Beating Treewidth for Average-Case Subgraph Isomorphism

Gregory Rosenthal

University of Toronto

IPEC 2019

Subgraph Isomorphism Problem

- Does X have a subgraph isomorphic to G?
- Parameterize by fixing G.

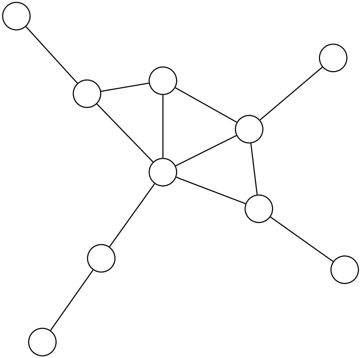

(a) X

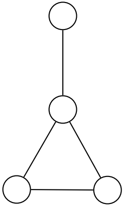

(b) G

Subgraph Isomorphism Problem

- Does X have a subgraph isomorphic to G?
- Parameterize by fixing G.

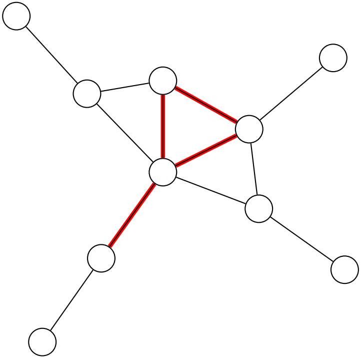

(a) X

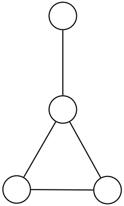

(b) G

Colored Subgraph Isomorphism Problem

Does X have a subgraph H such that the given coloring is an isomorphism from H to G?

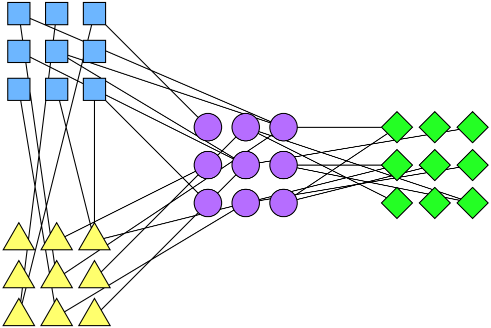

(a) X

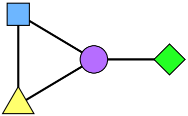

(b) G

Subgraph isomorphism efficiently reduces to colored subgraph isomorphism [Alon–Yuster–Zwick'95].

Colored Subgraph Isomorphism Problem

Does X have a subgraph H such that the given coloring is an isomorphism from H to G?

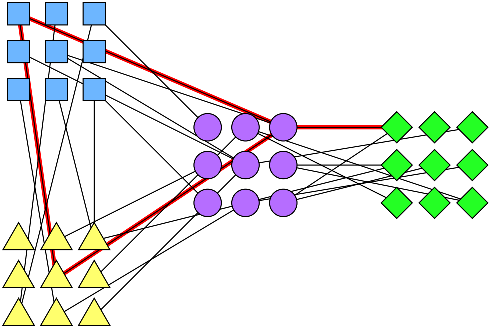

(a) X

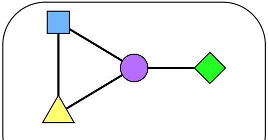

(b) G

Subgraph isomorphism efficiently reduces to colored subgraph isomorphism [Alon–Yuster–Zwick'95].

Previous Results

Time Complexity

- $O(n^{tw(G)+1})$ upper bound [AYZ'95]
- $n^{\Omega(tw(G)/\log tw(G))}$ lower bound assuming ETH [Marx'10]
  - "Substantially different techniques" required to close the gap when G is a 3-regular expander [Alon–Marx'11].

Previous Results

Time Complexity

- $O(n^{tw(G)+1})$ upper bound [AYZ'95]
- $n^{\Omega(tw(G)/\log tw(G))}$ lower bound assuming ETH [Marx'10]
  - "Substantially different techniques" required to close the gap when G is a 3-regular expander [Alon–Marx'11].

