Demand Aware Network (DAN) Design

Some Results and Open Questions

Chen Avin

Ben-Gurion University of the Negev

Joint work with Stefan Schmid, Kaushik Mondal, Alexandr Hercules, Andreas Loukas

Motivation

- Demand Aware Network Design?
  - "self-adjust" the networks' routing paths (topology) to routing requests
- Data Centres?
  - ProjecTor / Wireless technologies
  - Skype example?
- Peer-to-Peer Networks

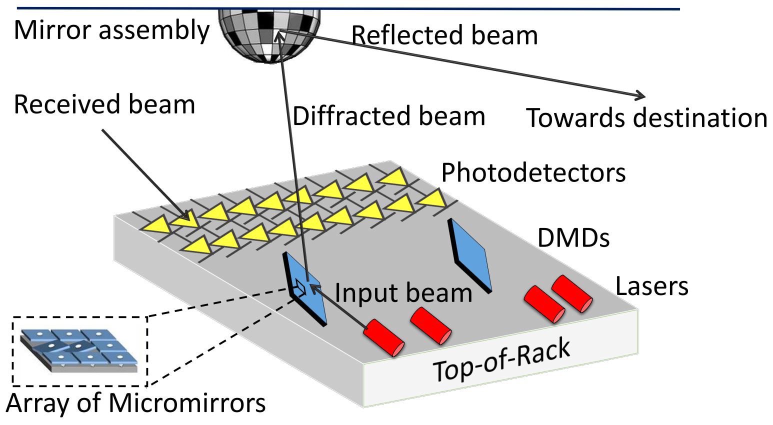

Outline

- Motivation
- Problem Settings
    - Relation to other problems
- Lower Bounds
- Bounded degree network design
- The continuous discrete approach
- Future work

Problem Settings

- Demand distribution, $\mathcal{D}$ over $V \times V$.
  - Pairwise communication demands
  - Can be represented as directed weighted graph

	1	2	3	4	5	6	7
1	0	$\frac{2}{65}$	$\frac{1}{13}$	$\frac{1}{65}$	$\frac{1}{65}$	$\frac{2}{65}$	$\frac{3}{65}$
2	$\frac{2}{65}$	0	$\frac{1}{65}$	0	0	0	$\frac{2}{65}$
3	$\frac{1}{13}$	$\frac{1}{65}$	0	$\frac{2}{65}$	0	0	$\frac{1}{13}$
4	$\frac{1}{65}$	0	$\frac{2}{65}$	0	$\frac{4}{65}$	0	0
5	$\frac{1}{65}$	0	$\frac{3}{65}$	$\frac{4}{65}$	0	0	0
6	$\frac{2}{65}$	0	0	0	0	0	$\frac{3}{65}$
7	$\frac{3}{65}$	$\frac{2}{65}$	$\frac{1}{13}$	0	0	$\frac{3}{65}$	0

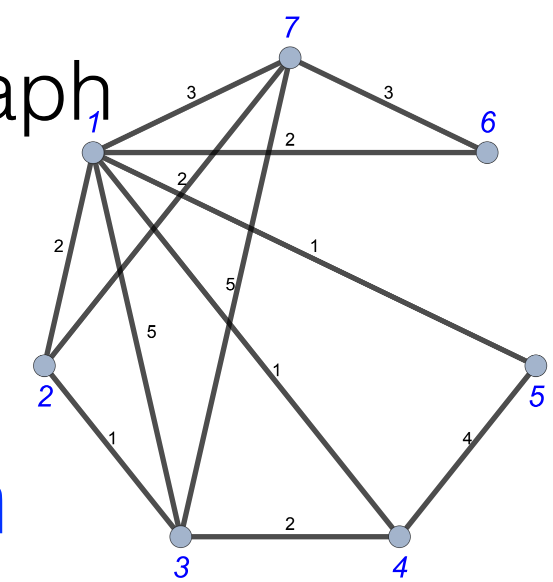

- A network $N = (V, E)$
- Metric of interest: **Expected Path Length**

$$\mathrm{EPL}(\mathcal{D}, N) = \mathbb{E}_{\mathcal{D}}[\mathrm{d}_N(\cdot,\cdot)] = \sum_{(u,v)\in\mathcal{D}} \mathrm{p}(u,v) \cdot \mathrm{d}_N(u,v)$$

$\mathrm{d}_N(u,v)$ - hop distance between *u,v* in *N*

Problem Settings

- Demand distribution, $\mathcal{D}$
- Expected path length

$$\mathrm{EPL}(\mathcal{D}, N) = \mathbb{E}_{\mathcal{D}}[\mathrm{d}_N(\cdot,\cdot)] = \sum_{(u,v)\in\mathcal{D}} \mathrm{p}(u,v)\cdot \mathrm{d}_N(u,v)$$

- Desired topology family $\mathcal{N}$
  - e.g., bounded degree, trees, sparse, etc.
- Optimal Demand Aware Network (DAN)

$$N^* = \arg\min_{N\in\mathcal{N}} \mathrm{EPL}(\mathcal{D}, N)$$

	1	2	3	4	5	6	7
1	0	$\frac{2}{65}$	$\frac{1}{13}$	$\frac{1}{65}$	$\frac{1}{65}$	$\frac{2}{65}$	$\frac{3}{65}$
2	$\frac{2}{65}$	0	$\frac{1}{65}$	0	0	0	$\frac{2}{65}$
3	$\frac{1}{13}$	$\frac{1}{65}$	0	$\frac{2}{65}$	0	0	$\frac{1}{13}$
4	$\frac{1}{65}$	0	$\frac{2}{65}$	0	$\frac{4}{65}$	0	0
5	$\frac{1}{65}$	0	$\frac{3}{65}$	$\frac{4}{65}$	0	0	0
6	$\frac{2}{65}$	0	0	0	0	0	$\frac{3}{65}$
7	$\frac{3}{65}$	$\frac{2}{65}$	$\frac{1}{13}$	0	0	$\frac{3}{65}$	0

