The Complexity of Satisfaction on Sparse Graphs

Anuj Dawar

University of Cambridge Computer Laboratory

IPEC, Chennai, 15 December 2010

Complexity of the Satisfaction Relation

We are interested in the complexity of the *satisfaction relation* for *first-order logic*.

That is, given a graph G

and a formula φ of first-order logic in the *language of graphs*:

$$E(x, y) \mid \varphi \wedge \psi \mid \varphi \vee \psi \mid \neg\varphi \mid \exists x\varphi \mid \forall x\varphi$$

decide whether $G \models \varphi$

What is the complexity of this, *parameterized* by the length of φ?

In the rest of this talk, we use n for the size of G, l for the length of φ and m for the depth of nesting of quantifiers in φ.

Complexity of Satisfaction

The satisfaction relation is AW$[\star]$-complete.

The naive algorithm takes time $O(ln^m)$.

Restricted to formulas with at most t alternations of quantifiers, the problem is $W[t]$-hard.

It subsumes many natural parameterized graph problems (where, fixing the parameter, we can express the problem in first-order logic).

- *Vertex Cover*;
- *Independent Set*;
- *Dominating Set*;
- *Network Centres* (distance-d dominating set).

Restricted Classes

One way to get a handle on the complexity of first-order satisfaction is to consider restricted graph classes.

Given: a first-order formula φ and a graph $G \in \mathcal{C}$
Decide: if $G \models \varphi$

For many classes $\mathcal{C}$, this problem has been shown to be FPT.

1. $\mathcal{T}_k$—the class of graphs of tree-width at most k.
   Courcelle (1990) shows that MSO satisfaction is fixed parameter linear time on this class.
2. $\mathcal{D}_k$—the class of graphs of *degree* bounded by k. **Seese (1996)** shows that FO satisfaction is fixed-parameter linear time.

Some Results

3. LTW_t—the class of graphs of *local tree-width* bounded by a function t. **Frick and Grohe (2001)** show that FO satisfaction is fixed parameter quadratic time.

4. $\mathcal{M}_k$—the class of graphs *excluding K_k as a minor*. **Flum and Grohe (2001)** show that FO satisfaction is $O(f(l)n^5)$.

5. LEM_t—the class of structures with *locally excluded minors* given by t. **D., Grohe and Kreutzer (2007)** show that FO satisfaction is $O(f(l)n^6)$.

6. On any class of *bounded expansion* **Dvořak, Král and Thomas; D. and Kreutzer (2010)** show that FO satisfaction is fixed parameter linear time.

7. On any class of graphs that is *nowhere dense*, **D. and Kreutzer (2009)** show that *some independence* and *domination* problems are decidable in fixed parameter linear time.