A Family of Input Distributions

- Vertices: $V(G) \times [n]$
- Let $\beta : E(G) \to \mathbb{R}_{\geq 0}$.
- Include each edge $\{(u,i),(v,j)\}$ independently with probability $n^{-\beta(\{u,v\})}$.

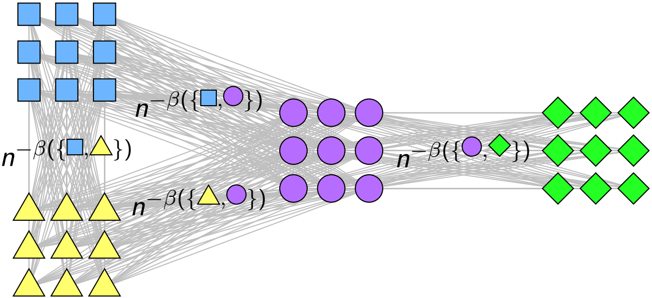

[Li–Razborov–Rossman'17]

The Average-Case Problem

- Fail with probability $o(1)$.
- "Lower Bound": for *some* edge weighting
- "Upper Bound": upper bounds for *all* edge weightings
- Input distribution is *nontrivial* if $P(\exists\ G\text{-colored subgraph})$ is bounded away from 0 and 1.

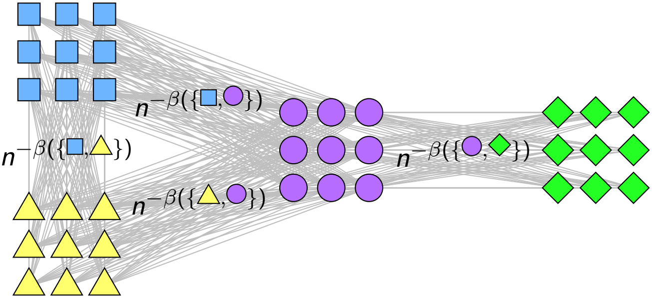

[Li–Razborov–Rossman'17]

AC^0 Circuits

- Constant-depth, unbounded fanin boolean circuits.
- The *size* of a circuit is the number of gates.

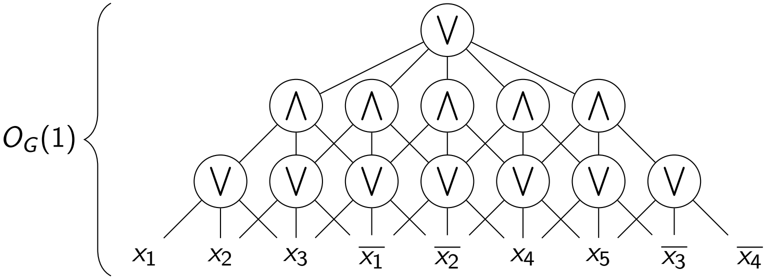

Previous Results

Time Complexity

- $O(n^{tw(G)+1})$ upper bound [AYZ'95]
- $n^{\Omega(tw(G)/\log tw(G))}$ lower bound assuming ETH [Marx'10]
  - "Substantially different techniques" required to close the gap when G is a 3-regular expander [Alon–Marx'11].

AC^0 Circuit Size

- $O(n^{tw(G)+1})$ upper bound [Amano'10]
- $n^{\kappa(G)-o(1)}$ *average-case* lower bound [LRR'17]
- $\kappa(G)$ is $\Omega(tw(G)/\log tw(G))$ [LRR'17]
  - For constant-degree expanders, $\kappa(G)$ is $\Omega(tw(G))$ [LRR'17].
- $n^{2\kappa(G)+O(1)}$ *average-case* upper bound [LRR'17]

Previous Results

Time Complexity

- $O(n^{tw(G)+1})$ upper bound [AYZ'95]
- $n^{\Omega(tw(G)/\log tw(G))}$ lower bound assuming ETH [Marx'10]
  - "Substantially different techniques" required to close the gap when G is a 3-regular expander [Alon–Marx'11].

