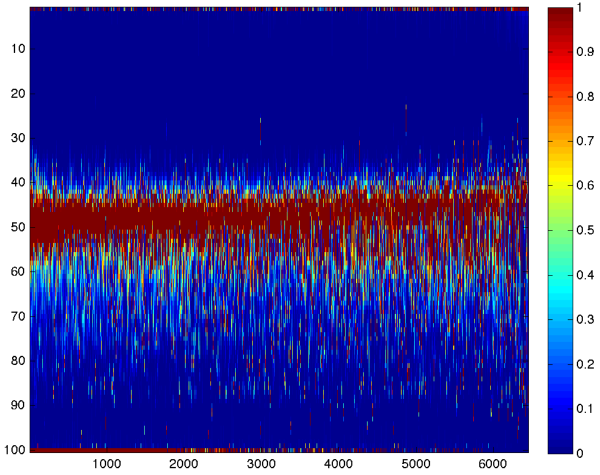

Marvel comics

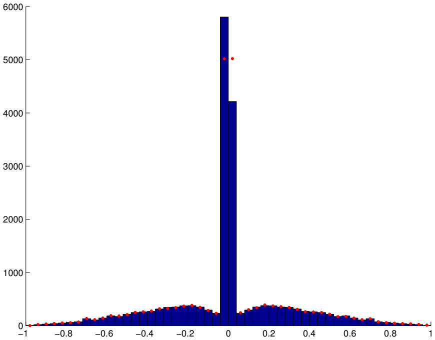

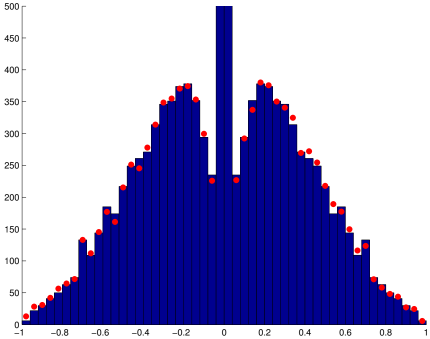

Marvel comics (local)

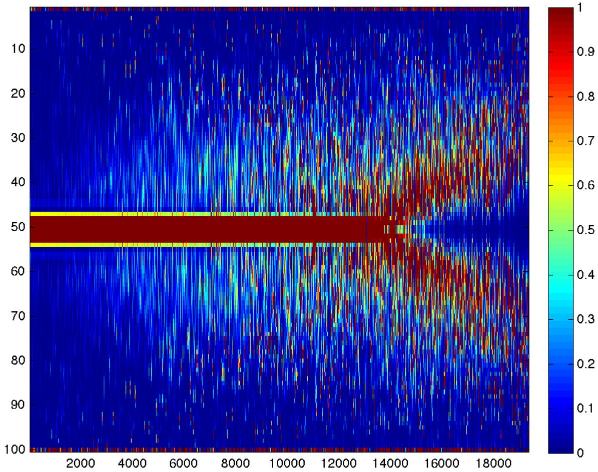

PGP

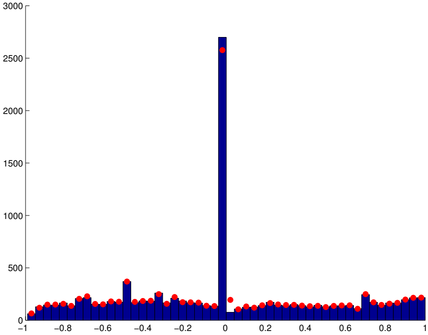

3000
2500
2000
1500
1000
500
0
−1
−0.8
−0.6
−0.4
−0.2
0
0.2
0.4
0.6
0.8
1

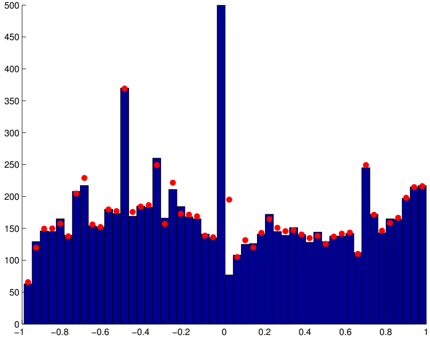

500
450
400
350
300
250
200
150
100
50
0
−1
−0.8
−0.6
−0.4
−0.2
0
0.2
0.4
0.6
0.8
1

