Expander Random Walks: The General Case and Limitations

Gil Cohen (Tel Aviv University)

Joint work with

Dor Minzer (MIT), **Shir Peleg** (Tel Aviv University),
Aaron Potechin (University of Chicago), **Amnon Ta-Shma** (Tel Aviv University)

June 28, 2022

Outline

Spectral expanders

Let $G = (V, E)$ be a d-regular undirected graph on n vertices.

$$(\mathbf{W}_G)_{u,v} = \begin{cases} \frac{1}{d}, & \{u, v\} \in E, \\ 0, & \text{otherwise.} \end{cases}$$

The eigenvalues of $\mathbf{W}_G$ are real, satisfying

$$-1 \le \lambda_n \le \cdots \le \lambda_2 \le \lambda_1 = 1.$$

The spectral expansion of G is given by

$$\lambda = \max\left(|\lambda_2|, |\lambda_n|\right).$$

The smaller λ is - the better. The "best" spectral expanders, dubbed Ramanujan graphs, satisfy

$$\lambda = \frac{2\sqrt{d-1}}{d}.$$

Expander random walks

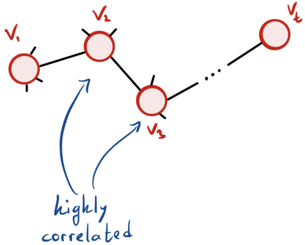

$$v_1 \sim V \qquad v_2 \sim N(v_1) \qquad \cdots \qquad v_t \sim N(v_{t-1})$$

We invest $\log(n) + (t-1) \cdot \log d \ll t \cdot \log n$ random bits in the process.

Meta question

How "random" are random walks on expanders?

Some pseudorandom properties

The Expander hitting property (Ajtai-Komlós-Szemerédi'87)

$\Pr\Big[$ 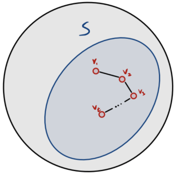 $\Big] \leq (\mu(S)+\lambda)^t$

The Expander Chernoff bound (AKS'87,Gillman'98,Healy'08)

$\Pr\Big[$ 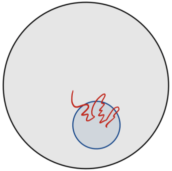 $\Big] < e^{-c(1-\lambda)\varepsilon^2 t}$

Formalizing the question

For $G = (V, E)$ and $\mathrm{val} : V \to \{\pm 1\}$ let

$$\mathrm{RW}_{G,\mathrm{val}} \in \{\pm 1\}^t$$

be the distribution $(\mathrm{val}(v_1), \ldots, \mathrm{val}(v_t))$ where $v_1, \ldots, v_t$ is a random walk on G.

Given $f : \{\pm 1\}^t \to \{\pm 1\}$ define

$$\mathcal{E}_{G,\mathrm{val}}(f) = \left| \mathbb{E}\left[f(\mathrm{RW}_{G,\mathrm{val}}) \right] - \mathbb{E}\left[f(\mathrm{val}(V)^t) \right] \right| .$$

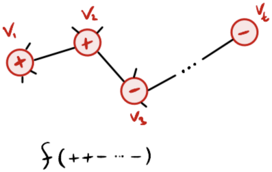

Formalizing the question

Recall

$$\mathcal{E}_{G,\mathrm{val}}(f) = \left| \mathbb{E}\left[f(\mathrm{RW}_{G,\mathrm{val}}) \right] - \mathbb{E}\left[f(\mathrm{val}(V)^t) \right] \right|.$$

Definition

For λ, μ and $f : \{\pm 1\}^t \to \{\pm 1\}$, define

$$\mathcal{E}_{\lambda,\mu}(f) = \sup_{G,\mathrm{val}} \mathcal{E}_{G,\mathrm{val}}(f),$$

where

- $G = (V, E)$ ranges over all λ-spectral expanders; and
- val ranges over all valuations with $\mathbb{E}\left[\mathrm{val}(V)\right] = \mu$.