AC^0 Circuit Size

- $O(n^{tw(G)+1})$ upper bound [Amano'10]
- $n^{\kappa(G)-o(1)}$ *average-case* lower bound [LRR'17]
- $\kappa(G)$ is $\Omega(tw(G)/\log tw(G))$ [LRR'17]
  - For constant-degree expanders, $\kappa(G)$ is $\Omega(tw(G))$ [LRR'17].
- $n^{2\kappa(G)+O(1)}$ *average-case* upper bound [LRR'17]

Previous Results

Time Complexity

- $O(n^{tw(G)+1})$ upper bound [AYZ'95]
- $n^{\Omega(tw(G)/\log tw(G))}$ lower bound assuming ETH [Marx'10]
  - "Substantially different techniques" required to close the gap when G is a 3-regular expander [Alon–Marx'11].

AC^0 Circuit Size

- $O(n^{tw(G)+1})$ upper bound [Amano'10]
- $n^{\kappa(G)-o(1)}$ *average-case* lower bound [LRR'17]
- $\kappa(G)$ is $\Omega(tw(G)/\log tw(G))$ [LRR'17]
  - For constant-degree expanders, $\kappa(G)$ is $\Omega(tw(G))$ [LRR'17].
- $n^{2\kappa(G)+O(1)}$ *average-case* upper bound [LRR'17]

Our Contributions

- The average-case $[\mathrm{AC}^0$, time$]$ complexity is at most $n^{\kappa(G)+O(1)}$.
- If G is a hypercube then $\kappa(G)$ is $\Theta\left(tw(G)/\sqrt{\log tw(G)}\right)$.
- $\kappa(G)$ is Ω(the exponent from Marx's ETH-hardness result).

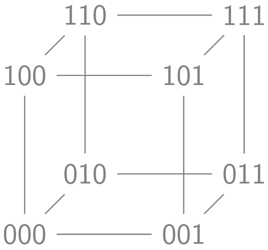

Our Contributions

- The average-case $[\mathrm{AC}^0, \text{time}]$ complexity is at most $n^{\kappa(G)+O(1)}$.
- If G is a hypercube then $\kappa(G)$ is $\Theta\left(tw(G)/\sqrt{\log tw(G)}\right)$.
- $\kappa(G)$ is Ω(the exponent from Marx's ETH-hardness result).

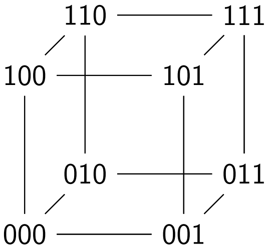

Our Contributions

- The average-case $[\mathrm{AC}^0, \text{time}]$ complexity is at most $n^{\kappa(G)+O(1)}$.
- If G is a hypercube then $\kappa(G)$ is $\Theta\left(tw(G)/\sqrt{\log tw(G)}\right)$.
- $\kappa(G)$ is Ω(the exponent from Marx's ETH-hardness result).

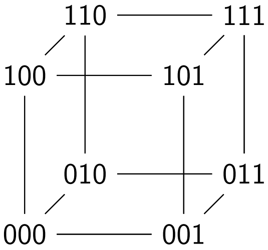

Our Contributions

- The average-case $[\mathrm{AC}^0, \text{time}]$ complexity is at most $n^{\kappa(G)+O(1)}$.
- If G is a hypercube then $\kappa(G)$ is $\Theta\left(tw(G)/\sqrt{\log tw(G)}\right)$.
- $\kappa(G)$ is Ω(the exponent from Marx's ETH-hardness result).

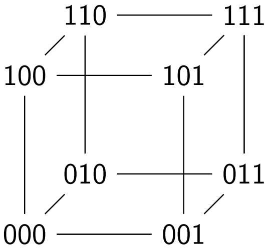

Union Sequences

Sequences $(H_1, \ldots, H_k)$ of subgraphs of G such that

(a) $H_k = G$, and

(b) each H is either an edge or the union of two previous graphs in the sequence.

