

Finding a Bounded-Degree **Expander** Inside a Dense One

HAL-02002377
archives-ouvertes.fr

Emanuele Natale



Joint work with L. Becchetti, A. Clementi,
F. Pasquale and L. Trevisan



13 March 2019

Outline

- Definitions: Graph Expansion
- Motivation for this work
- Our Results
- Crash Course on Encoding Arguments
- Some Proof Ideas

Graph Expansion I

What is a good measure of *connectedness* for a set of nodes S ?

Graph Expansion I

What is a good measure of *connectedness* for a set of nodes S ?

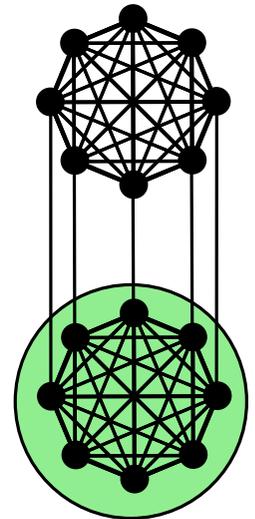
- Attempt 1. Number of edges going out of S :
$$e(S, V - S) = |\{(u, v) | u \in S, v \in V - S\}|$$

Graph Expansion I

What is a good measure of *connectedness* for a set of nodes S ?

- Attempt 1. Number of edges going out of S :
$$e(S, V - S) = |\{(u, v) | u \in S, v \in V - S\}|$$

Problem: big sets are better than small ones



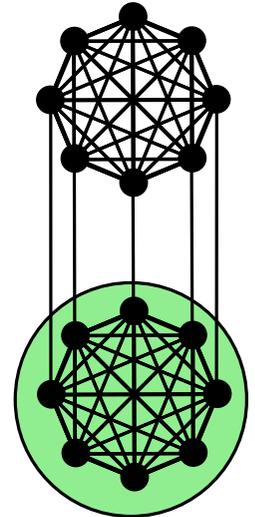
Graph Expansion I

What is a good measure of *connectedness* for a set of nodes S ?

- Attempt 1. Number of edges going out of S :
$$e(S, V - S) = |\{(u, v) | u \in S, v \in V - S\}|$$

Problem: big sets are better than small ones

- Attempt 2. We also divide by the sum of its degrees
degrees $vol(S) = \sum_{u \in S} d_u$: $\frac{e(S, V - S)}{vol(S)}$



Graph Expansion I

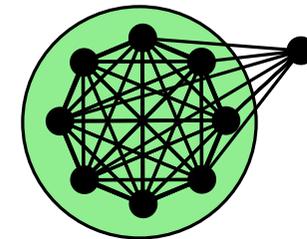
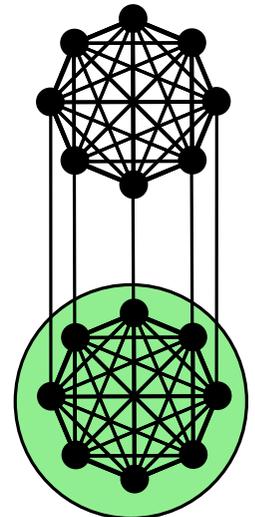
What is a good measure of *connectedness* for a set of nodes S ?

- Attempt 1. Number of edges going out of S :
$$e(S, V - S) = |\{(u, v) | u \in S, v \in V - S\}|$$

Problem: big sets are better than small ones

- Attempt 2. We also divide by the sum of its degrees
degrees $vol(S) = \sum_{u \in S} d_u$: $\frac{e(S, V - S)}{vol(S)}$

Problem: Very big sets have big $vol(S)$



Graph Expansion I

What is a good measure of *connectedness* for a set of nodes S ?

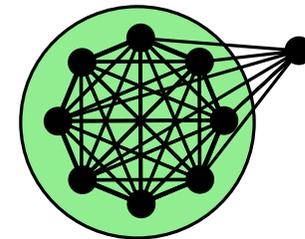
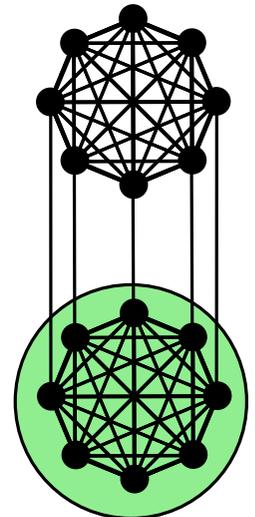
- Attempt 1. Number of edges going out of S :
$$e(S, V - S) = |\{(u, v) | u \in S, v \in V - S\}|$$

Problem: big sets are better than small ones

- Attempt 2. We also divide by the sum of its degrees
degrees $vol(S) = \sum_{u \in S} d_u$: $\frac{e(S, V - S)}{vol(S)}$

Problem: Very big sets have big $vol(S)$

- Attempt 3. We consider the “worst” between S and $V - S$:
$$\frac{e(S, V - S)}{\min\{vol(S), vol(V - S)\}}$$



Graph Expansion I

What is a good measure of *connectedness* for a set of nodes S ?

- Attempt 1. Number of edges going out of S :
$$e(S, V - S) = |\{(u, v) | u \in S, v \in V - S\}|$$

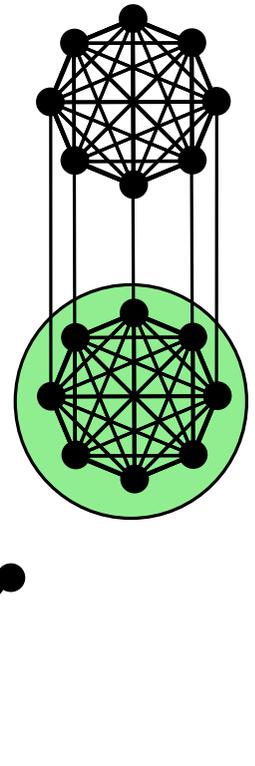
Problem: big sets are better than small ones

- Attempt 2. We also divide by the sum of its degrees $vol(S) = \sum_{u \in S} d_u$: $\frac{e(S, V - S)}{vol(S)}$

Problem: Very big sets have big $vol(S)$

- Attempt 3. We consider the “worst” between S and $V - S$: $\frac{e(S, V - S)}{\min\{vol(S), vol(V - S)\}}$ → **conductance**

Examples: Giakkoupis et al. JACM'18 and Giakkoupis SODA'14 for another expansion measure.