Pseudorandom properties revisited

The Expander hitting property (Ajtai-Komlós-Szemerédi'87)

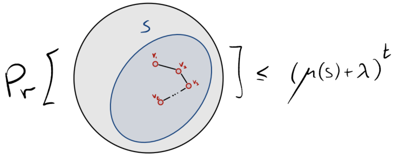

$$\mathcal{E}_{\lambda,\mu}(\mathrm{AND}_t) \leq (\mu + \lambda)^t$$

The Expander Chernoff bound (AKS'87,Gillman'98,Healy'08)

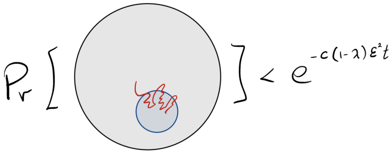

$$\mathcal{E}_{\lambda,\mu}\left(\mathbf{1}_{[(\mu-\varepsilon)t,(\mu+\varepsilon)t]}\right) \leq e^{-c(1-\lambda)\varepsilon^2 t}$$

Pseudorandom properties revisited

Expanders as parity samplers ([Ta-Shma'17 (see also Alon'93, Wigderson - Rozenman'04)])

$$\mathcal{E}_{\lambda,\mu}(\text{Parity}) \leq (\mu + 2\lambda)^{t/2}$$

A crucial ingredient in Ta-Shma's construction of near-optimal small bias sets (STOC 2017).

Other test functions?

What about other test functions?

This question was raised by Guruswami and Kumar (ITCS 2021) and independently by Cohen, Peri and Ta-Shma (STOC 2021).

One can consider

1. Symmetric functions
2. $\mathbf{AC^0}$ circuits
3. Bounded space test functions
4. Low query complexity

$\vdots$

Outline

Prior work

Theorem (Cohen-Peri-Ta Shma (CPTS))

For every symmetric function $f : \{\pm 1\}^t \to \{\pm 1\}$,

$$\mathcal{E}_\lambda(f) = O(\lambda \cdot \log^{3/2}(1/\lambda)).$$

For several specific symmetric functions a bound that vanishes with t has been obtained, e.g.,

$$\mathcal{E}_\lambda(\mathbf{1}_w) = O\left(\frac{\lambda}{\sqrt{t}}\right) \qquad \forall w \in [-t, t]$$

$$\mathcal{E}_\lambda(\text{Majority}) = O\left(\frac{\lambda^2}{\sqrt{t}}\right)$$

Prior work

Theorem (CPTS)

For every $f : \{\pm 1\}^t \to \{\pm 1\}$ that is computable by a size-s depth-d circuit,

$$\mathcal{E}_\lambda(f) = O\left(\sqrt{\lambda} \cdot (\log s)^{2(d-1)}\right).$$

This can be seen as an analog to Braverman's celebrated result (J. ACM 2010) which states that every

$$k = \left(\log \frac{s}{\varepsilon}\right)^{O(d^2)}$$

wise independent distribution ε-fools every function that is computable by a size-s depth-d circuit (see also Tal; CCC 2017).

Outline

Our contribution

The work of CPTS left four open problems:

1. $\mu \neq 0$
2. Is the $\log^{3/2}(1/\lambda)$ inherent?
3. Is $\mathcal{E}_\lambda(f) \xrightarrow[t\to\infty]{} 0$ for all symmetric functions?
4. Does $\lambda = \Omega(1)$ suffice to fool $\mathbf{AC^0}$?

In this work we resolve all four problems.

Our contribution

Theorem

For every symmetric function $f : \{\pm 1\}^t \to \{\pm 1\}$ and every $\mu \in (-1,1)$,

$$\mathcal{E}_{\lambda,\mu}(f) = O\left(\frac{\lambda}{\sqrt{1-|\mu|}}\right).$$

This resolves Items 1,2.

1. Holds for every μ; and
2. no $\log^{3/2}(1/\lambda)$ factor.

Our contribution

Theorem

For every t, set
$$w = \frac{t - \sqrt{t}}{2}.$$
Then,
$$\mathcal{E}_\lambda(\mathbf{1}_{>w}) = \Omega(\lambda).$$

This resolves Item 3.