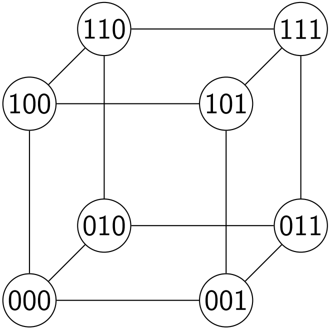

Union Sequences

Sequences $(H_1, \ldots, H_k)$ of subgraphs of G such that

(a) $H_k = G$, and

(b) each H is either an edge or the union of two previous graphs in the sequence.

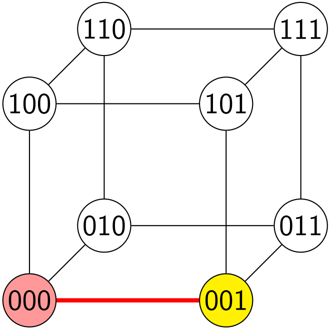

Union Sequences

Sequences $(H_1, \ldots, H_k)$ of subgraphs of G such that

(a) $H_k = G$, and

(b) each H is either an edge or the union of two previous graphs in the sequence.

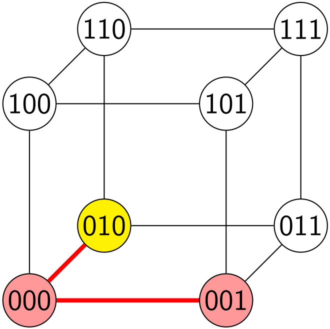

Union Sequences

Sequences $(H_1, \ldots, H_k)$ of subgraphs of G such that

(a) $H_k = G$, and

(b) each H is either an edge or the union of two previous graphs in the sequence.

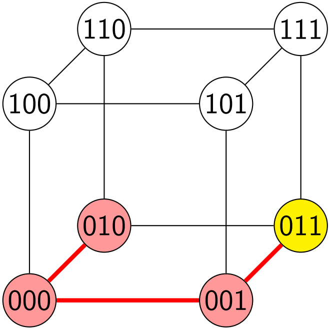

Union Sequences

Sequences $(H_1, \ldots, H_k)$ of subgraphs of G such that

(a) $H_k = G$, and

(b) each H is either an edge or the union of two previous graphs in the sequence.

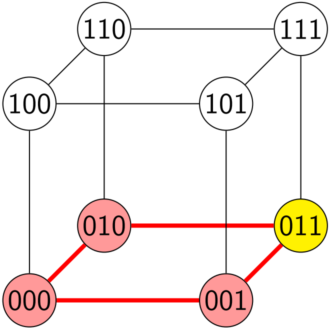

Union Sequences

Sequences $(H_1, \ldots, H_k)$ of subgraphs of G such that

(a) $H_k = G$, and

(b) each H is either an edge or the union of two previous graphs in the sequence.

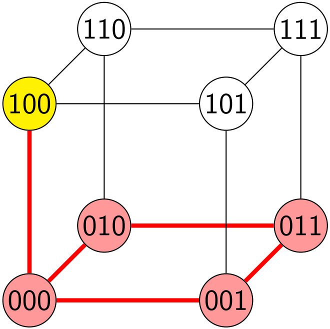

Union Sequences

Sequences $(H_1, \ldots, H_k)$ of subgraphs of G such that

(a) $H_k = G$, and

(b) each H is either an edge or the union of two previous graphs in the sequence.

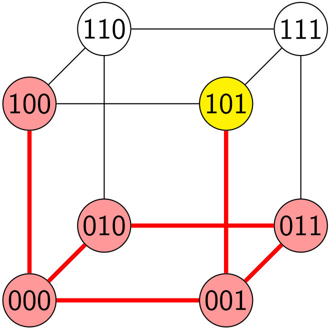

Union Sequences

Sequences $(H_1, \ldots, H_k)$ of subgraphs of G such that

(a) $H_k = G$, and

(b) each H is either an edge or the union of two previous graphs in the sequence.