Graph Expansion II

In regular graphs $\frac{e(S, V-S)}{\min\{\text{vol}(S), \text{vol}(V-S)\}}$ is equivalent to $\phi(S) = \frac{e(S, V-S)}{\text{vol}(S)}$ assuming $|S| \leq \frac{n}{2}$

Graph Expansion II

In regular graphs $\frac{e(S, V-S)}{\min\{\text{vol}(S), \text{vol}(V-S)\}}$ is equivalent to

$$\phi(S) = \frac{e(S, V-S)}{\text{vol}(S)} \text{ assuming } S \leq \frac{n}{2}$$

Interpretation. In regular graphs, $\phi(S) = \Pr(\text{random walk on random node of } S \text{ exits it})$

Graph Expansion II

In regular graphs $\frac{e(S, V-S)}{\min\{\text{vol}(S), \text{vol}(V-S)\}}$ is equivalent to

$$\phi(S) = \frac{e(S, V-S)}{\text{vol}(S)} \text{ assuming } S \leq \frac{n}{2}$$

Interpretation. In regular graphs, $\phi(S) = \Pr(\text{random walk on random node of } S \text{ exits it})$

Graph G is ϵ -expander if $\min_S \phi(S) \geq \epsilon$

Graph Expansion II

In regular graphs $\frac{e(S, V-S)}{\min\{\text{vol}(S), \text{vol}(V-S)\}}$ is equivalent to

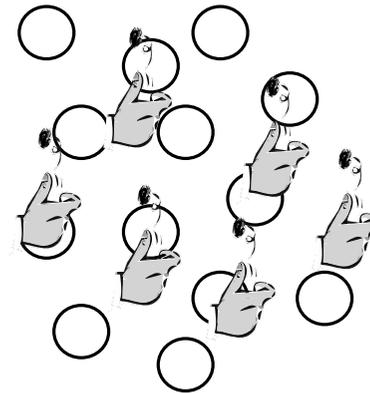
$$\phi(S) = \frac{e(S, V-S)}{\text{vol}(S)} \text{ assuming } S \leq \frac{n}{2}$$

Interpretation. In regular graphs, $\phi(S) = \Pr(\text{random walk on random node of } S \text{ exits it})$

Graph G is ϵ -expander if $\min_S \phi(S) \geq \epsilon$

Example:

In an Erős-Rényi graph $G_{n,p}$, include each edge with prob p .



Graph Expansion II

In regular graphs $\frac{e(S, V-S)}{\min\{\text{vol}(S), \text{vol}(V-S)\}}$ is equivalent to $\phi(S) = \frac{e(S, V-S)}{\text{vol}(S)}$ assuming $S \leq \frac{n}{2}$

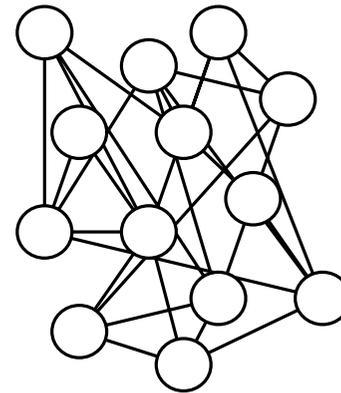
Interpretation. In regular graphs, $\phi(S) = \Pr(\text{random walk on random node of } S \text{ exits it})$

Graph G is ϵ -expander if $\min_S \phi(S) \geq \epsilon$

Example:

In an Erős-Rényi graph $G_{n,p}$, include each edge with prob p .

For any $p \gg \frac{\log n}{n}$, they are good expanders with high probability.



Expander Mixing Lemma

Expanders can be studied using **linear algebra**
(*Spectral Graph Theory*)

Expander Mixing Lemma

Expanders can be studied using **linear algebra**
(*Spectral Graph Theory*)

Lemma. For any subset S of nodes of a Δ -regular graph with 2nd-largest eigenvalue of adjacency matrix λ :

$$e(S, S) \leq |S| \left(\frac{|S|}{2} \frac{\Delta}{n} + \frac{\lambda}{2} \right)$$

Expander Mixing Lemma

Expanders can be studied using **linear algebra**
(*Spectral Graph Theory*)

Lemma. For any subset S of nodes of a Δ -regular graph with 2nd-largest eigenvalue of adjacency matrix λ :

$$e(S, S) \leq |S| \left(\frac{|S|}{2} \frac{\Delta}{n} + \frac{\lambda}{2} \right)$$

Proof. A adjacency matrix, 1_S indicator vector of S , J all-1 matrix. We observe

$$2e(S, S) = 1_S^T A 1_S \text{ and } 1_S^T \left(\frac{\Delta}{n} J \right) 1_S = \frac{\Delta}{n} |S|$$

Expander Mixing Lemma

Expanders can be studied using **linear algebra**
(*Spectral Graph Theory*)

Lemma. For any subset S of nodes of a Δ -regular graph with 2nd-largest eigenvalue of adjacency matrix λ :

$$e(S, S) \leq |S| \left(\frac{|S|}{2} \frac{\Delta}{n} + \frac{\lambda}{2} \right)$$

Proof. A adjacency matrix, 1_S indicator vector of S , J all-1 matrix. We observe

$$2e(S, S) = 1_S^T A 1_S \text{ and } 1_S^T \left(\frac{\Delta}{n} J \right) 1_S = \frac{\Delta}{n} |S|$$

Hence

$$2e(S, S) - \frac{\Delta}{n} |S| = 1_S^T \left(A - \frac{\Delta}{n} J \right) 1_S \leq \lambda \|1_S\|^2 = \lambda |S|$$



λ is the largest eigenvalue

Motivations for this Work I

Distributed construction of constant-degree expanders

- Corollary of Marcus-Spielman-Srivastava proof's of the Kadison-Singer conjecture [Ann. of Math. '15]:
Every dense expander has a *constant-degree subgraph* which is also an expander.

Motivations for this Work I

Distributed construction of constant-degree expanders

- Corollary of Marcus-Spielman-Srivastava proof's of the Kadison-Singer conjecture [Ann. of Math. '15]:
Every dense expander has a *constant-degree subgraph* which is also an expander.

But the proof is non-constructive.

Motivations for this Work I

Distributed construction of constant-degree expanders

- Corollary of Marcus-Spielman-Srivastava proof's of the Kadison-Singer conjecture [Ann. of Math. '15]:
Every dense expander has a *constant-degree subgraph* which is also an expander.

But the proof is non-constructive.