Map of Restrictions

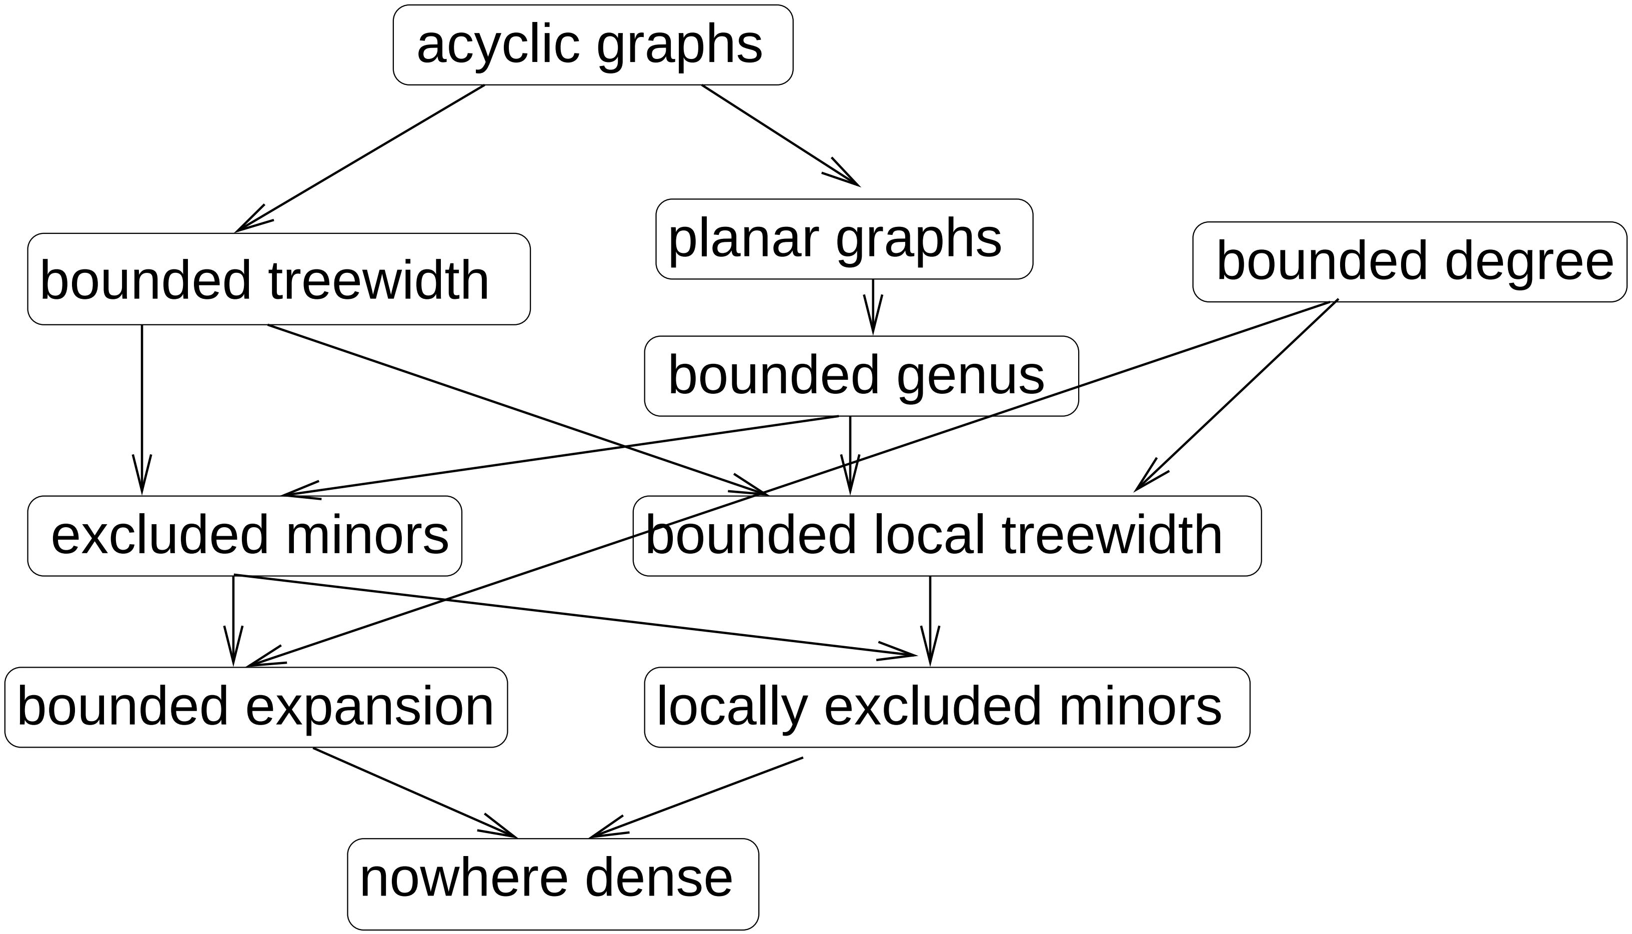

Automata and Locality

The methods of proof for all but the last two of the above results are combinations of two general techniques:

- Methods of *automata* or *decompositions*; and
- Methods based on the *locality* of first-order logic.