PGP (local)

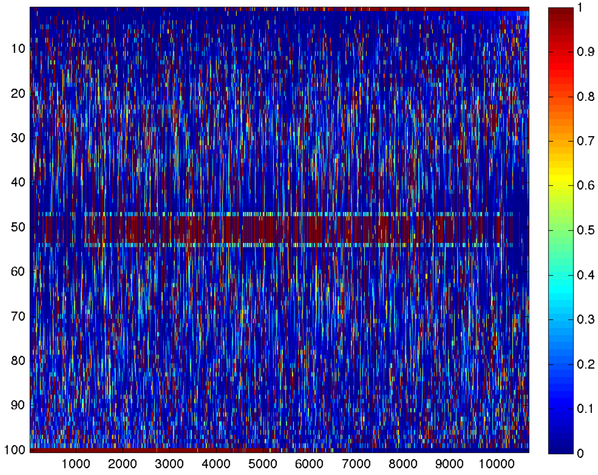

10
20
30
40
50
60
70
80
90
100
1000 2000 3000 4000 5000 6000 7000 8000 9000 10000
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0

Yeast

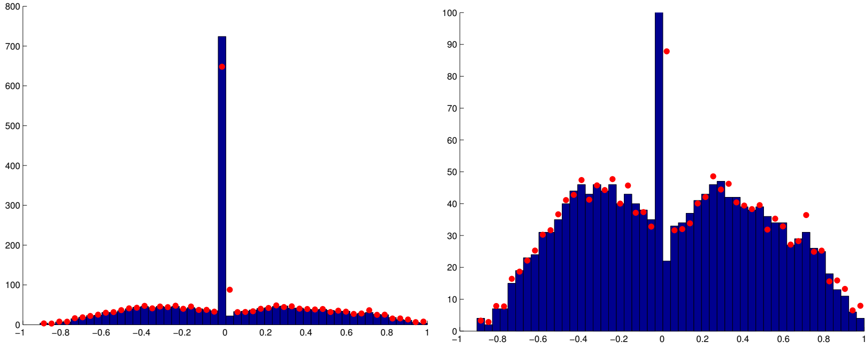
800
700
600
500
400
300
200
100
0
−1
−0.8
−0.6
−0.4
−0.2
0
0.2
0.4
0.6
0.8
1
100
90
80
70
60
50
40
30
20
10
0
−1
−0.8
−0.6
−0.4
−0.2
0
0.2
0.4
0.6
0.8
1

Yeast (local)

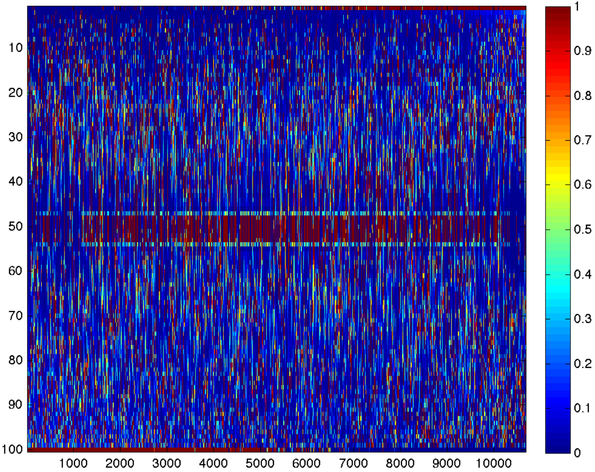
10
20
30
40
50
60
70
80
90
100
1000
2000
3000
4000
5000
6000
7000
8000
9000
10000
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0

And a few more...