- Several works propose complicated distributed construction of expanders:
 - Law and Siu [INFOCOM'03]: incremental construction using Hamiltonian cycles

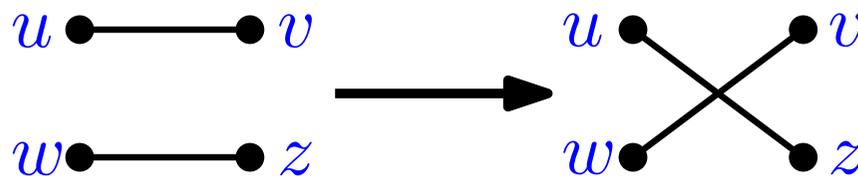
Motivations for this Work I

Distributed construction of constant-degree expanders

- Corollary of Marcus-Spielman-Srivastava proof's of the Kadison-Singer conjecture [Ann. of Math. '15]:
Every dense expander has a *constant-degree subgraph* which is also an expander.

But the proof is non-constructive.

- Several works propose complicated distributed construction of expanders:
 - Law and Siu [INFOCOM'03]: incremental construction using Hamiltonian cycles
 - Allen-Zhu et al. [SODA'16]: start with a $\Omega(\log n)$ -regular graph and increase its expansion

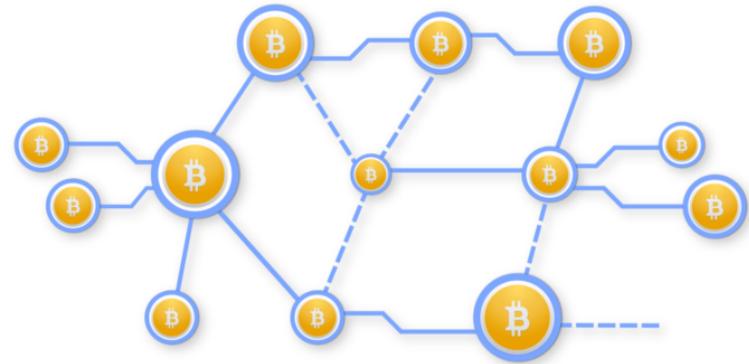


Bonus Motivation II

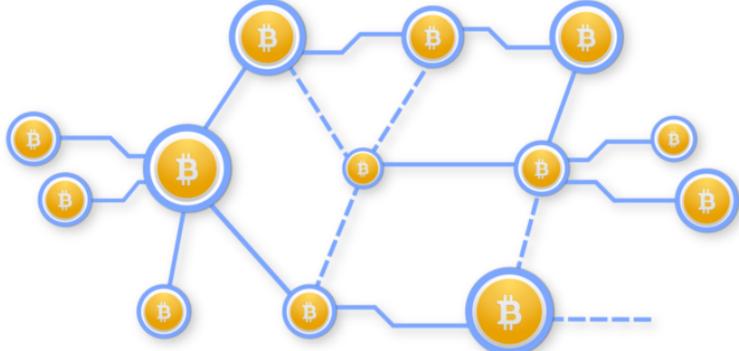
- Parallel algorithms for *sparsifying* a graph don't achieve sublogarithmic degree and assume weighted edges

Bonus Motivation II

- Parallel algorithms for *sparsifying* a graph don't achieve sublogarithmic degree and assume weighted edges
- Model the way nodes create bounded-degree overlay networks in real distributed protocols, such as in peer-to-peer protocols (BitTorrent) or in distributed ledger protocols (Bitcoin)



Bonus Motivation II

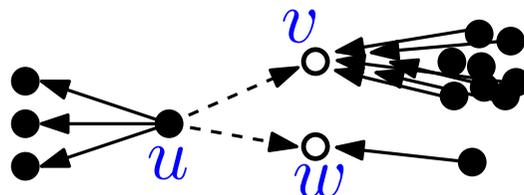
- Parallel algorithms for *sparsifying* a graph don't achieve sublogarithmic degree and assume weighted edges
- Model the way nodes create bounded-degree overlay networks in real distributed protocols, such as in peer-to-peer protocols (BitTorrent) or in distributed ledger protocols (Bitcoin)A network diagram consisting of 12 nodes, each represented by a yellow circle with a Bitcoin symbol (₿) inside. The nodes are interconnected by blue lines representing edges. Some edges are solid, while others are dashed. The network is roughly circular in shape, with nodes arranged in a ring-like pattern. The connections are sparse, with each node having a degree of 2 or 3. The overall structure is a bounded-degree graph.
- A distributed construction of a constant-degree graph implies a *constant-load balancing* algorithm. Previous works obtain almost-tight load balancing in polynomial time (Berenbrink et al., SPAA'14)

Our Algorithm: RAES

Algorithm 1 RAES(G, d, c)

```
1:  $H :=$  empty directed graph over the node set  $V$ 
2: while  $H$  has nodes of outdegree  $< d$  do
3:   PHASE 1: ▷  $d_v^{\text{out}}$ : current outdegree of  $v$  in  $H$ 
4:   for each node  $v \in V$  do
5:      $v \in V$  picks  $d - d_v^{\text{out}}$  neighbors in  $G$  uniformly at random
6:      $v$  submits a connection request to each of them
7:   end for
8:   PHASE 2: ▷  $d_v^{\text{in}}$ : current indegree of  $v$  in  $H$ 
9:   for each node  $v \in V$  do
10:    if  $v$  received  $\leq cd - d_v^{\text{in}}$  connection requests in the previous phase then
11:       $v$  accepts all of them and the corresponding directed edges are added to  $H$ 
12:    else
13:       $v$  rejects all connection requests received in Phase 1
14:    end if
15:  end for
16: end while
17: Replace each directed edge by an undirected one
18: return  $H$ 
```

Example
with $d = 5$



u is missing 2 connections.
 u asks to connect to w and v .
 v has already cd incoming connections
and refuses u 's requests.

Our Result

Theorem.

For every $d \gg 1$, $0 < \alpha \leq 1$, $c \gg \frac{1}{\alpha^2}$, and αn -regular graph G , w.h.p. $\text{RAES}(G, d, c)$ runs in $\mathcal{O}(\log n)$ parallel rounds with message complexity is $\mathcal{O}(n)$.

Moreover, if G 's 2nd-largest eigenvalue λ of normalized adjacency matrix is $\leq \epsilon \alpha^2$, then w.h.p. $\text{RAES}(G, d, c)$ creates a ϵ -expander with degrees between d and $d(c + 1)$.

Our Result

Theorem.