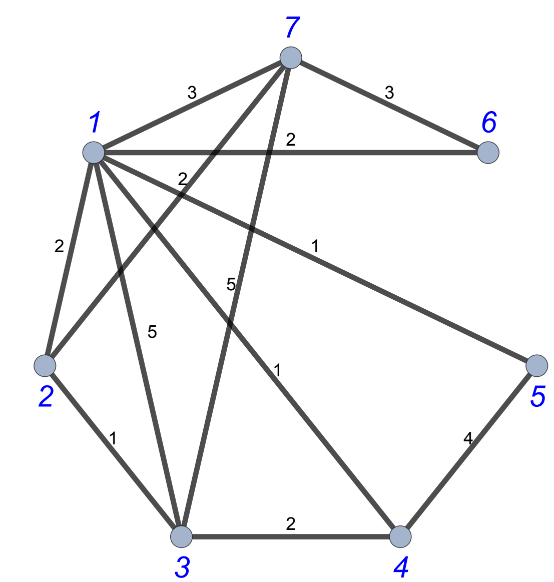

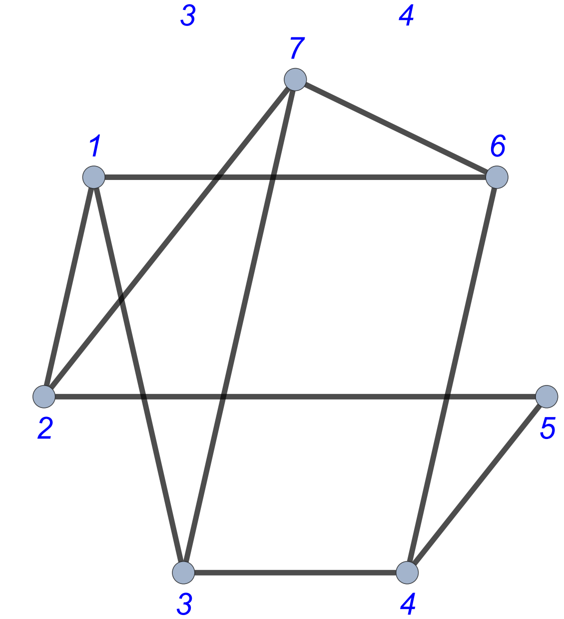

Relation to Other Problems

- Minimum Linear Arrangement (MLA)

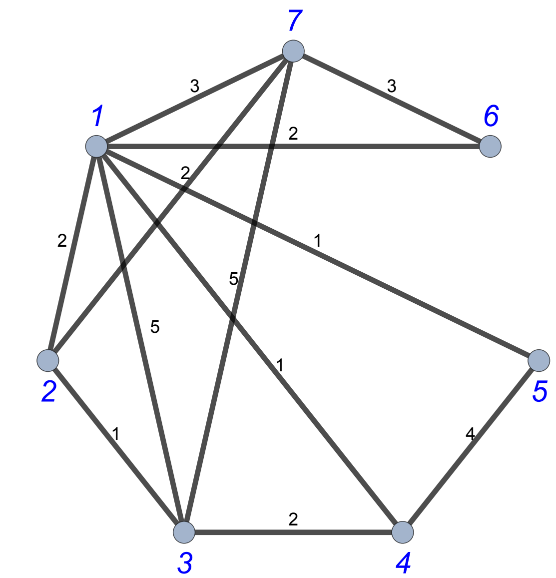

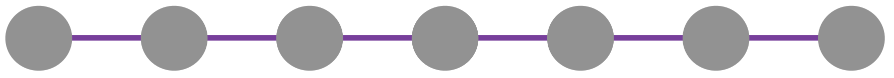

Relation to Other Problems

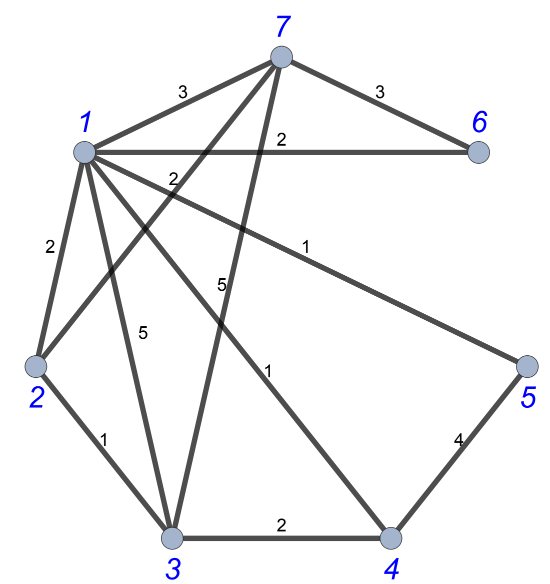

- Minimum Linear Arrangement (MLA)
- Embeddings (guest, host graphs)
- Spanners
- Information Theory / Coding
  - Entropy: $H(X) = \sum_{i=1}^{n} p(x_i) \log_2 \frac{1}{p(x_i)}$
  - Conditional Entropy: $H(X|Y) = \sum_{j=1}^{n} p(y_j) H(X|Y = y_j)$
  - Coding - Expected code length

Lower Bound

- For a Δ bounded degree **DAN**
- Theorem

$$N^* \geq \Omega(\max(H_\Delta(Y|X), H_\Delta(X|Y))$$

- Proof Idea (using coding):
  - Replacing each row with an optimal Δ-ary tree
  - Same for columns
  - Optimal code length is larger than row Entropy

	1	2	3	4	5	6	7
1	0	$\frac{2}{65}$	$\frac{1}{13}$	$\frac{1}{65}$	$\frac{1}{65}$	$\frac{2}{65}$	$\frac{3}{65}$
2	$\frac{2}{65}$	0	$\frac{1}{65}$	0	0	0	$\frac{2}{65}$
3	$\frac{1}{13}$	$\frac{1}{65}$	0	$\frac{2}{65}$	0	0	$\frac{1}{13}$
4	$\frac{1}{65}$	0	$\frac{2}{65}$	0	$\frac{4}{65}$	0	0
5	$\frac{1}{65}$	0	$\frac{3}{65}$	$\frac{4}{65}$	0	0	0
6	$\frac{2}{65}$	0	0	0	0	0	$\frac{3}{65}$
7	$\frac{3}{65}$	$\frac{2}{65}$	$\frac{1}{13}$	0	0	$\frac{3}{65}$	0