Our contribution

As for Item 4, we prove that the CPTS bound for $\mathbf{AC^0}$ is tight up to a quartic power.

Theorem

For every constant depth $d \geq 3$ there exists a function $f : \{\pm 1\}^t \to \{\pm 1\}$ computable by a depth-d poly(t)-size circuit s.t.

$$\mathcal{E}_\lambda(f) = \Omega(1),$$

where

$$\lambda = \Theta\left(\frac{1}{\log^{d-2} t}\right).$$

Outline

CPTS's approach

CPTS's Approach

1. Bound $\mathcal{E}_\lambda(\chi_S) = |\mathbb{E}[\chi_S(\mathrm{RW})]|$ for a general character
$$\chi_S(x_1,\ldots,x_t) = \prod_{i\in S} x_i$$
with $\emptyset \neq S \subseteq [t]$.
2. Expand f in the Fourier basis
$$f(x) = \sum_{S\subseteq[t]} \widehat{f}(S)\chi_S(x).$$
3. Conclude that
$$\mathcal{E}_\lambda(f) \le \sum_{\emptyset\neq S\subseteq[t]} \left|\widehat{f}(S)\right| \mathcal{E}_\lambda(\chi_S).$$

Degree 2 characters

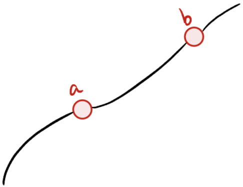

$$\mathcal{E}_\lambda(x_a x_b) \le \lambda^{b-a}.$$

Degree 3 characters

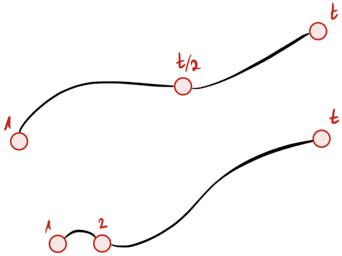

Test your intuition.

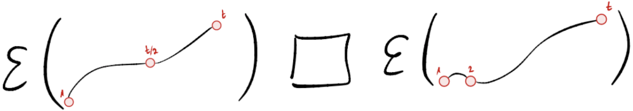

Degree 3 characters

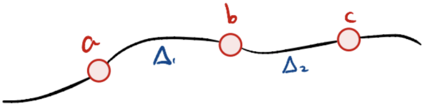

$$\mathcal{E}_\lambda(\chi_{a,b,c}) \le \lambda^{\Delta_1+\Delta_2}.$$

Degree 4 characters

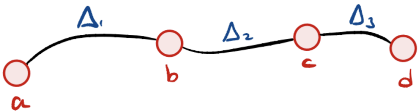

$$\mathcal{E}_\lambda(\chi_{a,b,c,d}) \leq \lambda^{\Delta_1+\Delta_3}.$$

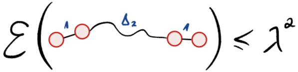

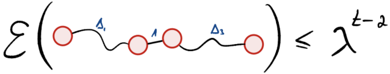

Outline

Our approach

Lemma

Let $G = (V, E)$ be a Cayley graph on the Boolean hypercube and consider a labelling of V by an eigenvector corresponding to λ_2.

Given $f : \{\pm 1\}^t \to \{\pm 1\}$ define $g : \{\pm 1\}^{2t} \to \{\pm 1\}$ by

$$g(x_1, \ldots, x_{2t}) = f(x_1 \cdot x_2,\ x_3 \cdot x_4,\ \ldots,\ x_{2t-1} \cdot x_{2t}).$$

Then,

$$\mathbb{E}[g(\mathrm{RW})] = (T_\lambda f)(\mathbf{1}).$$

Using this, we prove the tightness result for symmetric functions (Item 4). The result on $\mathbf{AC^0}$ circuits follows by applying the lemma to an iterated Tribes-like function.

Followup work and open problems

Follow-up work.

Golowich and Vadhan (CCC 2022; to appear) continued this line of work and obtained, among other results, bounds for non-binary labellings, and better dependence on μ.

Open problems.

- Tightness for $\mathbf{AC}^0$ circuits
- Lower bound for constant degree expanders
- Applications?

Thanks!