The last two are based on new techniques:

- A method based on *low-depth* colourings.
- A method based on *wideness* of the classes.

In the rest of the talk, I'll give a brief overview of the first two methods and then discuss classes of *bounded expansion* and *nowhere dense* classes.

Courcelle's Theorem

Theorem (Courcelle)

MSO (or MS_2) satisfaction is fixed parameter linear time on $\mathcal{T}_k$.

Write $G, \bar{u} \equiv_m H, \bar{v}$ to denote that any formula of FO that is true of the tuple of vertices $\bar{u}$ in G is true of $\bar{v}$ in H and *vice versa*.

For each k, there is a *finite* collection of operations such that any $G \in \mathcal{T}_k$ can be built up from graphs with $\leq k$ vertices using these operations. Moreover, $\equiv_m$ is a *congruence* for each of them.

Pointed Sum

$G_1 \oplus_X G_2$

$|X| \leq k$

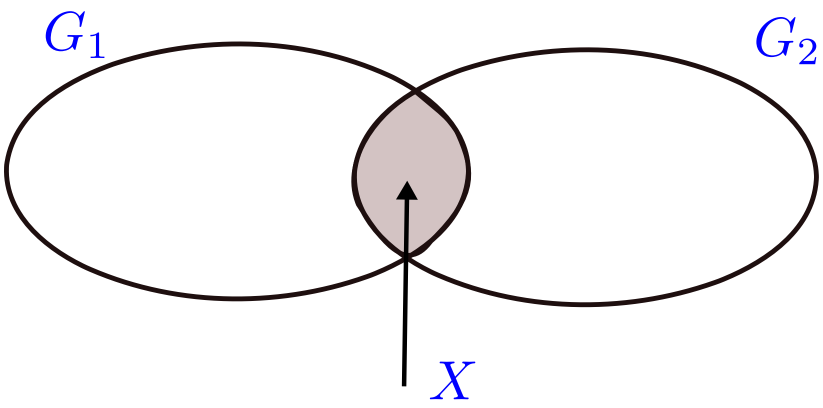

If $G_1, X \equiv_m H_1, Y$ and $G_2, X \equiv_m H_2, Y$ then

$$G_1 \oplus_X G_2, X \equiv_m H_1 \oplus_Y H_2, Y.$$

The Method of Decompositions

Suppose FO satisfaction is FPT on a class $\mathcal{B}$

and $\mathcal{C}$ is a class of graphs such that there is a finite collection Op of operations such that:

- $\mathcal{C}$ is contained in the closure of $\mathcal{B}$ under the operations in Op;
- there is a polynomial-time algorithm which computes, for any $G \in \mathcal{C}$, an Op-decomposition of G over $\mathcal{B}$; and
- for each m, the equivalence class $\equiv_m$ is an *effective* congruence with respect to to all operations $o \in$ Op (i.e., the $\equiv_m$-type of $o(G_1, \ldots, G_s)$ can be computed from the $\equiv_m$-types of $G_1, \ldots, G_s$).

Then, FO satisfaction is fixed-parameter tractable on $\mathcal{C}$.

Bounded Degree Graphs

$\mathcal{D}_k$—the class of graphs G in which every element has degree at most k.

Theorem (Seese)

For every sentence φ of FO and every k there is a linear time algorithm which, given a graph $G \in \mathcal{D}_k$ determines whether $G \models \varphi$.

Note: this is not true for MSO unless P $=$ NP.

A proof is based on *locality* of first-order logic.

Gaifman's Theorem

We write $\delta(x, y) > d$ for the formula of FO that says that the distance between x and y is greater than d.

We write $\psi^r(x)$ to denote the formula obtained from $\psi(x)$ by relativising all quantifiers to the set $N_r = \{y \mid \delta(x, y) < r\}$.