Bounded Degree `DAN`

- Bounded (e.g., Δ = constant) degree
- Theorem: Can design "optimal" network N, s.t

  $$\mathrm{EPL}(\mathcal{D}, N) \leq O(H(Y \mid X) + H(X \mid Y))$$

  for,
- Sparse distributions (weighted, directed)
- Local doubling dimension distribution
  - Possibly dense but uniform and regular

Sparse Distributions

- Proof idea

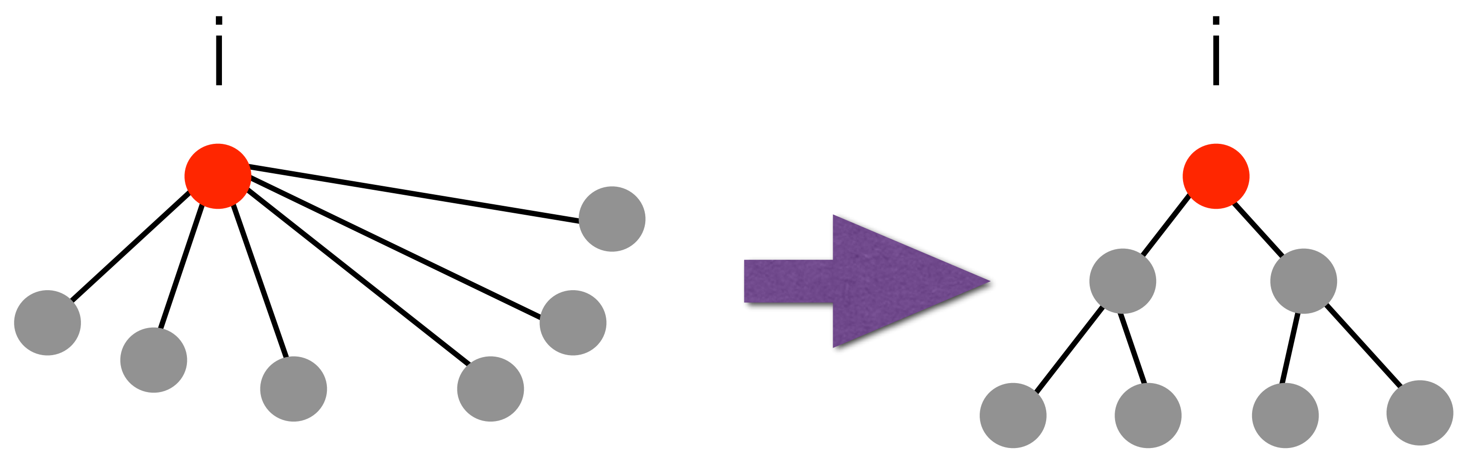

Optimal bounded degree tree

Sparse Distributions

- Proof idea