For every $d \gg 1$, $0 < \alpha \leq 1$, $c \gg \frac{1}{\alpha^2}$, and αn -regular graph G , w.h.p. $\text{RAES}(G, d, c)$ runs in $\mathcal{O}(\log n)$ parallel rounds with message complexity is $\mathcal{O}(n)$.

Moreover, if G 's 2nd-largest eigenvalue λ of normalized adjacency matrix is $\leq \epsilon \alpha^2$, then w.h.p. $\text{RAES}(G, d, c)$ creates a ϵ -expander with degrees between d and $d(c + 1)$.

Proof Technique: *Encoding Argument*

(omitted: message complexity using martingale theory)

Encoding Arguments

Encoding Lemma.

If X finite set and
 $C : X \rightarrow \{0, 1\}^*$ a (partial &
prefix-free) encoding of X then

$$\Pr_{x \sim \text{Unif}(X)} (|C(x)| \leq \log |X| - s) \leq 2^{-s}$$



Encoding Arguments

Encoding Lemma.

If X finite set and
 $C : X \rightarrow \{0, 1\}^*$ a (partial &
prefix-free) encoding of X then

$$\Pr_{x \sim \text{Unif}(X)} (|C(x)| \leq \log |X| - s) \leq 2^{-s}$$

Proof. $\frac{2^{\log |X| - s}}{|X|} \leq 2^{-s}.$



Encoding Arguments

Encoding Lemma.

If X finite set and
 $C : X \rightarrow \{0, 1\}^*$ a (partial &
prefix-free) encoding of X then

$$\Pr_{x \sim \text{Unif}(X)} (|C(x)| \leq \log |X| - s) \leq 2^{-s}$$

Proof. $\frac{2^{\log |X| - s}}{|X|} \leq 2^{-s}.$

Suggested reading: P. Morin et al. *Encoding Arguments*, ACM Comp. Surveys '17.



Encoding Argument Example

Flip a coin n times: 0110010...

Probability of $\log n + s$ consecutive heads?

Encoding Argument Example

Flip a coin n times: $0110010\dots$.

Probability of $\log n + s$ consecutive heads?

Call B a *bad substring* of $\log n + s$ consecutive heads. Consider encoding C_B for strings containing B :

(index i of first bit of B , all other bits of the string except those at entry $i, i + 1, \dots, i + \log n + s$)
 $\log n$ bits $n - (\log n + s)$ bits

Encoding Argument Example

Flip a coin n times: $0110010\dots$.
Probability of $\log n + s$ consecutive heads?



Call B a *bad substring* of $\log n + s$ consecutive heads.
Consider encoding C_B for strings containing B :

(index i of first bit of B , all other bits of the string except those at entry $i, i + 1, \dots, i + \log n + s$)
 $\log n$ bits $n - (\log n + s)$ bits

By the Encoding Lemma

$$\Pr(|C_B(x)| \leq \log |X| - s) = \Pr(|C_B(x)| \leq n - s) \leq 2^{-s}$$

Encoding for Always-Rejected v

We encode with the following bits

- v 's identity: $\log n$

Encoding for Always-Rejected v

We encode with the following bits

- v 's identity: $\log n$
- v 's request l_v : $2 \log l_v$

Encoding for Always-Rejected v

We encode with the following bits

- v 's identity: $\log n$
- v 's request ℓ_v : $2 \log \ell_v$
- v 's accepted requests: $2 \log d'$

Encoding for Always-Rejected v

We encode with the following bits

- v 's identity: $\log n$
- v 's request ℓ_v : $2 \log \ell_v$
- v 's accepted requests: $2 \log d'$
- position of v 's accepted requests in ℓ_v : $\log \binom{\ell_v}{d'}$

Encoding for Always-Rejected v

We encode with the following bits

- v 's identity: $\log n$
- v 's request ℓ_v : $2 \log \ell_v$
- v 's accepted requests: $2 \log d'$
- position of v 's accepted requests in ℓ_v : $\log \binom{\ell_v}{d'}$
- destinations of accepted requests: $d' \log \Delta$

Encoding for Always-Rejected v

We encode with the following bits

- v 's identity: $\log n$
- v 's request ℓ_v : $2 \log \ell_v$
- v 's accepted requests: $2 \log d'$
- position of v 's accepted requests in ℓ_v : $\log \binom{\ell_v}{d'}$
- destinations of accepted requests: $d' \log \Delta$
- destinations of rejected requests: $(\ell_v - d') \log \frac{n}{c}$

Encoding for Always-Rejected v

We encode with the following bits

- v 's identity: $\log n$
- v 's request ℓ_v : $2 \log \ell_v$
- v 's accepted requests: $2 \log d'$
- position of v 's accepted requests in ℓ_v : $\log \binom{\ell_v}{d'}$
- destinations of accepted requests: $d' \log \Delta$
- destinations of rejected requests: $(\ell_v - d') \log \frac{n}{c}$

Observation: at each round there are at most $\frac{n}{c}$ rejecting nodes



Encoding for Always-Rejected v

We encode with the following bits

- v 's identity: $\log n$
- v 's request ℓ_v : $2 \log \ell_v$
- v 's accepted requests: $2 \log d'$
- position of v 's accepted requests in ℓ_v : $\log \binom{\ell_v}{d'}$
- destinations of accepted requests: $d' \log \Delta$
- destinations of rejected requests: $(\ell_v - d') \log \frac{n}{c}$

Observation: at each round there are at most $\frac{n}{c}$ rejecting nodes



After calculations we see that we save

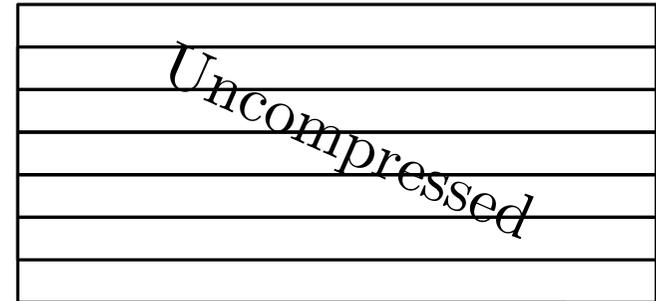
$$\frac{1}{2} \ell_v \log(\alpha c) - \log n = \Omega(\log n)$$

Compressing the Non-Expanding Set

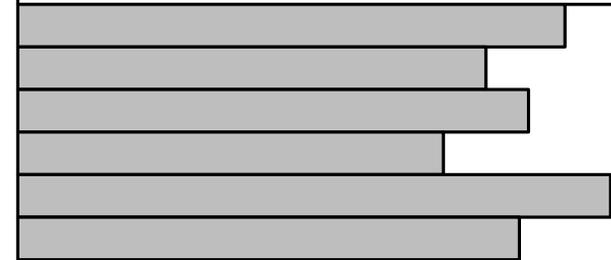
Encoding:

- Randomness of $V - S$
- Set S : $\log |S| + \log \binom{n}{s}$

Nodes in
 $V - S$



Nodes
in S

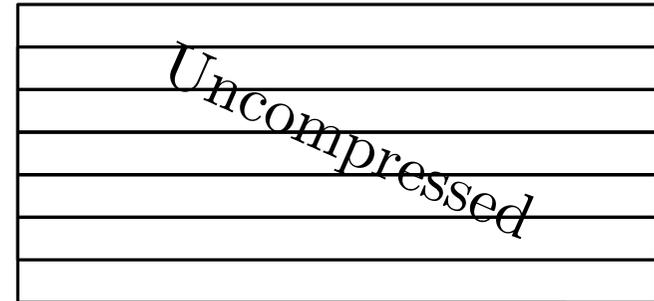


Compressing the Non-Expanding Set

Encoding:

- Randomness of $V - S$
- Set S : $\log |S| + \log \binom{n}{s}$
- Accepted connections:
 $\sum_{v \in S} 2 \log l_v + \log \binom{l_v}{d}$

Nodes in
 $V - S$



Nodes
in S

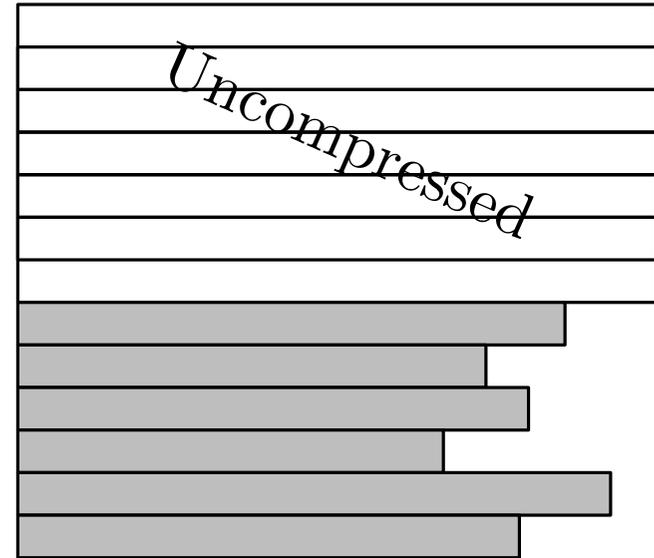


Compressing the Non-Expanding Set

Encoding:

- Randomness of $V - S$
- Set S : $\log |S| + \log \binom{n}{s}$
- Accepted connections:
 $\sum_{v \in S} 2 \log \ell_v + \log \binom{\ell_v}{d}$
- Accepted connections from S to
 $V - S$: $\sum_{v \in S} 2 \log(\epsilon_v d) + \log \binom{d}{\epsilon_v d}$
 ϵ_v : fraction of v 's accepted connections towards $V - S$

Nodes in
 $V - S$



Nodes
in S

Compressing the Non-Expanding Set

Encoding:

- Randomness of $V - S$
- Set S : $\log |S| + \log \binom{n}{s}$

- Accepted connections:
 $\sum_{v \in S} 2 \log \ell_v + \log \binom{\ell_v}{d}$

- Accepted connections from S to $V - S$: $\sum_{v \in S} 2 \log(\epsilon_v d) + \log \binom{d}{\epsilon_v d}$

ϵ_v : fraction of v 's accepted connections towards $V - S$

- Destinations of connections from S :

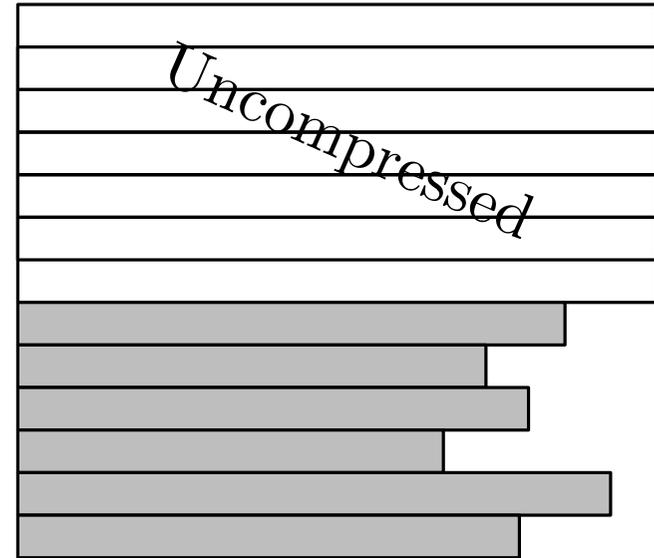
$$\sum_{v \in S} (1 - \epsilon_v) d \log((1 - \delta_v) \Delta) + \sum_{v \in S} \epsilon_v d \log \Delta$$

connections to S connections to $V - S$ (uncompressed)

δ_v : fraction of v 's edges towards $V - S$ in G

Nodes in $V - S$

Nodes in S



Compressing the Non-Expanding Set

Encoding:

- Randomness of $V - S$
- Set S : $\log |S| + \log \binom{n}{s}$

- Accepted connections:
 $\sum_{v \in S} 2 \log \ell_v + \log \binom{\ell_v}{d}$

- Accepted connections from S to $V - S$: $\sum_{v \in S} 2 \log(\epsilon_v d) + \log \binom{d}{\epsilon_v d}$

ϵ_v : fraction of v 's accepted connections towards $V - S$

- Destinations of connections from S :

$$\sum_{v \in S} (1 - \epsilon_v) d \log((1 - \delta_v) \Delta) + \sum_{v \in S} \epsilon_v d \log \Delta$$