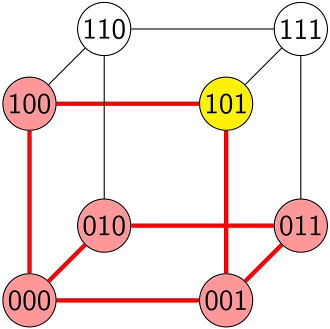

Union Sequences

Sequences $(H_1, \ldots, H_k)$ of subgraphs of G such that

(a) $H_k = G$, and

(b) each H is either an edge or the union of two previous graphs in the sequence.

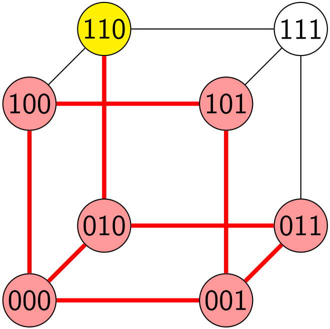

Union Sequences

Sequences $(H_1, \ldots, H_k)$ of subgraphs of G such that

(a) $H_k = G$, and

(b) each H is either an edge or the union of two previous graphs in the sequence.

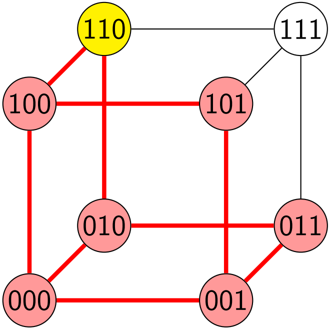

Union Sequences

Sequences $(H_1, \ldots, H_k)$ of subgraphs of G such that

(a) $H_k = G$, and

(b) each H is either an edge or the union of two previous graphs in the sequence.

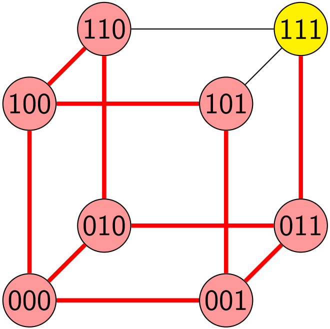

Union Sequences

Sequences $(H_1, \ldots, H_k)$ of subgraphs of G such that

(a) $H_k = G$, and

(b) each H is either an edge or the union of two previous graphs in the sequence.

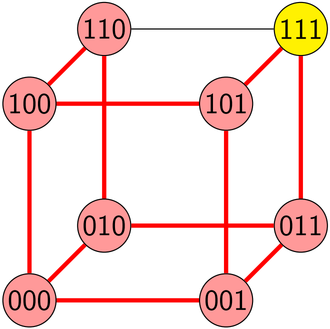

Union Sequences

Sequences $(H_1, \ldots, H_k)$ of subgraphs of G such that

(a) $H_k = G$, and

(b) each H is either an edge or the union of two previous graphs in the sequence.

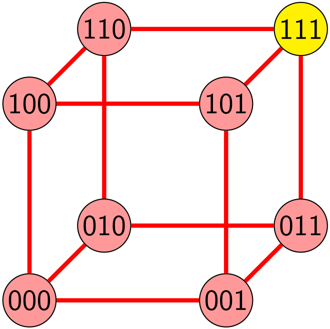

Union Sequences

Sequences $(H_1, \ldots, H_k)$ of subgraphs of G such that

(a) $H_k = G$, and

(b) each H is either an edge or the union of two previous graphs in the sequence.

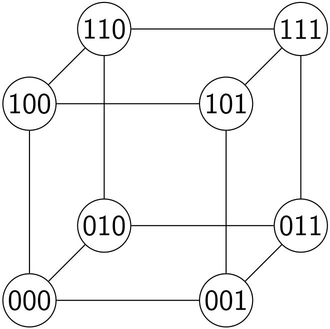

Union Sequences

Sequences $(H_1, \ldots, H_k)$ of subgraphs of G such that

(a) $H_k = G$, and

(b) each H is either an edge or the union of two previous graphs in the sequence.

