Expanders:

Definition

- bipartite
- n vertices, regular degree d
- $|\Gamma(S)| \geq (1 + c(1 - 2|S|/n))|S|$

factor c more neighbors, at least until S near $n/2$.
Take random walk on (n, d, c) expander with constant c

- add self loops (with probability 1/2 to deal with periodicity.
- uniform stationary distribution
- lemma: second eigenvalue $1 - O(1/d)$

$$\lambda_2 \leq 1 - \frac{c^2}{d(2048 + 4c^2)}$$

- Intuition on convergence: because neighborhoods grow, position becomes unpredictable very fast.
- proof: messy math

Deduce: mixing time in expander is $O(\log n)$ to get ϵ r.p.d. (since $\pi_i = 1/n$)
Converse theorem: if $\lambda_2 \leq 1 - \epsilon$, get expander with

$$c \geq 4(\epsilon - \epsilon^2)$$

Walks that mix fast are on expanders.
Gabber-Galil expanders:

- Do expanders exist? Yes! proof: probabilistic method.
- But in this case, can do better deterministically.
  - Gabber Galil expanders.
  - Let $n = 2m^2$. Vertices are (x, y) where $x, y \in Z_m$ (one set per side)
  - 5 neighbors: $(x, y), (x, x + y), (x, x + y + 1), (x + y, y), (x + y + 1, y)$ (add mod m)
  - or 7 neighbors of similar form.
- Theorem: this $d = 5$ graph has $c = (2 - \sqrt{3})/4$, degree 7 has twice the expansion.
- in other words, c and d are constant.
- meaning $\lambda_2 = 1 - \epsilon$ for some **constant** ϵ
- So random walks on this expander mix *very* fast: for polynomially small r.p.d., $O(\log n)$ steps of random walk suffice.
- Note also that n can be huge, since only need to store one vertex ($O(\log n)$ bits).

Application: conserving randomness.

- Consider an BPP algorithm (gives right answer with probability 99/100 (constant irrelevant) using n bits.
- t independent trials with majority rule reduce failure probability to $2^{-O(t)}$ (chernoff), but need tn bits
- in case of RP, used 2-point sampling to get error $O(1/t)$ with $2n$ bits and t trials.
- Use walk instead.
  - vertices are $N = 2^n$ (n-bit) random strings for algorithm.
  - edges as degree-7 expander
  - only 1/100 of vertices are bad.
  - what is probability majority of time spent there?
  - in limit, spend 1/100 of time there
  - how fast converge to limit? How long must we run?
  - Power the markov chain so $\lambda_2^\beta \leq 1/10$ (constant number of steps)
  - use random seeds encountered every β steps.
- number of bits needed:
  - $O(n)$ for stationary starting point
  - 3β more per trial,
- Theorem: after $7k$ samples, probability majority wrong is $1/2^k$. So error $1/2^n$ with $O(n)$ bits! (compare to naive)
  - Let B be powered transition matrix
  - let $p^{(i)}$ be distribution of sample i, namely $p^0 B^i$
  - Let W be indicator **matrix** for good witnesses, namely 1 at diagonal i if i is a witness. $\overline{W}$ completmentary set $I - W$.
  - $\|p^i W\|_1$ is probability p^i is witness set. similar for nonwitness.
  - Consider a sequence of $7k$ results "witness or not"
  - represent as matrices $S = (S_1, \ldots, S_{7k}) \in \{W, \overline{W}\}^{7k}$
  - claim

    $$\Pr[S] = \|p^{(0)}(BS_1)(BS_2)\cdots(BS_{7k})\|_1.$$

    (draw layered graph. sums prob. of paths through correct sequence of witness/nonwitness)
  - Use 2-norm since easier to work with.

    * Note $\|A\|_1 \leq \sqrt{N}\|A\|_2$
    * For fixed sum of values, minimize sum of square by setting all equal
    * ie, for sum α, set all equal to α/N
    * 2-norm $\alpha/\sqrt{N}$
  - defer: $\|pBW\|_2 \leq \|p\|_2$ and $\|pB\overline{W}\|_2 \leq \frac{1}{5}\|p\|_2$
  - deduce if more than $7k/2$ bad witnesses,

$$\begin{aligned} \|p^0 \prod BS_i\|_1 &\leq \sqrt{N}\|p^0 \prod BS_i\| \\ &\leq \sqrt{N}(\frac{1}{5})^{7k/2}\|p^0\| \\ &= (\frac{1}{5})^{7k/2} \end{aligned}$$

    (since $\|p_0\| = 1/\sqrt{N}$)
  - At same time, only 2^{7k} bad sequences, so error prob. $2^{7k}5^{-7k/2} \leq 2^{-k}$

- proof of lemma:
  - write $p = \sum a_i e_i$ with $e_1 = \pi$
  - obviously $\|pBW\| \leq \|pB\|$ since W just zeros some stuff out.
  - But $\|pB\| = \sqrt{\sum a_i^2 \lambda_i^2} \leq \sum a_i^2 = \|p\|$
  - Now write $p = a_1\pi + y$ where $y \cdot \pi = 0$
  - argue that $\|\pi B\overline{W}\| \leq \|\pi\|/10 \leq \|p\|/10$ and $\|yB\overline{W}\| \leq \|y\|/10 \leq \|p\|/10$, done by pythagorous
  - First π:
    * recall $\pi B = \pi$ is uniform vector, all coords $1/N$, so norm $1/\sqrt{N}$
    * $\overline{W}$ has only $1/100$ of coordintes nonzero, so
    * $\|e_1\overline{W}\| = \sqrt{(N/100)(1/N)} = 1/10$
  - Now y: just note $\|yB\| \leq \|y\|/10$ since $\lambda_2 \leq 1/10$. Then $\overline{W}$ zeros out.
  - summary: π part likely to be in good witness set, y part unlikely to be relevant.