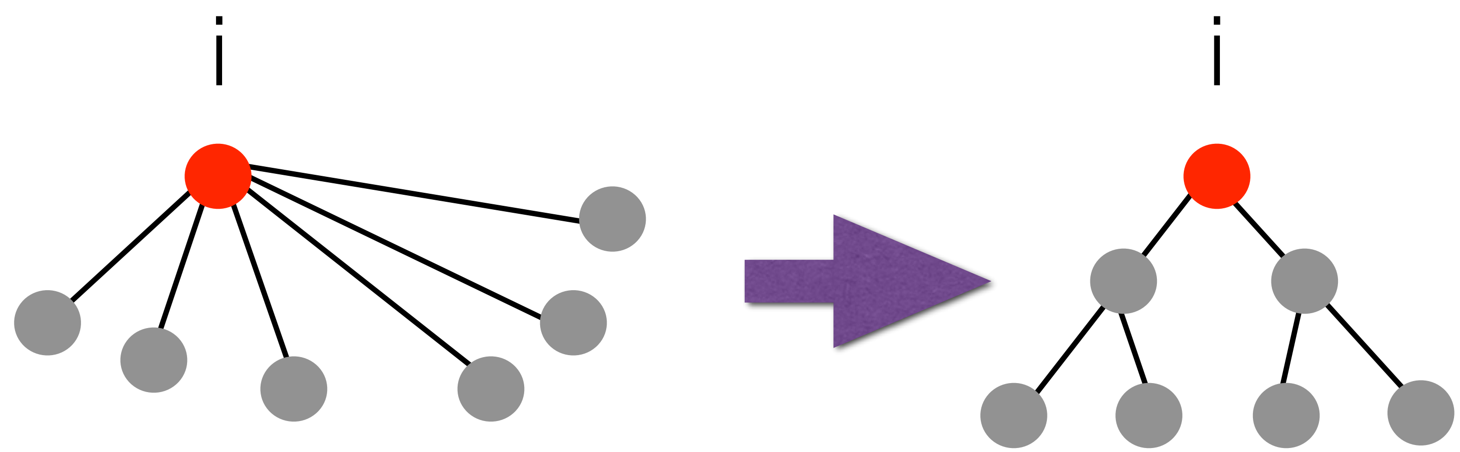

Optimal bounded degree tree

Problem

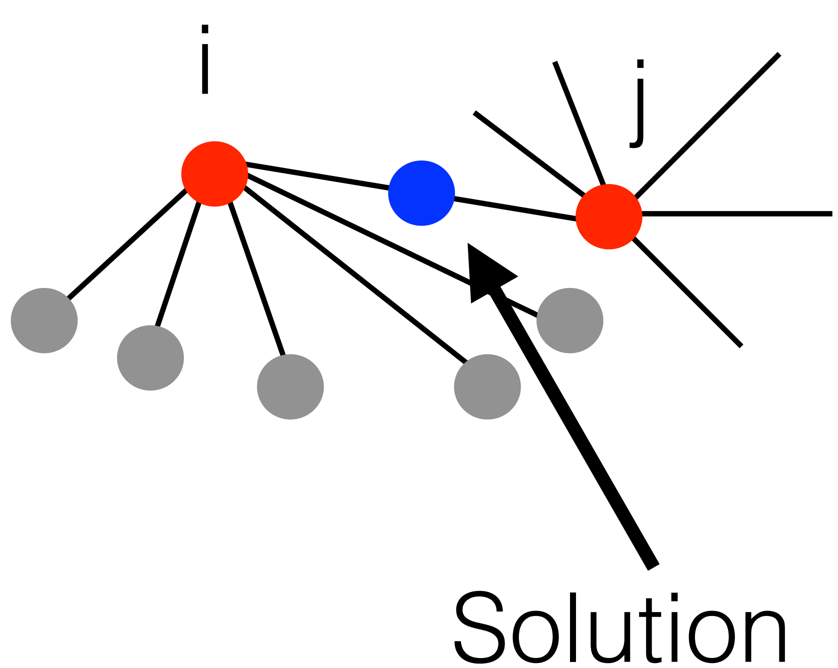

Sparse Distributions

- Proof idea

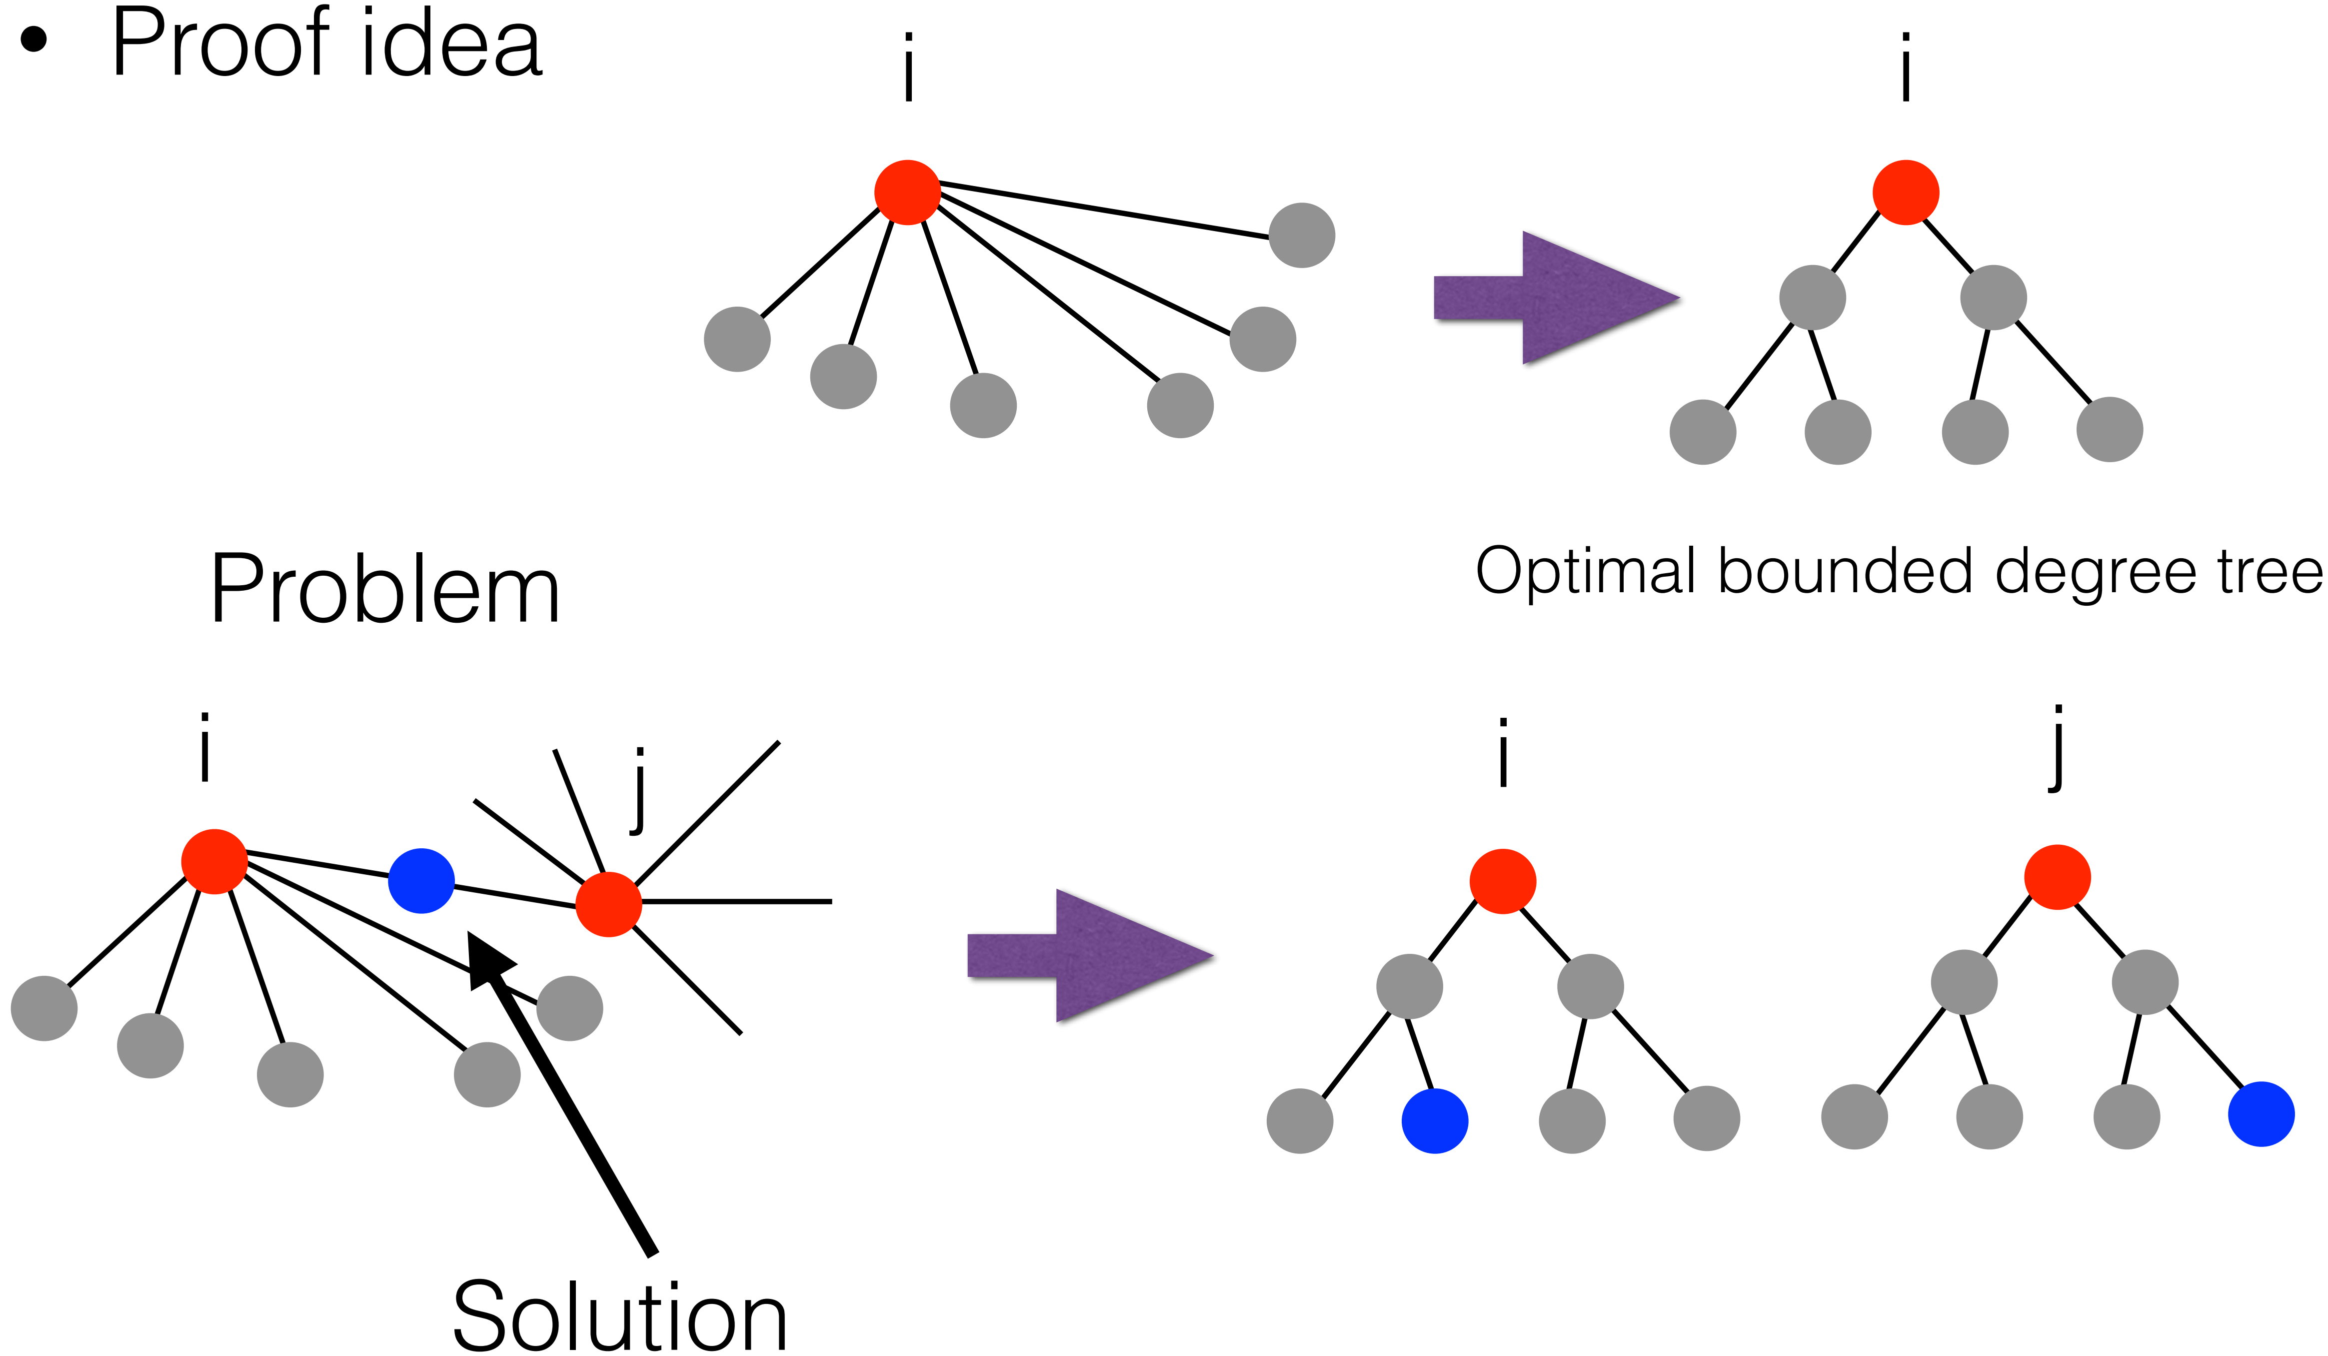

Doubling Dimensions Dist.