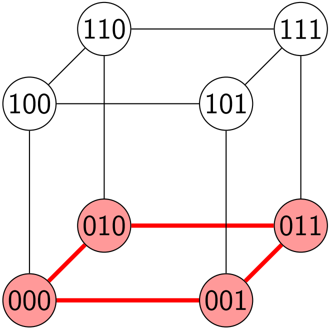

Union Sequences

Sequences $(H_1, \ldots, H_k)$ of subgraphs of G such that

(a) $H_k = G$, and

(b) each H is either an edge or the union of two previous graphs in the sequence.

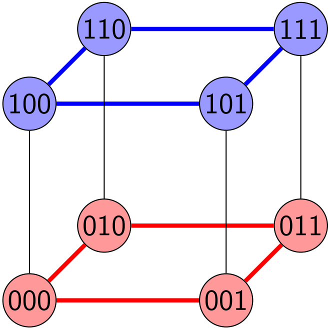

Union Sequences

Sequences $(H_1, \ldots, H_k)$ of subgraphs of G such that

(a) $H_k = G$, and

(b) each H is either an edge or the union of two previous graphs in the sequence.

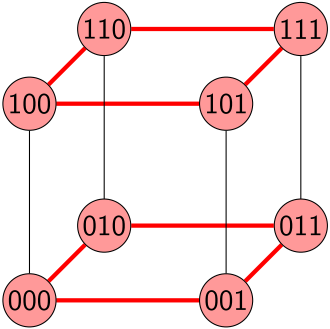

Union Sequences

Sequences $(H_1, \ldots, H_k)$ of subgraphs of G such that

(a) $H_k = G$, and

(b) each H is either an edge or the union of two previous graphs in the sequence.

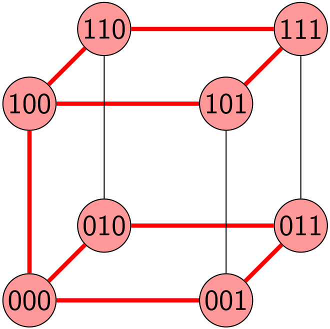

Union Sequences

Sequences $(H_1, \ldots, H_k)$ of subgraphs of G such that

(a) $H_k = G$, and

(b) each H is either an edge or the union of two previous graphs in the sequence.

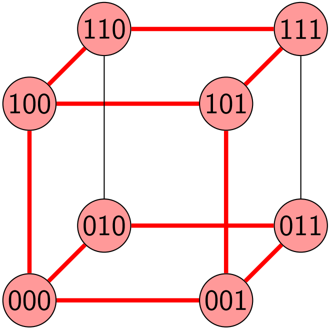

Union Sequences

Sequences $(H_1, \ldots, H_k)$ of subgraphs of G such that

(a) $H_k = G$, and

(b) each H is either an edge or the union of two previous graphs in the sequence.

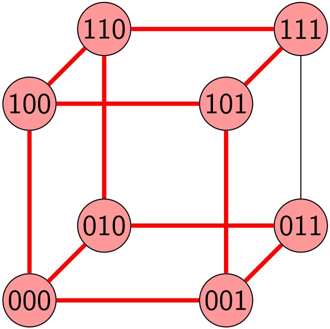

Union Sequences

Sequences $(H_1, \ldots, H_k)$ of subgraphs of G such that

(a) $H_k = G$, and

(b) each H is either an edge or the union of two previous graphs in the sequence.

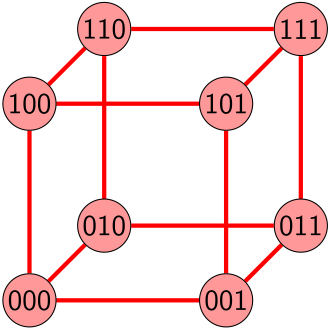

Average-Case Upper Bound

For a union sequence $(H_1, \ldots, H_k)$:

- For each successive $H = A \cup B$, find all H-colored subgraphs by considering all pairs of A-colored and B-colored subgraphs.
- Runtime is $\tilde{O}(\max_H \mathbb{E}[\#\ H\text{-colored subgraphs}]^2)$ w.h.p.