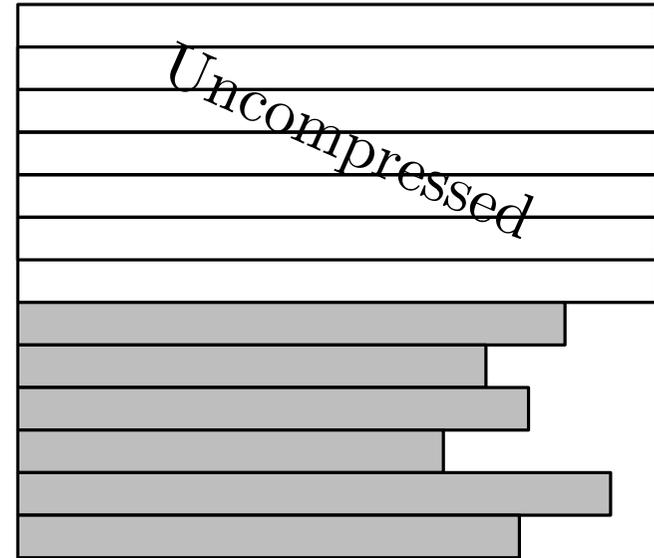
connections to S connections to $V - S$ (uncompressed)

- **Rejected requests**

δ_v : fraction of v 's edges towards $V - S$ in G

Nodes in $V - S$

Nodes in S



Compressing the Non-Expanding Set

Encoding:

- Randomness of $V - S$
- Set S : $\log |S| + \log \binom{n}{s}$

- Accepted connections:
 $\sum_{v \in S} 2 \log \ell_v + \log \binom{\ell_v}{d}$

- Accepted connections from S to $V - S$: $\sum_{v \in S} 2 \log(\epsilon_v d) + \log \binom{d}{\epsilon_v d}$

ϵ_v : fraction of v 's accepted connections towards $V - S$

- Destinations of connections from S :

$$\sum_{v \in S} (1 - \epsilon_v) d \log((1 - \delta_v) \Delta) + \sum_{v \in S} \epsilon_v d \log \Delta$$

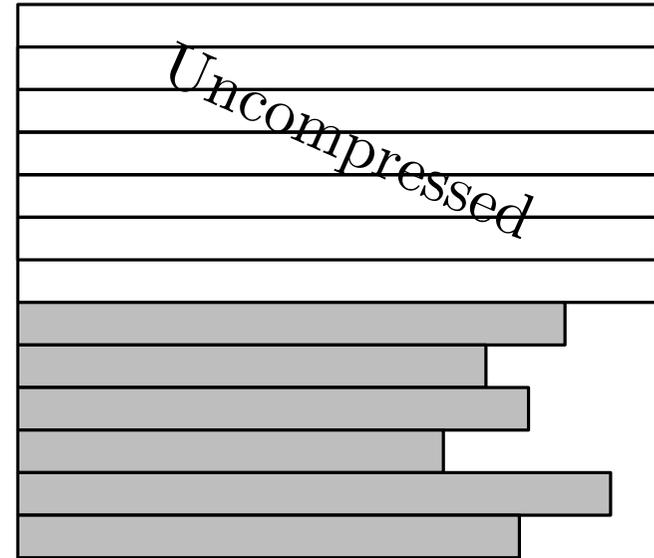
connections to S connections to $V - S$ (uncompressed)

- **Rejected requests**

- Unused randomness
(after node's termination)

Nodes in
 $V - S$

Nodes
in S



δ_v : fraction of v 's edges
towards $V - S$ in G

Compressing the Non-Expanding Set

Encoding:

- Randomness of $V - S$
- Set S : $\log |S| + \log \binom{n}{s}$

- Accepted connections:
 $\sum_{v \in S} 2 \log \ell_v + \log \binom{\ell_v}{d}$

- Accepted connections from S to $V - S$: $\sum_{v \in S} 2 \log(\epsilon_v d) + \log \binom{d}{\epsilon_v d}$

ϵ_v : fraction of v 's accepted connections towards $V - S$

- Destinations of connections from S :

$$\sum_{v \in S} (1 - \epsilon_v) d \log((1 - \delta_v) \Delta) + \sum_{v \in S} \epsilon_v d \log \Delta$$

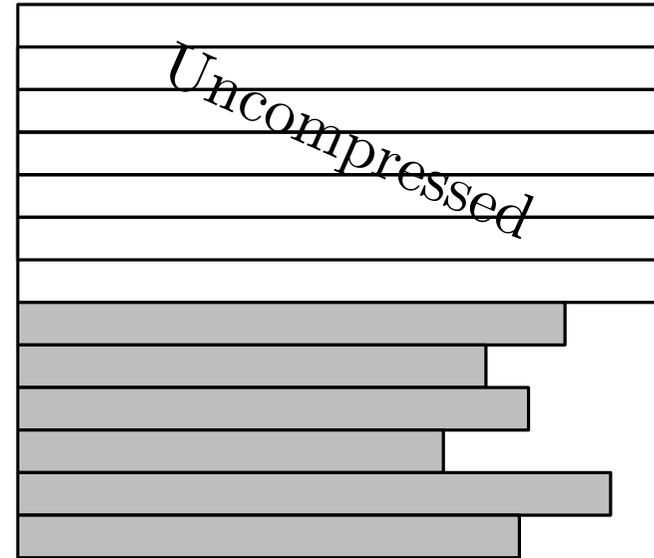
connections to S connections to $V - S$ (uncompressed)

- **Rejected requests**

- Unused randomness
(after node's termination)

Nodes in
 $V - S$

Nodes
in S



δ_v : fraction of v 's edges
towards $V - S$ in G

Compressing Accepted Connections I

To represent accepted requests from S we need

$$\sum_{v \in S} (1 - \epsilon_v) d \log((1 - \delta_v) \Delta) + \sum_{v \in S} \epsilon_v d \log \Delta$$
$$\leq sd \log \Delta - \frac{1 - \epsilon}{2} sd \log \frac{n}{s} + 2\epsilon ds$$

where $\epsilon = \frac{1}{s} \sum_{v \in S} \epsilon_v$

Compressing Accepted Connections I

To represent accepted requests from S we need

$$\begin{aligned} & \sum_{v \in S} (1 - \epsilon_v) d \log((1 - \delta_v) \Delta) + \sum_{v \in S} \epsilon_v d \log \Delta \\ & \leq sd \log \Delta - \frac{1 - \epsilon}{2} sd \log \frac{n}{s} + 2\epsilon ds \end{aligned}$$

where $\epsilon = \frac{1}{s} \sum_{v \in S} \epsilon_v$

With simple calculations

$$\begin{aligned} sd \log \Delta - (\sum_{v \in S} (1 - \epsilon_v) d \log((1 - \delta_v) \Delta) + \sum_{v \in S} \epsilon_v d \log \Delta) \\ \geq d \sum_{v \in S} (1 - \epsilon_v) \log \frac{1}{1 - \delta_v} \end{aligned}$$