- **Local** Doubling Dimension distribution

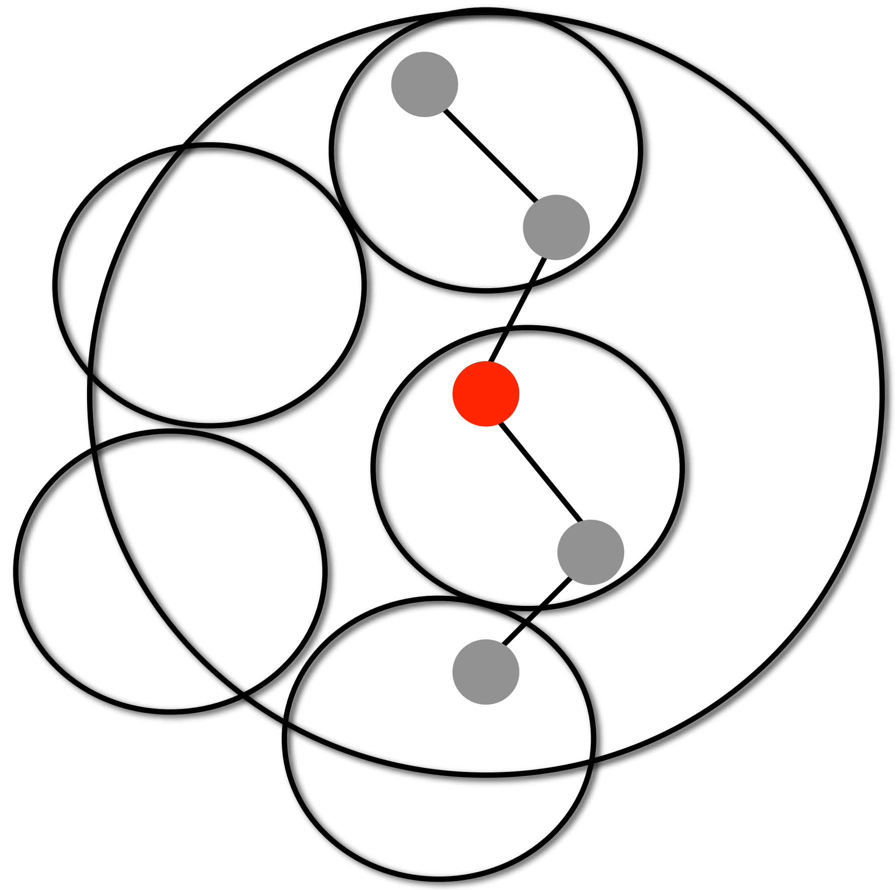

2-hops balls
can be covered by
1-hop balls

Doubling Dimensions Dist.

- **Local** Doubling Dimension distribution

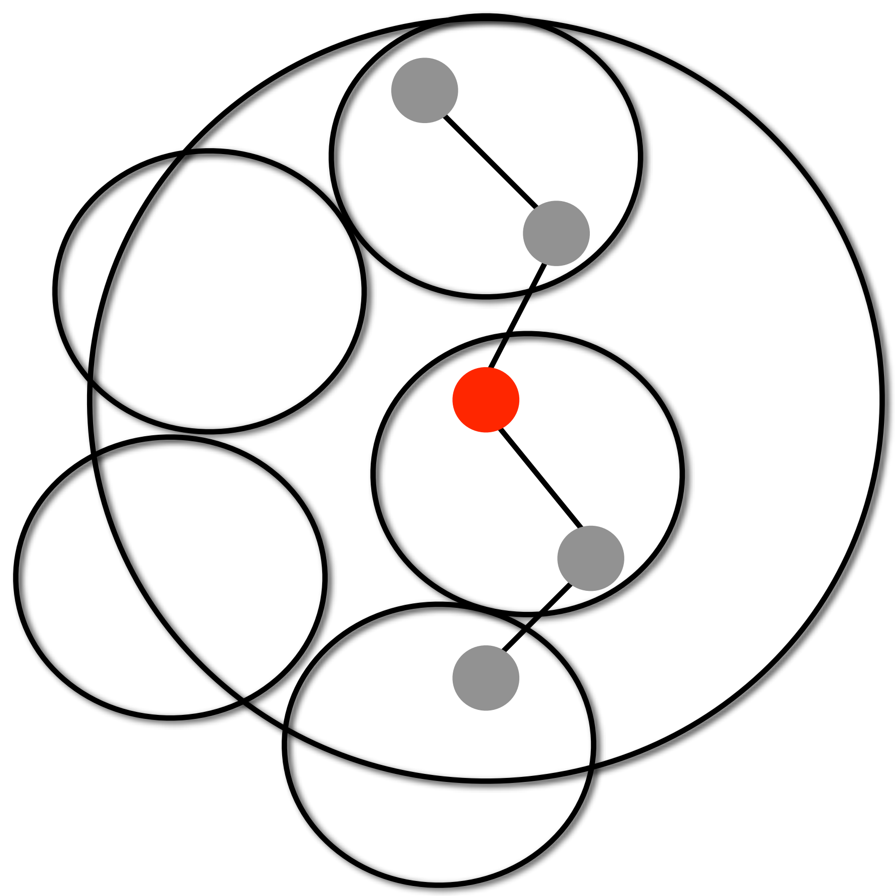

2-hops balls
can be covered by
1-hop balls

- Can be a dense graph

Continuous-Discrete Design

- Greedy routing

	1	2	3	4	5	6	7
1	0	$\frac{2}{65}$	$\frac{1}{13}$	$\frac{1}{65}$	$\frac{1}{65}$	$\frac{2}{65}$	$\frac{3}{65}$
2	$\frac{2}{65}$	0	$\frac{1}{65}$	0	0	0	$\frac{2}{65}$
3	$\frac{1}{13}$	$\frac{1}{65}$	0	$\frac{2}{65}$	0	0	$\frac{1}{13}$
4	$\frac{1}{65}$	0	$\frac{2}{65}$	0	$\frac{4}{65}$	0	0
5	$\frac{1}{65}$	0	$\frac{3}{65}$	$\frac{4}{65}$	0	0	0
6	$\frac{2}{65}$	0	0	0	0	0	$\frac{3}{65}$
7	$\frac{3}{65}$	$\frac{2}{65}$	$\frac{1}{13}$	0	0	$\frac{3}{65}$	0

Continuous-Discrete Design

- Greedy routing

	1	2	3	4	5	6	7
1	0	$\frac{2}{65}$	$\frac{1}{13}$	$\frac{1}{65}$	$\frac{1}{65}$	$\frac{2}{65}$	$\frac{3}{65}$
2	$\frac{2}{65}$	0	$\frac{1}{65}$	0	0	0	$\frac{2}{65}$
3	$\frac{1}{13}$	$\frac{1}{65}$	0	$\frac{2}{65}$	0	0	$\frac{1}{13}$
4	$\frac{1}{65}$	0	$\frac{2}{65}$	0	$\frac{4}{65}$	0	0
5	$\frac{1}{65}$	0	$\frac{3}{65}$	$\frac{4}{65}$	0	0	0
6	$\frac{2}{65}$	0	0	0	0	0	$\frac{3}{65}$
7	$\frac{3}{65}$	$\frac{2}{65}$	$\frac{1}{13}$	0	0	$\frac{3}{65}$	0