Quadratic improvement with *sort-merge-join*

- Challenge: sorting is not in AC^0 [Håstad'86].

> $n^{\kappa(G)}$ is the maximum over nontrivial input distributions, of the minimum over union sequences $(H_1, \ldots, H_k)$, of $\max_H \mathbb{E}[\#\ H\text{-colored subgraphs}]$.

Average-Case Upper Bound

For a union sequence $(H_1, \ldots, H_k)$:

- For each successive $H = A \cup B$, find all H-colored subgraphs by considering all pairs of A-colored and B-colored subgraphs.
- Runtime is $\tilde{O}(\max_H \mathbb{E}[\#\ H\text{-colored subgraphs}]^2)$ w.h.p.

Quadratic improvement with *sort-merge-join*

- Challenge: sorting is not in AC^0 [Håstad'86].

$n^{\kappa(G)}$ is the maximum over nontrivial input distributions, of the minimum over union sequences $(H_1, \ldots, H_k)$, of $\max_H \mathbb{E}[\#\ H\text{-colored subgraphs}]$.

Average-Case Upper Bound

For a union sequence $(H_1, \ldots, H_k)$:

- For each successive $H = A \cup B$, find all H-colored subgraphs by considering all pairs of A-colored and B-colored subgraphs.
- Runtime is $\tilde{O}(\max_H \mathbb{E}[\#\ H\text{-colored subgraphs}]^2)$ w.h.p.

Quadratic improvement with *sort-merge-join*

- Challenge: sorting is not in AC^0 [Håstad'86].

> $n^{\kappa(G)}$ is the maximum over nontrivial input distributions, of the minimum over union sequences $(H_1, \ldots, H_k)$, of $\max_H \mathbb{E}[\#\ H\text{-colored subgraphs}]$.

Average-Case Upper Bound

For a union sequence $(H_1, \ldots, H_k)$:

- For each successive $H = A \cup B$, find all H-colored subgraphs by considering all pairs of A-colored and B-colored subgraphs.
- Runtime is $\tilde{O}(\max_H \mathbb{E}[\#\ H\text{-colored subgraphs}]^2)$ w.h.p.

Quadratic improvement with *sort-merge-join*

- Challenge: sorting is not in AC^0 [Håstad'86].

> $n^{\kappa(G)}$ is the maximum over nontrivial input distributions, of the minimum over union sequences $(H_1, \ldots, H_k)$, of $\max_H \mathbb{E}[\#\ H\text{-colored subgraphs}]$.

Our Contributions

- The average-case $[\mathrm{AC}^0, \text{time}]$ complexity is at most $n^{\kappa(G)+O(1)}$.
- If G is a hypercube then $\kappa(G)$ is $\Theta\left(tw(G)/\sqrt{\log tw(G)}\right)$.
- $\kappa(G)$ is Ω(the exponent from Marx's ETH-hardness result).

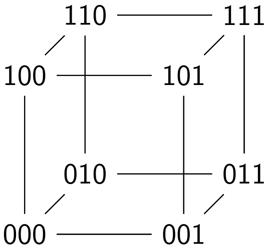

$\kappa(Q_d) = O(2^d/d)$

$\tilde{O}(n^{\kappa(G)})$ is the maximum over nontrivial input distributions, of the minimum over union sequences $(H_1, \ldots, H_k)$, of $\max_H \mathbb{E}[\#\ H\text{-colored subgraphs}]$.

- *Special Case*: Edge density $= n^{-2/d}$ uniformly.
- *General Case*: Reduce to special case via averaging argument.