A *basic local sentence* is a sentence of the form

$$\exists x_1 \cdots \exists x_s \left(\bigwedge_{i \neq j} \delta(x_i, x_j) > 2r \wedge \bigwedge_i \psi^r(x_i) \right)$$

Theorem (Gaifman)

Every first-order sentence is equivalent to a Boolean combination of basic local sentences.

Method of Locality

- Suppose we have a function, associating a parameter $k_G \in \mathbb{N}$ with each graph G.
- Suppose we have an algorithm which, given G and φ decides $G \models \varphi$ in time

$$g(l, k_G)n^c$$

for some computable function g and some constant c.

- Let $\mathcal{C}$ be a class of graphs of *bounded local* k, i.e.

  there is a computable function $t : \mathbb{N} \to \mathbb{N}$ such that for every $G \in \mathcal{C}$ and $v \in G$, $k_{N_r(a)} < t(r)$.

Then, there is an algorithm which, given $G \in \mathcal{C}$ and φ decides whether $G \models \varphi$ in time

$$f(l)n^{c+1}$$

for some computable function f.

Bounded Local Treewidth

Let $t : \mathbb{N} \to \mathbb{N}$ be a non-decreasing function.

LTW_t—the class of graphs G such that for every $v \in V(G)$:

$N_r^G(v)$ has tree-width at most $t(r)$. **(Eppstein; Frick-Grohe)**.

We say that $\mathcal{C}$ has *bounded local tree-width* if there is some function t such that $\mathcal{C} \subseteq \mathsf{LTW}_t$.

Examples:

1. $\mathcal{T}_k$ has local tree-width bounded by the constant function $t(r) = k$.
2. $\mathcal{D}_k$ has local tree-width bounded by $t(r) = k^r + 1$.
3. Planar graphs have local tree-width bounded by $t(r) = 3r$.

Excluded Minor Classes

Write $\mathcal{M}_k$ for the class of graphs G such that $K_k \not\preceq G$.

Grohe shows that graphs in $\mathcal{M}_k$ can be decomposed over graphs of *almost bounded local tree-width*.

By suitable definitions of congruences (*non-trivial*) we get an application of the *method of decompositions*.

First-order logic is *fixed-parameter tractable* on $\mathcal{M}_k$.

(Flum-Grohe)

Locally Excluded Minors

Say a class of graphs $\mathcal{C}$ *locally excludes minors* if there is some non-decreasing function $t : \mathbb{N} \to \mathbb{N}$ such that for any $G \in \mathcal{C}$ and any $v \in G$,

$$K_{t(r)} \not\preceq N_r(v)$$

Theorem (D., Grohe, Kreutzer)

First-order logic is fixed-parameter tractable on every class $\mathcal{C}$ that locally excludes minors.

The result would be an easy application of the *locality method* except that the proof of Flum and Grohe relies on non-constructive methods.

A new decomposition of $\mathcal{M}_k$ is needed.

