Small minors in dense graphs

Samuel Fiorini, Gwenaël Joret
Université Libre de Bruxelles

Dirk O. Theis
Otto-von-Guericke-Universität Magdeburg

David R. Wood
University of Melbourne

Minors and models

H minor of G if H can be obtained from a subgraph of G by contracting edges

H-model in G: collection $\{S_u : u \in V(H)\}$ of vertex-disjoint connected subgraphs of G s.t. $\exists$ edge between S_u and S_v in G for every edge $uv \in E(H)$

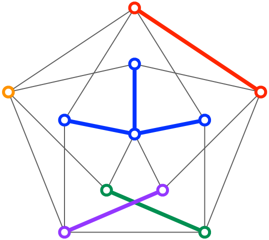

A K_5-model

The S_u's are called vertex images

Small models

$\mathcal{F}$ some class of graphs

H fixed graph

$\mathcal{F}$ has small H-models if $\exists\ c = c(\mathcal{F})$ s.t. every n-vertex graph $G \in \mathcal{F}$ has an H-model with $\leqslant c \cdot \log n$ vertices

Example: $H = K_3$ and $\mathcal{F} = \{G : G \text{ has min. deg.} \geqslant 3\}$

Small models

$\mathcal{F}$ some class of graphs

H fixed graph

$\mathcal{F}$ has small H-models if $\exists\ c = c(\mathcal{F})$ s.t. every n-vertex graph $G \in \mathcal{F}$ has an H-model with $\leqslant c \cdot \log n$ vertices

Example: $H = K_3$ and $\mathcal{F} = \{G : G \text{ has min. deg. } \geqslant 3\}$

This talk: average degree vs. small H-models

Mader's theorem

Theorem (Mader)

If G has average degree $\geqslant 2^{t-2}$ then G contains a K_t-minor

Mader's theorem

Theorem (Mader)

If G has average degree $\geqslant 2^{t-2}$ then G contains a K_t-minor

$\lambda(t) :=$ mininum real r s.t. avg degree $\geqslant r$ implies a K_t-minor

$\lambda(t)$ known for small values of t

t	3	4	5	6	7
$\lambda(t)$	2	4	6	8	10

Theorem (Kostochka, Thomason)

$\lambda(t) = \Theta(t\sqrt{\log t})$

A 'small model' version of Mader's theorem

Theorem

Let $\varepsilon > 0$. If G has average degree $\geqslant 2^{t-1} + \varepsilon$ then G contains a K_t-model with $O_{t,\varepsilon}(\log n)$ vertices.

A 'small model' version of Mader's theorem

Theorem

Let $\varepsilon > 0$. *If* G *has average degree* $\geqslant 2^{t-1} + \varepsilon$ *then* G *contains a* K_t*-model with* $O_{t,\varepsilon}(\log n)$ *vertices.*

Motivation: packing and covering H-models

Packing and covering H-models

H fixed graph

$\nu(G) :=$ packing number
$\quad = $ max. number of vertex-disjoint H-models in G

$\tau(G) :=$ covering number
$\quad = $ min. number of vertices hitting all H-models in G

$\nu(G) \leqslant \tau(G) \quad \forall G$

Packing and covering H-models

H fixed graph

$\nu(G) :=$ packing number
$\quad = \text{max. number of vertex-disjoint } H\text{-models in } G$

$\tau(G) :=$ covering number
$\quad = \text{min. number of vertices hitting all } H\text{-models in } G$

$\nu(G) \leqslant \tau(G) \quad \forall G$

Erdős-Pósa property of H-minors (Robertson and Seymour):

τ bounded from above by a function of $\nu \Leftrightarrow H$ planar

Packing and covering H-models

H fixed graph

$\nu(G)$:= packing number
= max. number of vertex-disjoint H-models in G

$\tau(G)$:= covering number
= min. number of vertices hitting all H-models in G

$\nu(G) \leqslant \tau(G) \quad \forall G$

Erdős-Pósa property of H-minors (Robertson and Seymour):

τ bounded from above by a function of $\nu \Leftrightarrow H$ planar

Known upper bounds $\tau \leqslant f(\nu)$ are big: $f(\nu) \in \Omega(2^{\nu^2})$

A variant of the Erdős-Pósa property

Let's introduce a small dependence in n

Question (“log n property”)

For which graphs H do we have $\tau \leqslant (c \log n)\nu$?

