21-Apr-2025 Fast Laplacian Solvers - $\|x_t\|_L \le \varepsilon \|t\|_L$ $t \in \mathbb{R}^n$ m edges

Last lecture: Time $\le O(m^{4/3} \ln 1/\varepsilon)$

Best known: $\le O(m (\lg\lg m)^c \ln 1/\varepsilon)$ $\tilde{O}$ = ignore terms of smaller order

Today: $\le \tilde{O}(m (\lg m)^2 \ln 1/\varepsilon)$ than those written here, like $(\lg\lg m)$

Computing eigenvectors & eigenvalues: for largest pick x at random, repeat

$x = Lx / \|x\|$ if $x_0 = \sum c_i \psi_i$ $x_t \sim \sum c_i (\lambda_i)^t \psi_i$ so, largest λ_n dominates

iterations to approach ψ_n depends on $\frac{\lambda_n - \lambda_{n-1}}{\lambda_n}$

λ_2? Could try $x = (\lambda_n I - L) x / \|x\|$. But takes a long time to distinguish close eigvals

Say, to tell apart $\lambda_2 \sim \frac{1}{n}$ from $\lambda_3 \sim \frac{4}{n}$

Ideal would be $x = L^+ x / \|x\|$. Largest eigval is $\frac{1}{\lambda_2}$ gap is $\frac{1}{\lambda_2} - \frac{1}{\lambda_3} \cdots$

Not clear approx solver works here. It does if solver is linear.

If can compute Zb so that $Z \approx L^+$. $\| ZL - \Pi \| \le \varepsilon$ $\left(\lambda_{max}(Z)\right)^{-1} \in (1 \pm \varepsilon)\lambda_2(L)$

Need Algorithm like Chebyshev, or preconditioned Chebyshev,

applying a fixed polynomial. Like $Z = P_t(L)$ or $P_t(B^{-1}L)$ B a precon.

Need to fix # of iterations.

CG is non-linear. poly depends on b.

Elimination. Produce $U^T U = L$ U is upper Δ-ar

Or, for Schur complement onto B, eliminating S, get $L = U^T \begin{array}{cc} S & B \\ \downarrow & \downarrow \end{array} \begin{pmatrix} I & \\ & L_B \end{pmatrix} U$

# of off-diagonal entries in row a of U = degree of a when eliminate it.

As U is upper Δ, can apply U or U^{-1} in time $O(\overbrace{nz(U)}^{\text{number non-zeros}})$

$$L^+ = U^{-T} \begin{pmatrix} I & \\ & L_B^+ \end{pmatrix} U^{-1}$$

Ultrasparsifier: tree plus edges. $G=(V,E,w)$ $H=(V,F,u)$ $F \subseteq E$ $|F| \le n-1+k$

Can get $\kappa(L_G, L_H) \le O(n/k^2)$

Lem If H has $n-1+k$ edges, can elim vertices of degree 1 or 2 until end up with L_B having $\le 4k$ vertices and $\le 5k$ edges

How solve systems in L_B? Recursively.

Towards ultra-sparsifier by random sampling

Recall if include edges with prob $p_{a,b} \ge \min\left(1, \frac{4 \ln n}{\varepsilon^2} w_{a,b} \mathrm{Reff}_G(a,b)\right)$ get ε-approx with $\mathbb{E}$ # edges $\sum p_{a,b}$

Want fewer edges. OK with worse approx

Let T be a low-stretch spanning tree of G. $\hat{G} = G + sT$ s sort of big $\approx (\ln n)^4$

Sample to sparsify $\hat{G}$. Most edges chosen will be in T. Call result H

$$\frac{1}{|E|}\sum_{(a,b)\in E} \mathrm{stretch}_T(a,b) = \frac{1}{|E|}\sum_{a,b} \mathrm{Reff}_T(a,b) \lesssim m \ln n$$

for $a,b \notin T$, $\mathrm{Reff}_{\hat{G}}(a,b) \le \mathrm{Reff}_{ST}(a,b) = \frac{1}{S}\mathrm{Reff}_T(a,b)$ set $\varepsilon = \frac{1}{2}$

So, $\mathbb{E}$ # non-tree edges $= \sum_{(a,b)\notin T} p_{a,b} \le \frac{4 \ln n}{S\varepsilon^2} \sum_{(a,b)\notin T} \mathrm{Reff}_T(a,b) \le \frac{4 m \ln^2 n}{S\varepsilon^2} \le \frac{16 m \ln^2 n}{S}$

$$\kappa(L_G, L_H) \le \kappa(L_G, L_{\hat{G}}) \cdot \kappa(L_{\hat{G}}, L_H) \le S\,\frac{1+\frac{1}{2}}{1-\frac{1}{2}} = 3S$$

Will need to do $O(\sqrt{S})$ solves in L_H, using precon-Cheby.

So, want # off-tree edges in $H \le \frac{m}{\sqrt{S}}$. Then cost of L_H^{-1} is like a multiply by L_G

Need $\frac{16 m \ln^2 n}{S} \le \frac{m}{\sqrt{S}} \longleftrightarrow 16 \ln^2 n \le \sqrt{S}$ $S \sim 256 (\lg n)^4$ $\kappa \le O(\ln^2 n)$

L_H is precon for L_G $L_H = U^T \begin{pmatrix} I & \\ & L_B \end{pmatrix} U$ solve sys in L_B recursively, by lin op Z s.t. $\| L_B Z - \Pi \| \le \varepsilon$, so

$$U^T \begin{pmatrix} I & \\ & Z^+ \end{pmatrix} U \approx L_H.$$

Total time $\tilde{O}(m \lg^2 n)$

Big picture. let $G = G_0$ Solve system in G_i by solving in smaller G_{i+1}

1. T_i low stretch tree of G_i
2. $\hat{G}_i = G_i + sT_i$
3. Reff sample $(\hat{G}_i) = H_i$
4. elim deg 1 & 2 vertices of H_i, are left with G_{i+1}. smaller.

Can save a log by keeping the same tree throughout KMP2