Map of Restrictions

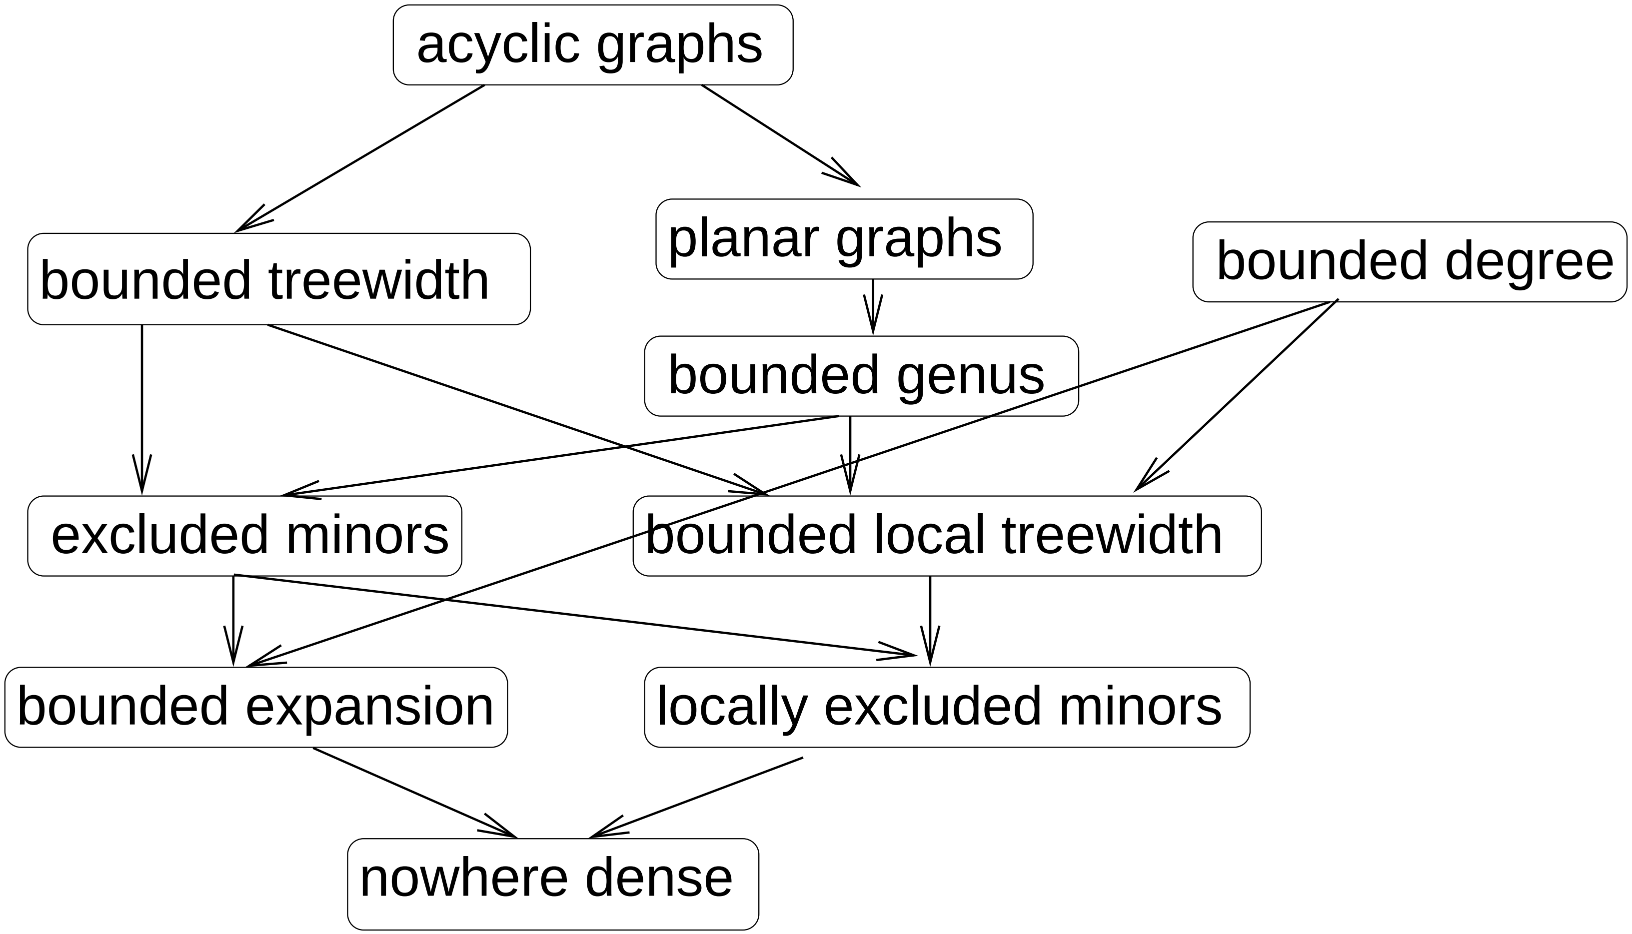

Sparse Classes

For a graph G, we write $|G|$ for the number of *vertices* in G and $||G||$ for the numer of *edges* in G.

