1. Definition of t-spanner: A subgraph H of G s.t. $d_H(u,v) \le t\, d_G(u,v)$.
2. Greedy centralized algorithm: Iterate over edges $\{u,v\}$ in order.
   Add $\{u,v\}$ to H if $d_H(u,v) \ge 2k$.
   ↳ This is a $(2k-1)$-spanner.
   ↳ No cycles of length $\le 2k$ (i.e. girth $\ge 2k+1$) $\rightarrow \le n^{1+1/k} + n = O(n^{1+1/k})$ edges (Moore bound)
3. Erdős Girth Conjecture: There exist arbitrarily large graphs with girth $\ge 2k+1$ and $\Omega(n^{1+1/k})$ edges.
   ↳ Implies greedy spanner has optimal size.

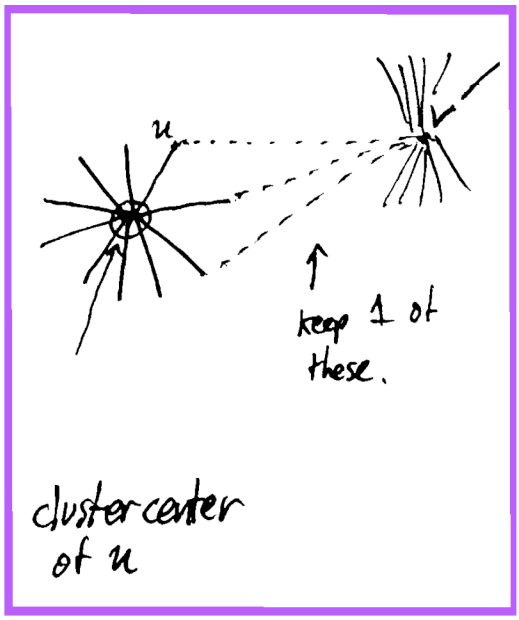
4. Today: LCA for a 3-spanner with $\tilde{O}(n^{3/2})$ edges, using $\tilde{O}(n^{1/2})$ probes per query.
   ↳ Probes: $\deg(v)$,
   i^{th} neighbor of v,
   $(u,v) \rightarrow$ Are u and v adjacent? And if so, return i and j s.t. v is the i^{th} neighbor of u and u is the j^{th} neighbor of v.
   ↳ Note: There is an $\Omega(n^{1/2})$ probe lower bound for this problem.
5. Algorithm Design.

- Divide vertices into "low-degree" ($\deg \le \sqrt{n}$) and "high-degree" ($\deg > \sqrt{n}$).
  Keep all edges incident to low-degree vertices. (only $n^{3/2}$)
- "Pick $\sqrt{n}\log n$ centers at random" (actually, pick vertices independently w.p. $\frac{\sqrt{n}\log n}{n}$).

For each vertex v, keep exactly one edge connecting v to each adjacent cluster (star surrounding a center).
↳ Every high-degree vertex is adjacent to ≥ 1 center; for now assume it's exactly 1.
↳ Use random edge ranks to tie break.
↳ #Probes = (#centers) to find cluster of u
+ $\deg(v)$ to determine other edges connecting v to cluster of u.

$\le \sqrt{n}\log n + n$
(bad!)

↳ #Edges = (vertices v) · (clusters)
$= n \cdot \sqrt{n}\log n = \tilde{O}(n^{3/2})$ (good!)



- Fix: Divide edges of v into buckets of size $\sqrt{n}$.
  For each vertex v, keep exactly one edge
  in each bucket connecting v to each adjacent center.
  ↳ #Probes = (#centers) + (bucket size) $= \sqrt{n}\log n + \sqrt{n}$
  (good!)
  ↳ #Edges = (vertices v) · (buckets of v) · (clusters)
  $\leq n \cdot \sqrt{n} \cdot \sqrt{n}\log n$ (bad!)

- Fix 2: Divide high-degree vertices further into classes by degree.
  $V_i = \{v : \deg(v) \in [2^i\sqrt{n}, 2^{i+1}\sqrt{n})\}$ for $i = 0..\log\sqrt{n}$.

- Use a separate set of centers S_i for each V_i.
  $\mathbb{E}|S_i| = 2^{-i}\sqrt{n}\log n$.
  ↳ Each vertex in V_i is adjacent to ≥ 1 and $\leq O(\log n)$ centers in S_i, whp.

- For each vertex v, keep 1 edge in each bucket connecting v to
  each cluster of equal or greater class #.
  ↳ #Probes = (#centers) + (bucket size)·(# of clusters u is in)
  $= \tilde{O}(\sqrt{n}) + \sqrt{n} \cdot \log n$ (good!)

  ↳ #Edges = For each vertex v, (buckets of v)·(# of clusters of class $\geq$ that of v)
  $= \sum_{i=0}^{\log\sqrt{n}} |V_i| \cdot 2^{i+1} \cdot 2^{-i+1}\sqrt{n}\log n$
  $= \tilde{O}(n^{3/2})$ (good!)

So we are done.