Compressing Accepted Connections I

To represent accepted requests from S we need

$$\begin{aligned} & \sum_{v \in S} (1 - \epsilon_v) d \log((1 - \delta_v) \Delta) + \sum_{v \in S} \epsilon_v d \log \Delta \\ & \leq sd \log \Delta - \frac{1 - \epsilon}{2} sd \log \frac{n}{s} + 2\epsilon ds \end{aligned}$$

where $\epsilon = \frac{1}{s} \sum_{v \in S} \epsilon_v$

With simple calculations

$$\begin{aligned} & sd \log \Delta - (\sum_{v \in S} (1 - \epsilon_v) d \log((1 - \delta_v) \Delta) + \sum_{v \in S} \epsilon_v d \log \Delta) \\ & \geq d \sum_{v \in S} (1 - \epsilon_v) \log \frac{1}{1 - \delta_v} \end{aligned}$$

Two cases: $s < \alpha \Delta$ and $\alpha \Delta \leq s \leq \frac{n}{2} \dots$

Compressing Accepted Connections II

Goal: bound $d \sum_{v \in S} (1 - \epsilon_v) \log \frac{1}{1 - \delta_v}$

Case $s < \alpha \Delta$

Use $\Delta(1 - \delta_v) \leq s$ and $(\frac{\Delta}{s})^2 > \frac{\Delta}{s} \frac{1}{\alpha} = \frac{\Delta}{s} \frac{n}{\Delta} = \frac{n}{s}$

hence $d \sum_{v \in S} (1 - \epsilon_v) \log \frac{1}{1 - \delta_v} > \frac{1 - \epsilon}{2} s d \log \frac{n}{s}$

Compressing Accepted Connections II

Goal: bound $d \sum_{v \in S} (1 - \epsilon_v) \log \frac{1}{1 - \delta_v}$

Case $s < \alpha \Delta$

Use $\Delta(1 - \delta_v) \leq s$ and $(\frac{\Delta}{s})^2 > \frac{\Delta}{s} \frac{1}{\alpha} = \frac{\Delta}{s} \frac{n}{\Delta} = \frac{n}{s}$

hence $d \sum_{v \in S} (1 - \epsilon_v) \log \frac{1}{1 - \delta_v} > \frac{1 - \epsilon}{2} s d \log \frac{n}{s}$

Case $\alpha \Delta \leq s \leq \frac{n}{2}$

Rewrite $-(1 - \epsilon) s d \sum_{v \in S} \frac{1 - \epsilon_v}{(1 - \epsilon)s} \log \frac{1}{1 - \delta_v}$

use Jensen's inequality to get $(1 - \epsilon) s d \log \frac{1 - \epsilon}{1 - \delta}$

Compressing Accepted Connections II

Goal: bound $d \sum_{v \in S} (1 - \epsilon_v) \log \frac{1}{1 - \delta_v}$

Case $s < \alpha \Delta$

Use $\Delta(1 - \delta_v) \leq s$ and $(\frac{\Delta}{s})^2 > \frac{\Delta}{s} \frac{1}{\alpha} = \frac{\Delta}{s} \frac{n}{\Delta} = \frac{n}{s}$

hence $d \sum_{v \in S} (1 - \epsilon_v) \log \frac{1}{1 - \delta_v} > \frac{1 - \epsilon}{2} s d \log \frac{n}{s}$

Case $\alpha \Delta \leq s \leq \frac{n}{2}$

Rewrite $-(1 - \epsilon) s d \sum_{v \in S} \frac{1 - \epsilon_v}{(1 - \epsilon)^s} \log \frac{1}{1 - \delta_v}$

use Jensen's inequality to get $(1 - \epsilon) s d \log \frac{1 - \epsilon}{1 - \delta}$

To bound $1 - \delta$ we use the **Expander Mixing Lemma**:

$$(1 - \delta) \leq \frac{s}{n} + \lambda$$

Compressing Accepted Connections II

Goal: bound $d \sum_{v \in S} (1 - \epsilon_v) \log \frac{1}{1 - \delta_v}$

Case $s < \alpha \Delta$

Use $\Delta(1 - \delta_v) \leq s$ and $(\frac{\Delta}{s})^2 > \frac{\Delta}{s} \frac{1}{\alpha} = \frac{\Delta}{s} \frac{n}{\Delta} = \frac{n}{s}$

hence $d \sum_{v \in S} (1 - \epsilon_v) \log \frac{1}{1 - \delta_v} > \frac{1 - \epsilon}{2} s d \log \frac{n}{s}$

Case $\alpha \Delta \leq s \leq \frac{n}{2}$

Rewrite $-(1 - \epsilon) s d \sum_{v \in S} \frac{1 - \epsilon_v}{(1 - \epsilon)^s} \log \frac{1}{1 - \delta_v}$

use Jensen's inequality to get $(1 - \epsilon) s d \log \frac{1 - \epsilon}{1 - \delta}$

To bound $1 - \delta$ we use the **Expander Mixing Lemma**:

$$(1 - \delta) \leq \frac{s}{n} + \lambda$$

together with hypothesis on s and λ , it implies

$$(1 - \epsilon) s d \log \frac{1 - \epsilon}{1 - \delta} > (1 - \epsilon) s d \log \frac{n}{s} - 2\epsilon d s$$

Compressing the Non-Expanding Set

Encoding:

- Randomness of $V - S$
- Set S : $\log |S| + \log \binom{n}{s}$

- Accepted connections:
 $\sum_{v \in S} 2 \log \ell_v + \log \binom{\ell_v}{d}$

- Accepted connections from S to $V - S$: $\sum_{v \in S} 2 \log(\epsilon_v d) + \log \binom{d}{\epsilon_v d}$

ϵ_v : fraction of v 's accepted connections towards $V - S$

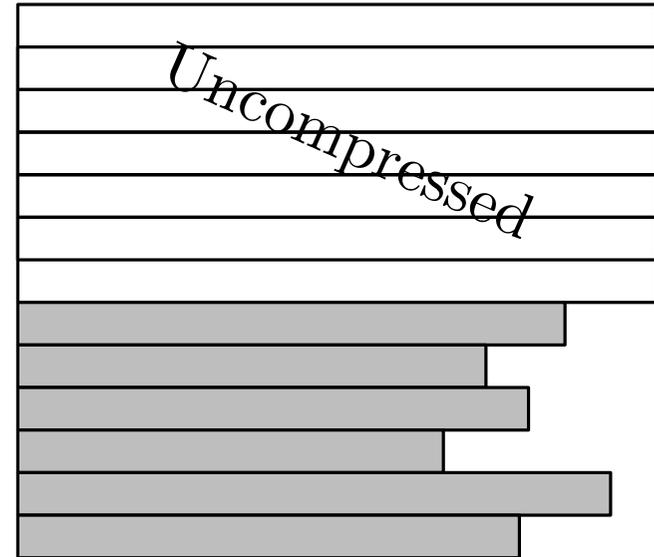
- Destinations of connections from S :

$$\sum_{v \in S} (1 - \epsilon_v) d \log((1 - \delta_v) \Delta) + \sum_{v \in S} \epsilon_v d \log \Delta$$

connections to S connections to $V - S$ (uncompressed)

- **Rejected requests**
- Unused randomness
(after node's termination)

Nodes in
 $V - S$



Nodes
in S

Compressing Rejected Requests (Idea)

With $\ell_v - d'$ bits we encode which requests are rejected.