If G has degree at most d, then $||G|| \leq \frac{d}{2}|G|$.

If G has tree-width at most k, then $||G|| \leq k|G|$.

Theorem

There exists a function f such that, for every r, every graph G of average degree $d \geq f(r)$ contains K_r as a minor.

Sparse Classes

Say a class $\mathcal{C}$ of graphs is *sparse* if there is a c so that for all $G \in \mathcal{C}$, $||G|| \leq c|G|$.

Equivalently, $\mathcal{C}$ has *bounded average degree*.

There are pathological sparse classes.

Take the class that contains, for every finite graph G the graph consisting of the union of G with $||G||$ isolated vertices.

Hereditarily Sparse Classes

Say a class $\mathcal{C}$ of graphs is *hereditarily sparse* if there is a c so that for all $G \in \mathcal{C}$, and every subgraph $H \subset G$ we have $||H|| \leq c|H|$.

With any graph $G = (V, E)$, we associate its *incidence graph* $I(G) = (V \cup E, F)$ where

$$F = \{(v, e) \mid v \in V, e \in E \text{ and } e \text{ is incident with } v\}.$$

The collection of all incidence graphs is hereditarily sparse but has all the complexity of the class of all graphs since the map $G \mapsto I(G)$ is an easy *first-order* reduction.

Shallow Minors

Recall that $H = (U, F)$ is a minor of $G = (V, E)$, if we can find a collection of *disjoint, connected* subgraphs of G: $(B_u \mid u \in U)$ such that whenever $(u_1, u_2) \in F$, there is an edge between some vertex in B_{u_1} and some vertex in B_{u_2}.

The graphs B_u are called *branch sets* witnessing that $H \preceq G$.

If the branch sets can be chosen so that for each u there is $b \in B_u$ and $B_u \subseteq N_r^G(b)$, we say that H is a minor *at depth* r of G and write $H \preceq_r G$

Bounded Expansion

A class $\mathcal{C}$ of graphs is said to have *bounded expansion* if, for each r, there is a ∇_r such that if $H \preceq_r G$ for some $G \in \mathcal{C}$ then $||H|| \leq \nabla_r |H|$.

In other words, $\mathcal{C}$ has bounded expansion if, for every r, the collection of depth-r minors of graphs in $\mathcal{C}$ is *sparse*.

$\mathcal{M}_k$ has bounded expansion, with ∇_r depending only on k.

$\mathcal{D}_k$ has bounded expansion by taking $\nabla_r = k^r$.

FO on Bounded Expansion Classes

A very recent result (FOCS 2010) by **Dvořak, Král and Thomas** states:

> If $\mathcal{C}$ is a class of bounded expansion then FO satisfaction on $\mathcal{C}$ is in *fixed parameter linear time*.
>
> **Note:** this improves the $O(n^5)$ bound for excluded minor classes.

The technique used is quite different to the locality and decomposition techniques we have seen. It relies on suitable *graph colourings*.

Tree Depth

The *tree-depth* of a graph G is defined to be the smallest k such that there is *directed forest* F of height k and

> G is a subgraph of the *undirected* graph underlying the *transitive closure* of F.

For any graph G, $\mathrm{tw}(G) \leq \mathrm{td}(G)$,

> where $\mathrm{tw}(G)$ is the tree-width of G and $\mathrm{td}(G)$ is the tree-depth of G.

Low Tree-Depth Colourings

Nešetřil and Ossona de Mendez prove a remarkable colouring property of classes of graphs of *bounded expansion*.

Theorem: (Nešetřil, Ossona de Mendez)

Let $\mathcal{C}$ be a class of graphs of bounded expansion. For any p there is an N such that any graph $G \in \mathcal{C}$ can be coloured using N colours in such a way that if $C_1, \ldots, C_p$ is any set of p colours then $G[C_1 \cup \cdots \cup C_p]$ has *tree-depth* less than p.