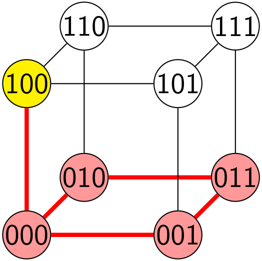

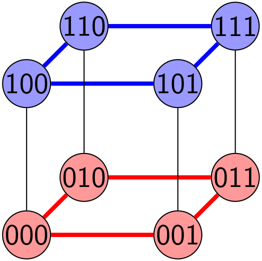

$\kappa(Q_d) = O(2^d/d)$

$\tilde{O}(n^{\kappa(G)})$ is the maximum over nontrivial input distributions, of the minimum over union sequences $(H_1, \ldots, H_k)$, of $\max_H \mathbb{E}[\#\ H\text{-colored subgraphs}]$.

- *Special Case*: Edge density $= n^{-2/d}$ uniformly.
- *General Case*: Reduce to special case via averaging argument.

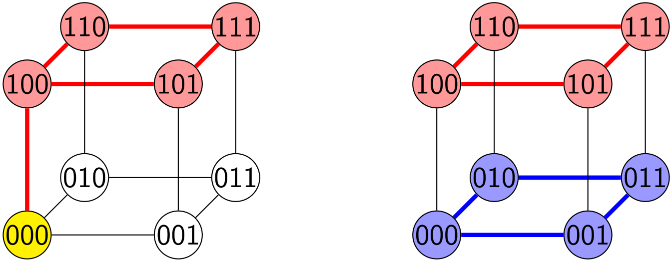

$\kappa(Q_d) = O(2^d/d)$

$\tilde{O}(n^{\kappa(G)})$ is the maximum over nontrivial input distributions, of the minimum over union sequences $(H_1, \ldots, H_k)$, of $\max_H \mathbb{E}[\#\ H\text{-colored subgraphs}]$.

- *Special Case*: Edge density $= n^{-2/d}$ uniformly.
- *General Case*: Reduce to special case via averaging argument.

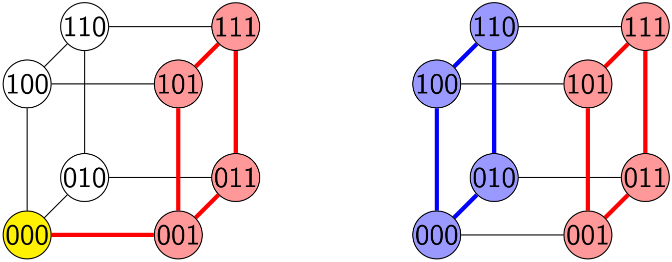

Treewidth of the Hypercube

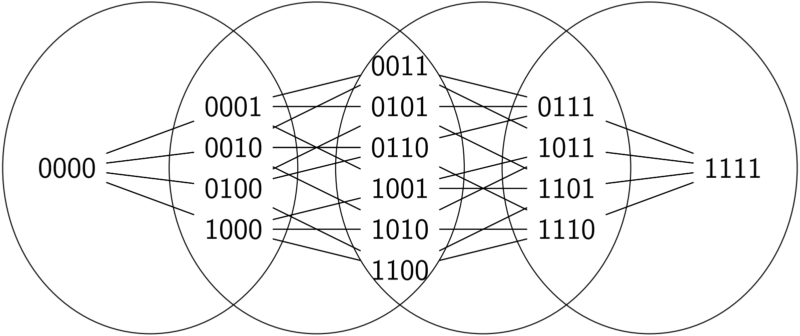

- $tw(Q_d) \lessapprox 2\binom{d}{d/2} = O(2^d/\sqrt{d})$.
- $tw(Q_d)$ is $\Theta(2^d/\sqrt{d})$ [Chandran–Kavitha'06].

Our Contributions

- The average-case $[\mathrm{AC}^0, \text{time}]$ complexity is at most $n^{\kappa(G)+O(1)}$.
- If G is a hypercube then $\kappa(G)$ is $\Theta\left(tw(G)/\sqrt{\log tw(G)}\right)$.
- $\kappa(G)$ is Ω(the exponent from Marx's ETH-hardness result).