The hard part is compressing their *destinations*, for which we use the following notions:

Semi-saturated nodes ss_t : accepted connections until time $t - 1$ + requests from $V - S$ are $> \frac{dc}{2}$

Critical nodes c_t : not semi-saturated at time t but accepted + rejected connections are $> cd$

Compressing Rejected Requests (Idea)

With $\ell_v - d'$ bits we encode which requests are rejected.

The hard part is compressing their *destinations*, for which we use the following notions:

Semi-saturated nodes ss_t : accepted connections until time $t - 1$ + requests from $V - S$ are $> \frac{dc}{2}$

Critical nodes c_t : not semi-saturated at time t but accepted + rejected connections are $> cd$

Claim. semi-saturated nodes $\leq \frac{n}{2c}$ and critical nodes $\leq \frac{n}{c}$.

Compressing Rejected Requests (Idea)

With $\ell_v - d'$ bits we encode which requests are rejected.

The hard part is compressing their *destinations*, for which we use the following notions:

Semi-saturated nodes ss_t : accepted connections until time $t - 1$ + requests from $V - S$ are $> \frac{dc}{2}$

Critical nodes c_t : not semi-saturated at time t but accepted + rejected connections are $> cd$

Claim. semi-saturated nodes $\leq \frac{n}{2n}$ and critical nodes $\leq \frac{n}{c}$.

We can then write

$$ss(v) \log \frac{2n}{c} + \sum_1^T rc_t(v) \log c_t$$

Where $rss(v)$ is the number of rejected connections from v to semisaturated nodes and $rc_t(v)$ is the number of rejected connections from v to critical nodes at time t

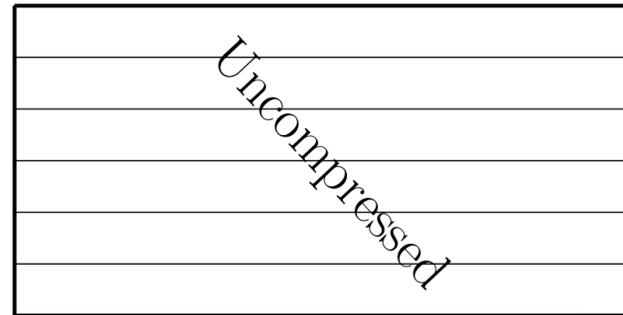
Compression Summary

Set S

Size	Index of the set
------	------------------

$$2 \log |S| + \log \binom{n}{|S|}$$

Nodes in $V \setminus S$

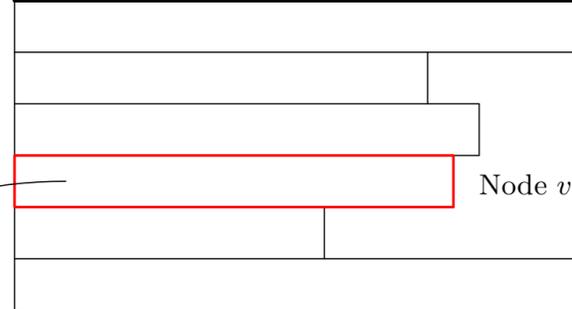


Critical Nodes

Sizes	Indices of sets
-------	-----------------

$$\sum_{t=1}^T \left[\log c_t + \log \binom{n}{c_t} \right]$$

Nodes in S



9 bronze badges

Subset of accepted requests	Subset of accepted requests in S	Destinations of accepted requests outside S (uncompressed) + + inside S (compressed)	Destinations of rejected requests
-----------------------------	------------------------------------	---	--

$$2 \log l_v + \log \binom{l_v}{d}$$

$$2 \log(\varepsilon_v d) + \log \binom{d}{\varepsilon_v d}$$

$$\varepsilon_v d \log \Delta + (1 - \varepsilon_v) d \log((1 - \delta)\Delta)$$

Semi-saturated / Critical	S.-sat. dest.	Crit. dest.	Crit. dest.	S.-sat. dest.	S.-sat. dest.	Crit. dest.
---------------------------	---------------	-------------	-------------	---------------	---------------	-------------

$$l_v - d$$

$$\log(n/c)$$

$$\log c_{t_1}$$

$$\log c_{t_2}$$

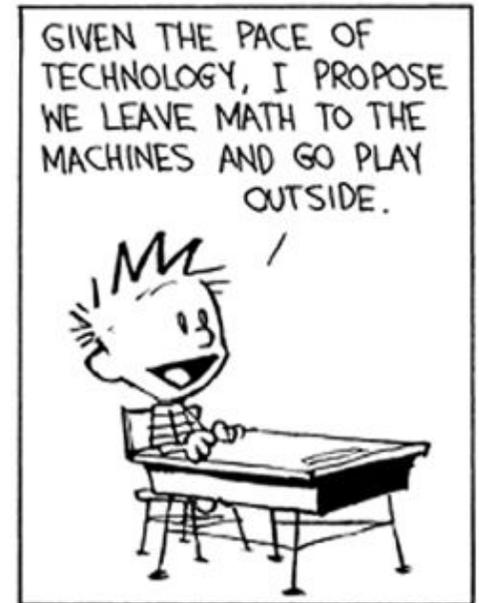
$$\log(n/c)$$

$$\log(n/c)$$

$$\log c_{t_k}$$

Open Problems

- Generalizing to non-dense expanders.
E.g., not clear if all nodes can achieve d connections if $\Delta = o(n)$ (if $\Delta = O(\log n)$, this happens w.h.p.).
- Extending analysis to non-regular graphs.
- Investigate robustness of RAES when nodes join or leave the network.



Thank You!