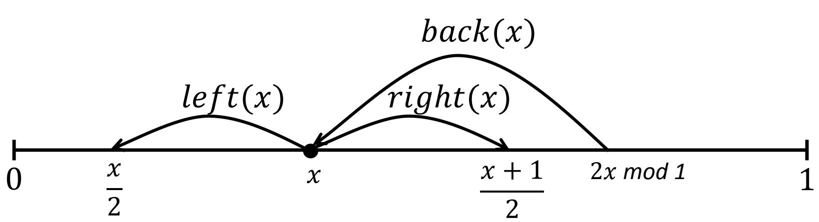

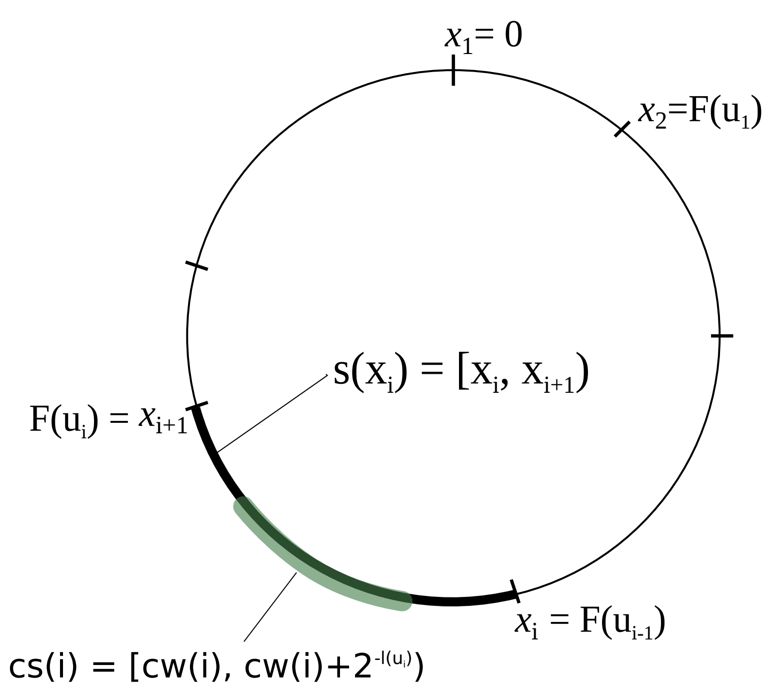

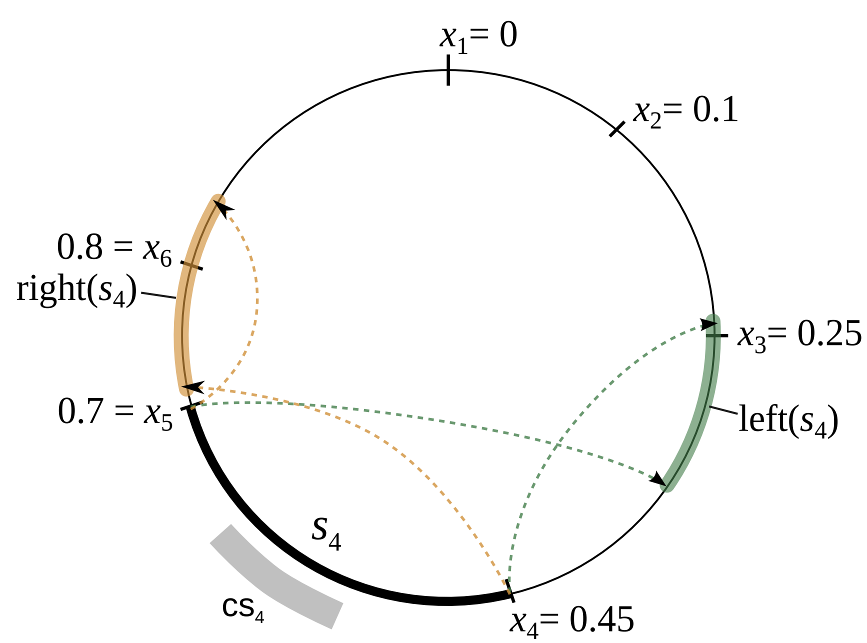

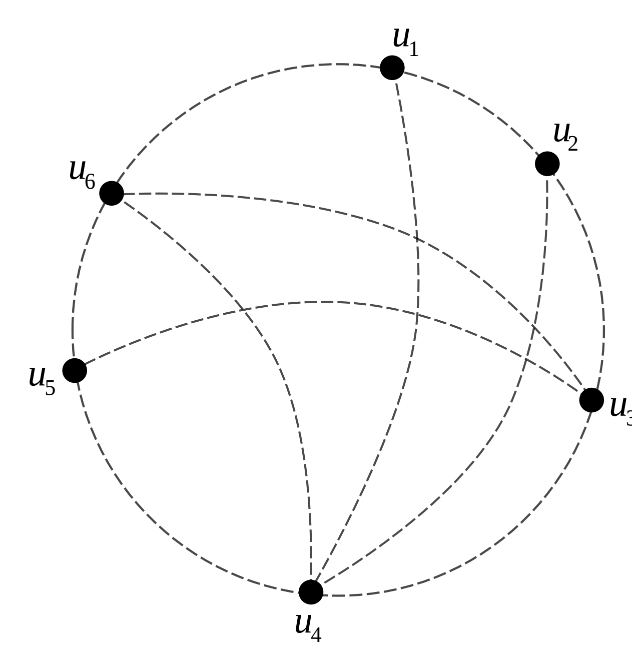

Continuous-Discrete Design

- Greedy routing

	1	2	3	4	5	6	7
1	0	$\frac{2}{65}$	$\frac{1}{13}$	$\frac{1}{65}$	$\frac{1}{65}$	$\frac{2}{65}$	$\frac{3}{65}$
2	$\frac{2}{65}$	0	$\frac{1}{65}$	0	0	0	$\frac{2}{65}$
3	$\frac{1}{13}$	$\frac{1}{65}$	0	$\frac{2}{65}$	0	0	$\frac{1}{13}$
4	$\frac{1}{65}$	0	$\frac{2}{65}$	0	$\frac{4}{65}$	0	0
5	$\frac{1}{65}$	0	$\frac{3}{65}$	$\frac{4}{65}$	0	0	0
6	$\frac{2}{65}$	0	0	0	0	0	$\frac{3}{65}$
7	$\frac{3}{65}$	$\frac{2}{65}$	$\frac{1}{13}$	0	0	$\frac{3}{65}$	0