$H = K_2$ ✓

$\tau \leqslant 2\nu$

$H = K_3$ ✓

$\tau \leqslant (2 \log_2 n)\nu$

$H = K_4 - e$ ✓

$\tau \leqslant (6 \log_{3/2} n + 8)\nu$

(Fiorini, J, Pietropaoli)

$H = K_4$?

open

$H = K_5$ ✗

$\sqrt{n} \times \sqrt{n}$ triangulated toroidal grid:

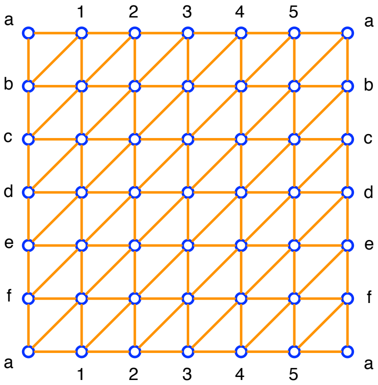

$H = K_5$ ✗

$\sqrt{n} \times \sqrt{n}$ triangulated toroidal grid:

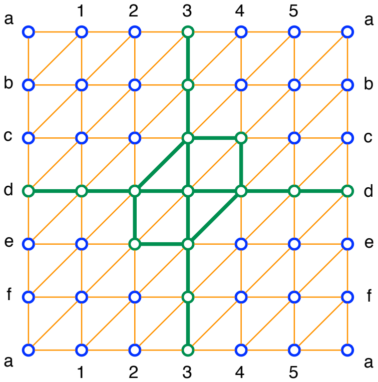

$\nu(G) = 1$ but $\tau(G) = \Theta(\sqrt{n})$

$H = K_5$ ✗

H not planar ✗

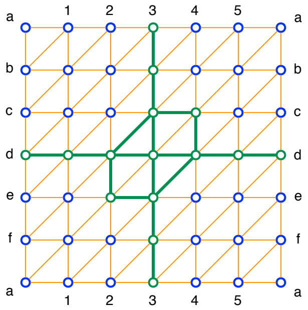

Question

H has the $\log n$ property $\Leftrightarrow$ H planar?

Related optimization problems

H-Minor Packing
INPUT: G
SOLUTION: collection $\mathcal{C}$ of vertex-disjoint H-models
GOAL: maximize $|\mathcal{C}|$

H-Minor Covering
INPUT: G
SOLUTION: $X \subseteq V(G)$ s.t. $G - X$ is H-minor free
GOAL: minimize $|X|$

H-Minor Covering

$H = K_2$

vertex cover problem

- 2-approx.

$H = K_4 - e$

- $O(\log n)$-approx.
- 9-approx.

(Fiorini, J, Pietropaoli)

$H = K_3$

feedback vertex set problem

- $O(\log n)$-approx. (Bar-Yehuda et al.)
- 8-approx. (Even et al.)
- 2-approx. (Bafna et al., Becker et al., Chudak et al.),

H not planar

open

Last week:

H planar

$O(\log^{3/2} n)$-approx. (Fomin et al.)

in fact: $O(\log^{3/2} \tau)$-approx.

Main ingredients:

- treewidth lower bound on τ (excluded wall theorem)
- $\sqrt{\mathrm{tw}(G)}$-approx. alg. for treewidth (Feige et al.)