Moreover, this colouring can be found efficiently—*in linear time*.

Evaluating Existential Formulas

Suppose φ is an *existential* first-order formula.

That is, it is of the form

$$\exists x_1 \cdots \exists x_q \theta$$

where θ is quantifier-free.

If $\mathcal{C}$ has bounded expansion, we can evaluate such formulas on graphs in $\mathcal{C}$ by the following process.

Find a colouring of G which guarantees that any q colours induce a graph of tree-depth at most q.

For each set of q colours, check whether φ can be evaluated in the subgraph induced by these colours.

This establishes that for *existential* formulas, satisfaction is FPT on $\mathcal{C}$.

Quantifier Alternation

Moreover, within a graph of *tree-depth* at most d, the $\equiv_m$-type of a tuple $v_1, \ldots, v_k$ ($k + m < d$) is determined by the types of the individual vertices v_i and, for each pair i, j, the *height* of the *least common ancestor* of v_i and v_j in the witnessing forest.

This allows us to turn a *universal* formula

$$\exists x_1 \cdots \exists x_p \varphi(\bar{x}, \bar{y})$$

into an equivalent existential formula

$$\bigwedge_{S \in [\bar{C}]^p} \exists \bar{z} \psi_S(\bar{y}, \bar{z})$$

whose length depends on N^p.

Nowhere Dense Graphs

We say that a class $\mathcal{C}$ of graphs is *nowhere-dense* if, for every r, the collection of graphs

$$\{H \mid H \preceq_r G \text{ for some } G \in \mathcal{C}\}$$

is *not* the class of all graphs.

In other words, for each r, there is a K_k that cannot be obtained as a depth-r minor of any graph in $\mathcal{C}$.

This clearly generalizes bounded expansion classes.

It also generalizes locally excluded minor classes because if $K_k \preceq_r G$ then there is a v in G such that $K_k \preceq N^G_{3r+1}(v)$.

Trichotomy Theorem

Associate with any infinite class $\mathcal{C}$ of graphs the following parameter:

$$d_{\mathcal{C}} = \lim_{r \to \infty} \limsup_{G \in \mathcal{C}_r} \frac{\log ||G||}{\log |G|},$$

where $\mathcal{C}_r$ is the collection of graphs obtained as minors of a graph in $\mathcal{C}$ by contracting neighbourhoods of radius at most r.

The *trichotomy theorem* of Nešetřil and Ossona de Mendez states that $d_{\mathcal{C}}$ can only take values 0, 1 and 2.

The nowhere-dense classes are exactly the ones where $d_{\mathcal{C}} \neq 2$.

This shows that these classes are a *natural limit* to one notion of sparseness.

FO on Nowhere Dense Classes

It is still an open question whether FO satisfaction is fixed-parameter tractable on nowhere-dense classes.

Some problems, defined by families of FO formulas, have been shown to be FPT on such classes.

- *Independent Set*;
- *Dominating Set*;
- *distance-d dominating set*

The proof for these is based on a technique distinct from those we have seen so far.

Wide Classes

A set of vertices A in a graph G is said to be r-scattered if for any $u, v \in A$, $N_r(u) \cap N_r(v) = \emptyset$.

Definition

A class of graphs $\mathcal{C}$ is said to be *wide* if for every r and m there is an N such that any graph in $\mathcal{C}$ with more than N vertices contains a r-scattered set of size m.

Example: Classes of graphs of bounded degree.

Non-Example: Trees

Almost Wide Classes

Definition

A class of graphs $\mathcal{C}$ is *almost wide* if there is an s such that for every r and m there is an N such that any graph in $\mathcal{C}$ with more than N vertices contains s elements whose removal leaves a r-scattered set of size m.

Example: Trees.

Examples: planar graphs?