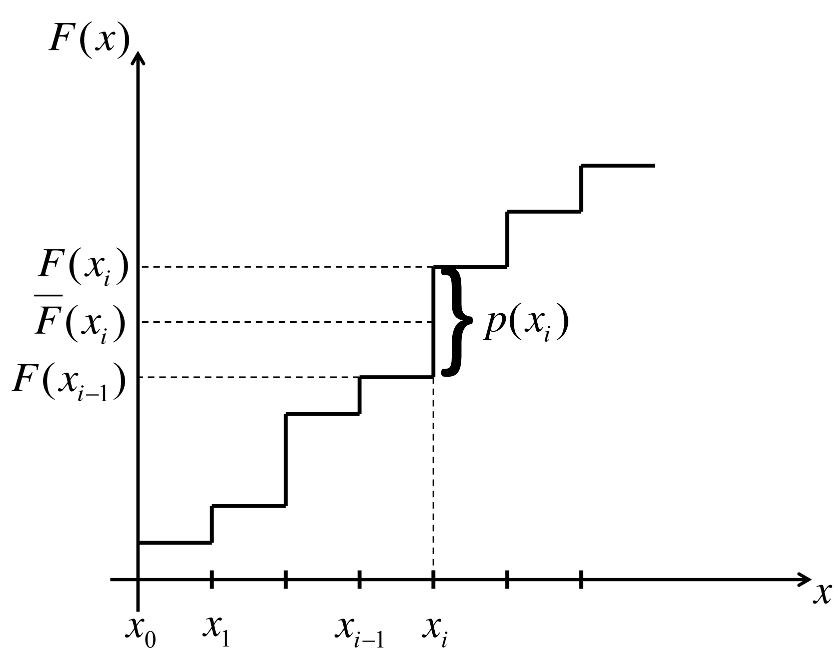

Shannon-Fano-Elias Coding

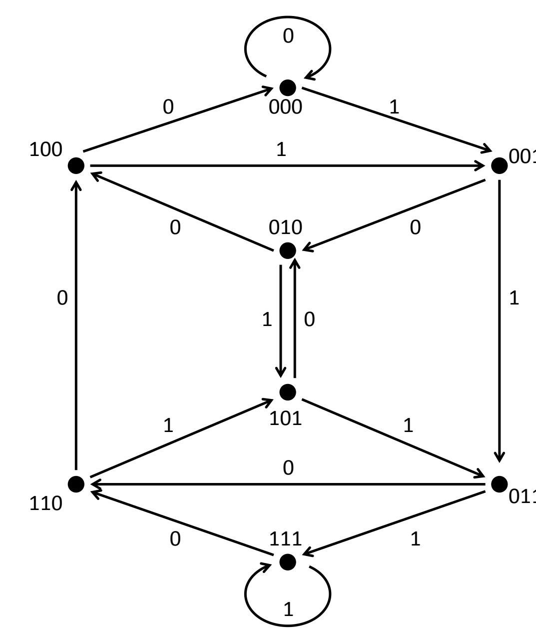

De-Bruijn Graph

Continuous-Discrete Design

- Greedy routing
- Theorem:
  - Linear size
  - Fair (please explain)
  - Robust to failures
  - Expected path length:

$$EPL(\mathcal{R}, G, \mathcal{A}) < \min\{H(\boldsymbol{p_s}), H(\boldsymbol{p_d})\} + 2.$$

	1	2	3	4	5	6	7
1	0	$\frac{2}{65}$	$\frac{1}{13}$	$\frac{1}{65}$	$\frac{1}{65}$	$\frac{2}{65}$	$\frac{3}{65}$
2	$\frac{2}{65}$	0	$\frac{1}{65}$	0	0	0	$\frac{2}{65}$
3	$\frac{1}{13}$	$\frac{1}{65}$	0	$\frac{2}{65}$	0	0	$\frac{1}{13}$
4	$\frac{1}{65}$	0	$\frac{2}{65}$	0	$\frac{4}{65}$	0	0
5	$\frac{1}{65}$	0	$\frac{3}{65}$	$\frac{4}{65}$	0	0	0
6	$\frac{2}{65}$	0	0	0	0	0	$\frac{3}{65}$
7	$\frac{3}{65}$	$\frac{2}{65}$	$\frac{1}{13}$	0	0	$\frac{3}{65}$	0

Future Work / Discussion

- New “Graph Entropy” measure for networks
- Online algorithms - Amortize analysis
  - Splay-nets example
- Distributed algorithms?
- Practical use ???

Thank you

avin@cse.bgu.ac.il

See papers:

- **Demand-Aware Network Designs of Bounded Degree**. Chen Avin, Kaushik Mondal, and Stefan Schmid.. ArXiv Technical Report, May 2017. https://arxiv.org/abs/1705.06024
- **Towards Communication-Aware Robust Topologies**. Chen Avin, Alexandr Hercules, Andreas Loukas, and Stefan Schmid. https://arxiv.org/abs/1705.07163
- **SplayNet: Towards Locally Self-Adjusting Networks**. Stefan Schmid, Chen Avin, Christian Scheideler, Michael Borokhovich, Bernhard Haeupler, and Zvi Lotker. IEEE/ACM Transactions on Networking (ToN). http://ieeexplore.ieee.org/document/7066977/