H-Minor Packing

$\boxed{H = K_2}$

maximum matching

$\boxed{H = K_3}$ $\boxed{H = K_4 - e}$

- $O(\log n)$-approx.

follows from proof of $\tau \leqslant (c \log n)\nu$

$\boxed{H = K_4, \ldots}$

open

H-Minor Packing

$H = K_2$

maximum matching

$H = K_3$ $H = K_4 - e$

- $O(\log n)$-approx.

follows from proof of $\tau \leqslant (c \log n)\nu$

$H = K_4, \ldots$

open

Algorithmic proof of $\tau \leqslant (c \log n)\nu$
$\Rightarrow$ $O(\log n)$-approx. algorithms for H-Minor Packing and H-Minor Covering

A 'small model' version of Mader's theorem

Theorem

Let $\varepsilon > 0$. If G has average degree $\geqslant 2^{t-1} + \varepsilon$ then G contains a K_t-model with $O_{t,\varepsilon}(\log n)$ vertices.

A 'small model' version of Mader's theorem

Theorem

Let $\varepsilon > 0$. If G has average degree $\geqslant 2^{t-1} + \varepsilon$ then G contains a K_t-model with $O_{t,\varepsilon}(\log n)$ vertices.

Proof. Induction on t. True for $t = 2$, assume $t \geqslant 3$.

Ideal situation:

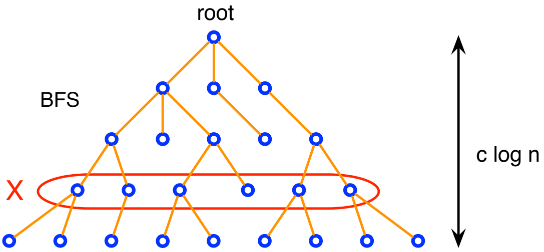

with X inducing a dense subgraph

A ‘small model’ version of Mader’s theorem

Theorem

Let $\varepsilon > 0$. If G has average degree $\geqslant 2^t + \varepsilon$ then G contains a K_t-model with $O_{t,\varepsilon}(\log n)$ vertices.

Proof. Induction on t. True for $t = 2$, assume $t \geqslant 3$.

Ideal situation:

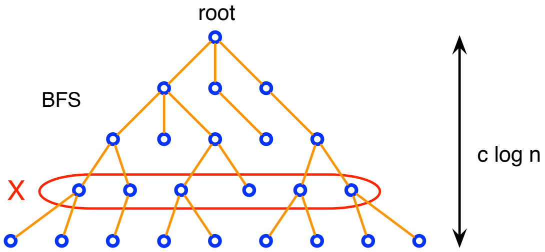

with X inducing a dense subgraph

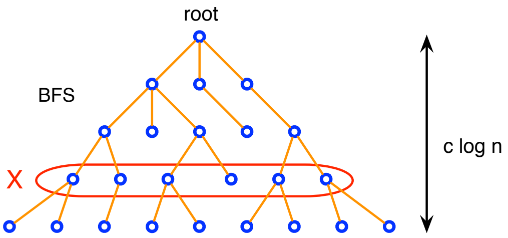

What if G has large diameter?

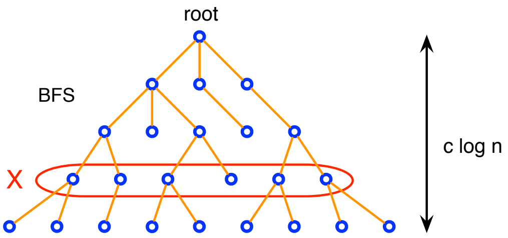

What if G has large diameter?

can find dense subgraph with small diameter:

Lemma

Let $d > d' > 2$. Then every n-vertex graph with average degree $\geqslant d$ has a subgraph with average degree $\geqslant d'$ and diameter $O_{d,d'}(\log n)$

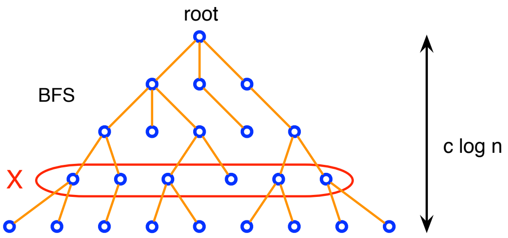

What if G has large diameter?
can find dense subgraph with small diameter:

Lemma

Let $d > d' > 2$. Then every n-vertex graph with average degree $\geqslant d$ has a subgraph with average degree $\geqslant d'$ and diameter $O_{d,d'}(\log n)$

$\Rightarrow$ G has a subgraph with average degree $\geqslant 2^t + \varepsilon/2$ and diameter $O_{t,\varepsilon}(\log n)$

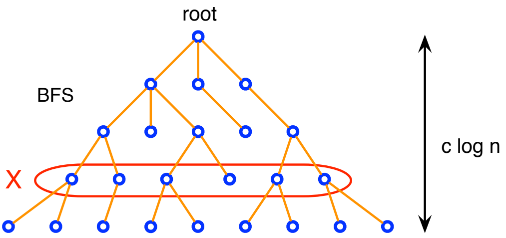

What if G has large diameter?
can find dense subgraph with small diameter:

Lemma

Let $d > d' > 2$. Then every n-vertex graph with average degree $\geqslant d$ has a subgraph with average degree $\geqslant d'$ and diameter $O_{d,d'}(\log n)$

$\Rightarrow$ G has a subgraph with average degree $\geqslant 2^t + \varepsilon/2$ and diameter $O_{t,\varepsilon}(\log n)$ $\rightarrow$ let's focus on such a subgraph

Now G has avg degree $\geqslant 2^t + \varepsilon/2$ and small diameter

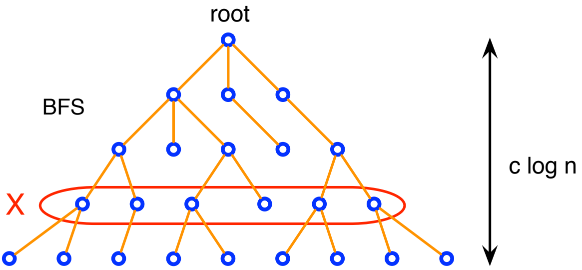

What if $G[X]$ is too sparse for every layer X?

What if $G[X]$ is too sparse for every layer X?

could happen if e.g. G is bipartite:

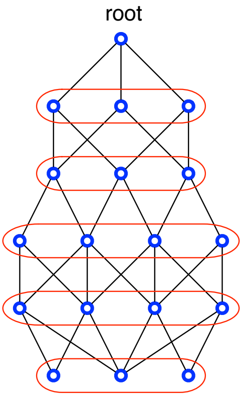

→ let's consider two consecutive layers

$xy \in E(G)$ is even / odd if distance between root and $\{x, y\}$ is even / odd

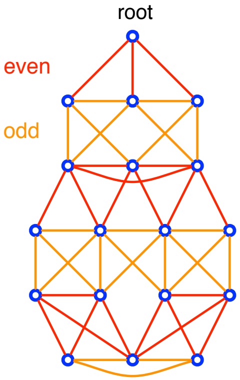

$xy \in E(G)$ is even / odd if distance between root and $\{x, y\}$ is even / odd

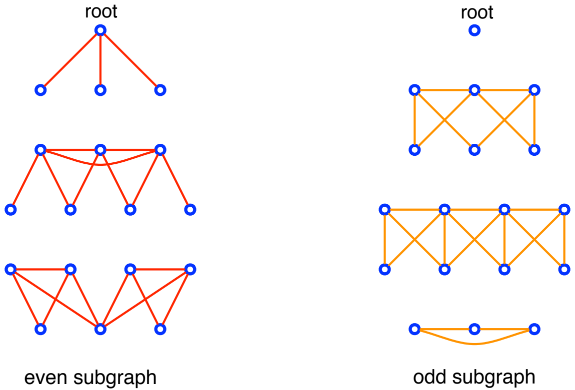

one of these contains $\geqslant$ half of the edges

$\Rightarrow$ some *component* has avg degree $\geqslant \frac{1}{2}(2^t + \varepsilon/2) = 2^{t-1} + \varepsilon/4$

$xy \in E(G)$ is even / odd if distance between root and $\{x, y\}$ is even / odd