Quasi-Wide Classes

Let $s : \mathbb{N} \to \mathbb{N}$ be a function. A class $\mathcal{C}$ of graphs is *quasi-wide with margin s* if for all $r \geq 0$ and $m \geq 0$ there exists an $N \geq 0$ such that if $G \in \mathcal{C}$ and $|G| > N$ then there is a set S of vertices with $|S| < s(r)$ such that $G - S$ contains an r-scattered set of size at least m.

We can show that any class of nowhere-dense graphs is quasi-wide.

The proof also shows that any class that excludes K_k as a minor is *almost wide* with margin $k - 2$.

Equivalence

A class $\mathcal{C}$ of graphs is *uniformly quasi-wide* with margin s if for all $r \geq 0$ and all $m \geq 0$ there exists an $N \geq 0$ such that if $G = (V, E) \in \mathcal{C}$ and $W \subseteq V$ with $|W| > N$ then there is a set $S \subseteq V$ with $|S| < s(r)$ such that W contains an r-scattered set of size at least m in $G - S$.

Theorem: (Nešetřil, Ossona de Mendez)

The following are equivalent for any class $\mathcal{C}$ that is closed under taking induced subgraphs:

1. $\mathcal{C}$ is nowhere dense
2. $\mathcal{C}$ is quasi-wide
3. $\mathcal{C}$ is uniformly quasi-wide

Review

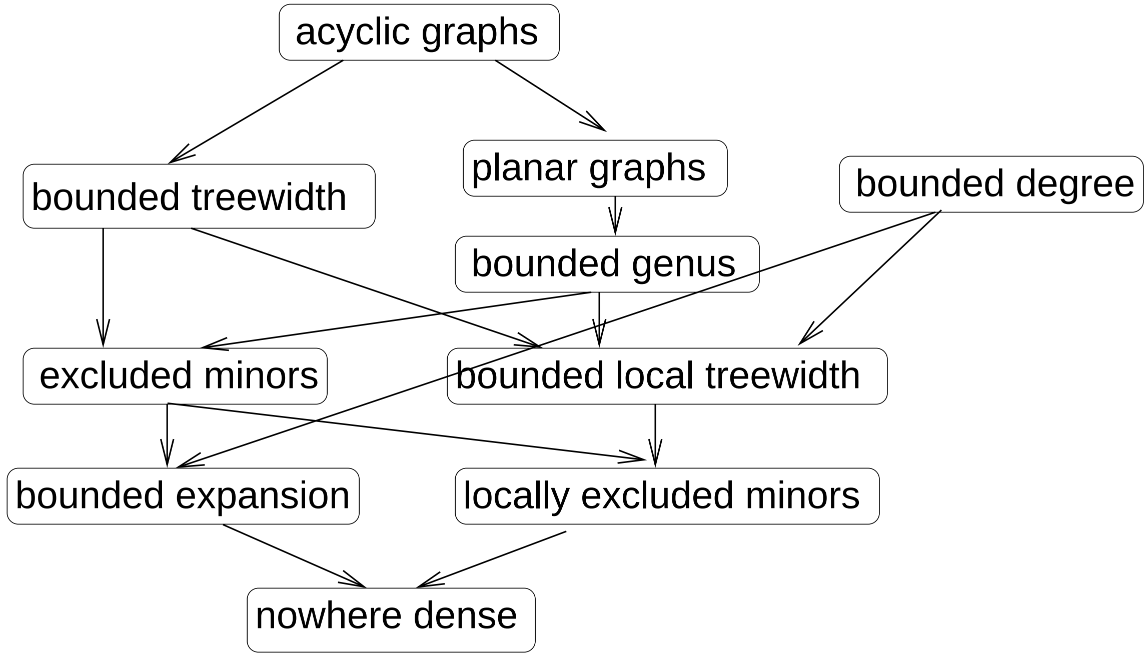

For all classes except the last one the picture, it has been established that FO satisfaction is FPT.

For *nowhere dense classes* this remains an open question.

Techniques deployed use: *locality*, *decompositions*, *low tree-depth colourings* and *wideness*.