Enron emails (SNAP)

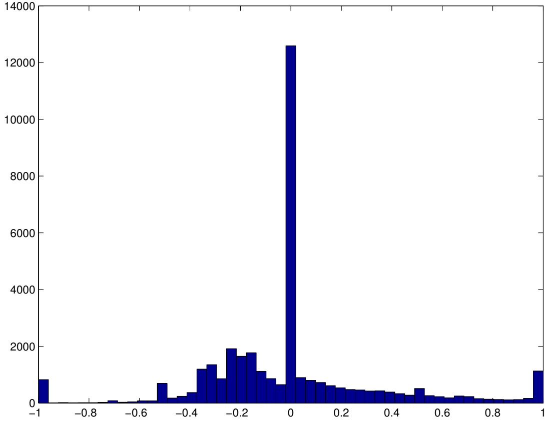

Reuters911 (Pajek)

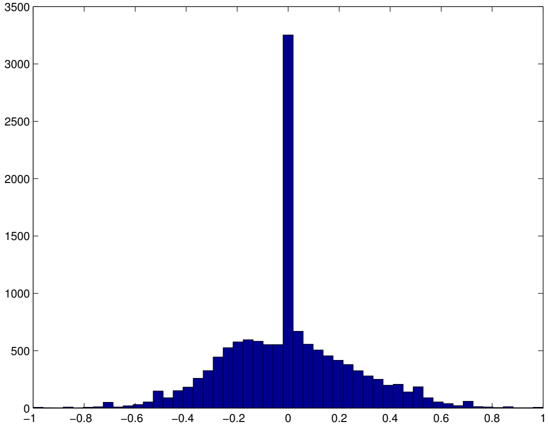

US power grid (Pajek)

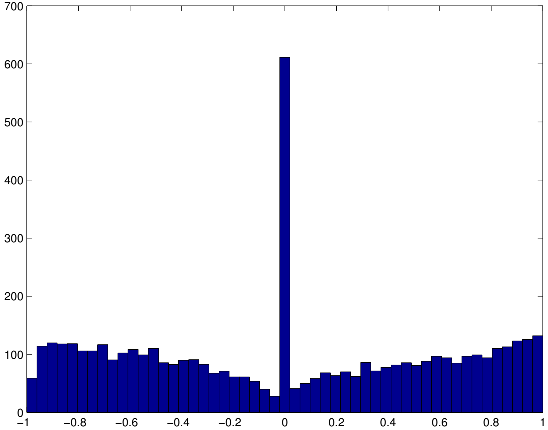

DBLP 2010 (LAW)

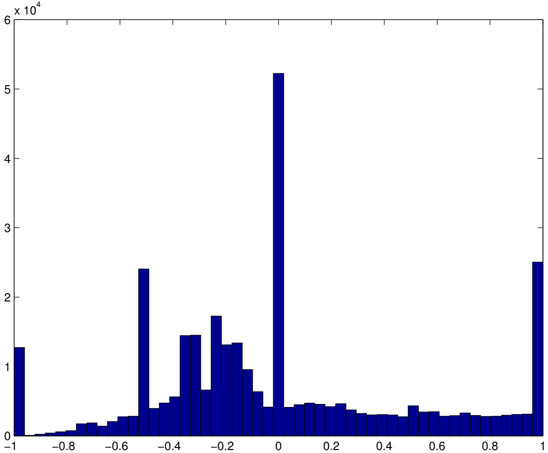

$N = 326186$, $nnz = 1615400$, 80 s (1000 moments, 10 probes)

Hollywood 2009 (LAW)

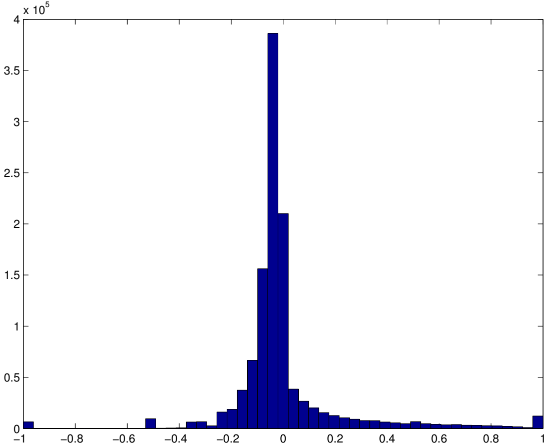

$N = 1139905$, $nnz = 113891327$, 2093 s (1000 moments, 10 probes)

What Do You Hear?