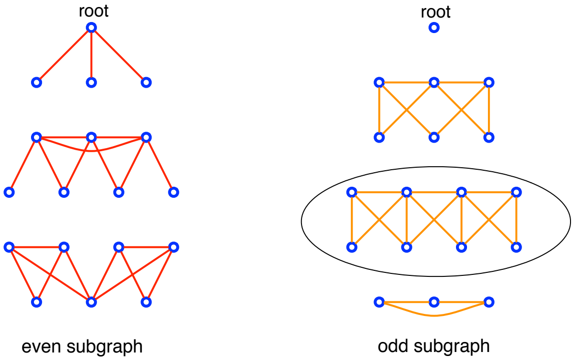

one of these contains $\geqslant$ half of the edges

$\Rightarrow$ some *component* has avg degree $\geqslant \frac{1}{2}(2^t + \varepsilon/2) = 2^{t-1} + \varepsilon/4$

Ideally:

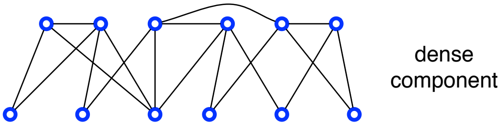

Ideally:

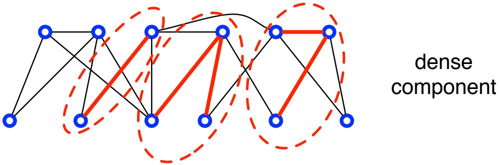

a small K_{t-1}-model

Ideally:

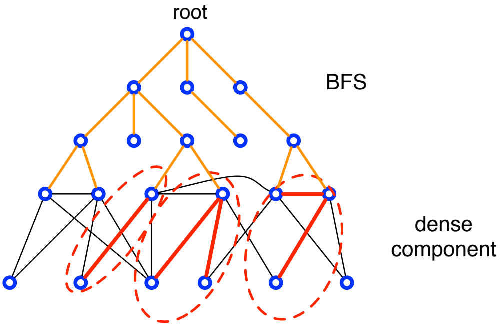

a small K_{t-1}-model

Ideally:

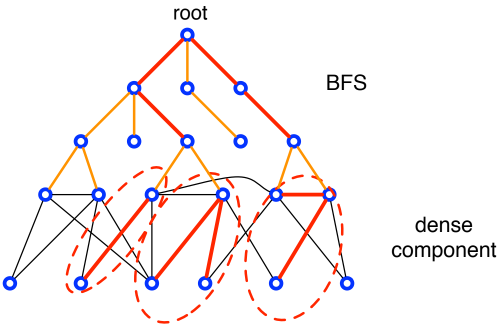

a small K_{t-1}-model
$\Rightarrow$ a small K_t-model

What if some vertex image is contained in the bottom layer?

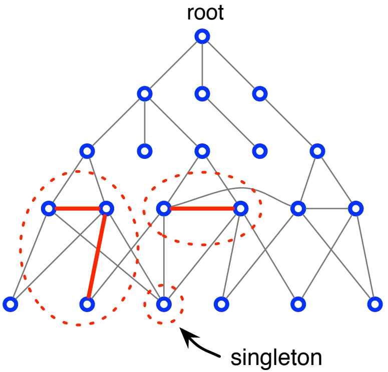

What if some vertex image is contained in the bottom layer?

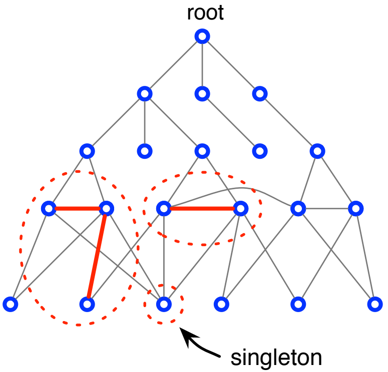

$\longrightarrow$ let's prove a stronger statement:

"avg degree $2^t + \varepsilon \Rightarrow K_t$-model with $O_{t,\varepsilon}(\log n)$ vertices
s.t. no vertex image is a singleton"

What if some vertex image is contained in the bottom layer?

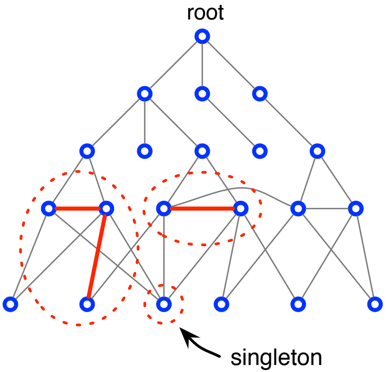

$\longrightarrow$ let's prove a stronger statement:

"avg degree $2^t + \varepsilon \Rightarrow K_t$-model with $O_{t,\varepsilon}(\log n)$ vertices
s.t. no vertex image is a singleton"

⚠ new vertex image from BFS must have $\geqslant 2$ vertices

Root at an edge:

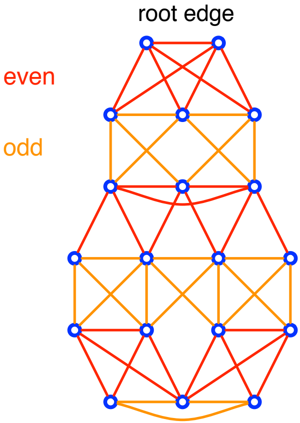

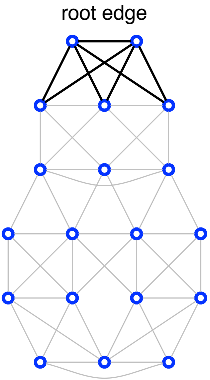

What if our 'dense component' contains the root?

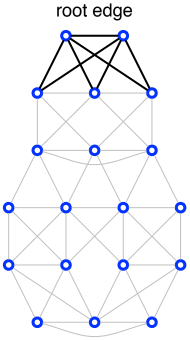

What if our 'dense component' contains the root?

each non-root vertex has degree $\leqslant 2$ in component

$\Rightarrow$ avg degree $\leqslant 4 < 2^{t-1} + \varepsilon/4$ (since $t \geqslant 3$)

$\square$

Reducing the bound to $2^{t-1} + \varepsilon$

Same proof, i.e. show:

"avg degree $2^{t-1} + \varepsilon \Rightarrow K_t$-model with $O_{t,\varepsilon}(\log n)$ vertices s.t. no vertex image is a singleton"

Base case $t = 2$ ✓

Inductive step:

- $t \geqslant 4$ ✓
- $t = 3$ ✗
  because dense component could contain the root edge

$\rightarrow$ prove claim for $t = 3$ separately

Tight bounds for $t \leqslant 4$

- Average degree $\geqslant \varepsilon$ guarantees a small K_2-model
- Average degree $\geqslant 2 + \varepsilon$ guarantees a small K_3-model

Theorem

Average degree $\geqslant 4 + \varepsilon$ guarantees a small K_4-model

Best possible:

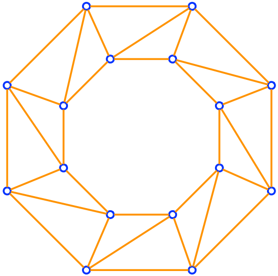

Overview of the proof:

- consider subgraph with avg degree $\geqslant 4 + \varepsilon/2$ and diameter $O_\varepsilon(\log n)$
- remove edge-set of a BFS tree
- resulting graph H has avg degree $\geqslant 2 + \varepsilon/2$
  $\Rightarrow$ H has a $O_\varepsilon(\log n)$ cycle
- consider how the cycle is attached to the BFS tree

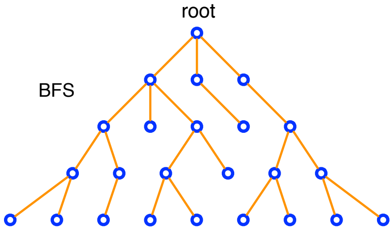

Overview of the proof:

- consider subgraph with avg degree $\geqslant 4 + \varepsilon/2$ and diameter $O_\varepsilon(\log n)$
- remove edge-set of a BFS tree
- resulting graph H has avg degree $\geqslant 2 + \varepsilon/2$
  $\Rightarrow$ H has a $O_\varepsilon(\log n)$ cycle
- consider how the cycle is attached to the BFS tree

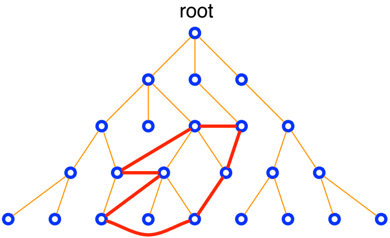

Constant-size K_4-models in planar graphs

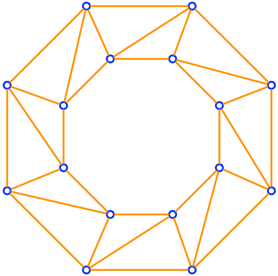

G is 4-connected, planar, and has no small K_4-model

Theorem

Every planar graph with average degree $\geqslant 4 + \varepsilon$ has a K_4-model with $O(1/\varepsilon)$ vertices

Theorem

Every planar graph with minimum degree 5 has a K_4-model with at most 8 vertices

Tight:

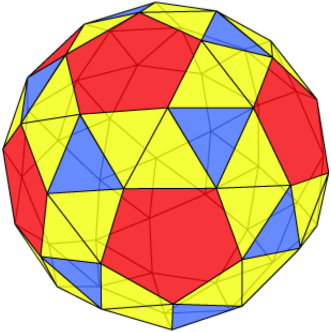

Higher genus surfaces

Theorem

Suppose

- *G has average degree* $\geqslant 4 + \varepsilon$,
- *G is* 3*-connected, and*
- *G has an embedding in a surface of Euler genus g with facewidth* $\geqslant 3$

Then G has a K_4-model with $\leqslant f(\varepsilon) \cdot \log(g + 2)$ *vertices*

Tight up to the value of $f(\varepsilon)$

Open problems

- is there a polynomial function $f(t)$ s.t. average degree $\geqslant f(t)$ guarantees a small K_t-model?

Open problems

- is there a polynomial function $f(t)$ s.t. average degree $\geqslant f(t)$ guarantees a small K_t-model?
- what about $t = 5$? $2^{t-1} + \varepsilon = 16 + \varepsilon$ is enough, 6 is not:

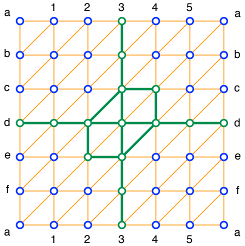

Open problems

- is there a polynomial function $f(t)$ s.t. average degree $\geqslant f(t)$ guarantees a small K_t-model?
- what about $t = 5$? $2^{t-1} + \varepsilon = 16 + \varepsilon$ is enough, 6 is not:

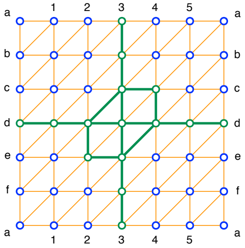

- characterize graphs H s.t. $\tau \leqslant (c \log n)\nu$

Open problems

- is there a polynomial function $f(t)$ s.t. average degree $\geqslant f(t)$ guarantees a small K_t-model?
- what about $t = 5$? $2^{t-1} + \varepsilon = 16 + \varepsilon$ is enough, 6 is not:

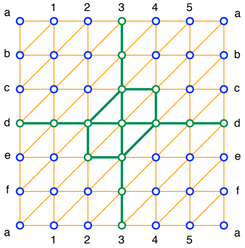

- characterize graphs H s.t. $\tau \leqslant (c \log n)\nu$
- For which H does H-MINOR PACKING / H-MINOR COVERING admits
  - an $O(\log n)$-approx. algorithm?
  - a constant-factor approx. algorithm?