Evaluating Formulas on Sparse Graphs

Part 2

Anuj Dawar

University of Cambridge Computer Laboratory

PhD Open, Warsaw, 2 October 2010

Quick Review

We consider the complexity of the problem of deciding,

> Given a graph G and a formula φ
>
> whether $G \models \varphi$

when φ is either in FO or MSO.

In general the problem is PSPACE-complete and AW$[\star]$-hard.

When we consider words instead of graphs it is FPT.

We now aim to identify classes of *sparse* graphs where the problem becomes tractable.

Graph Structure Theory

Graph Structure Theory has developed rapidly since the 1980s through the work of **Robertson, Seymour** and their collaborators on *graph minors*.

One important contribution is the notion of *treewidth*.

This associates a number $\mathrm{tw}(G) = k$ with every graph G which measures how densely interconnected G is.

The measure has many equivalent definitions, and has arisen independently in many contexts.

Treewidth

The *treewidth* of an undirected graph is a measure of how tree-like the graph is.

A graph has treewidth k if it can be covered by subgraphs of at most $k + 1$ nodes in a tree-like fashion.

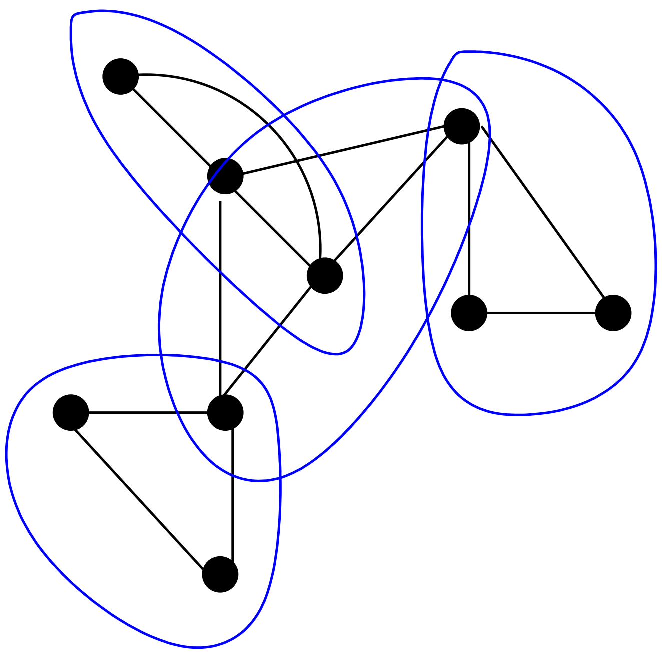

This gives a *tree decomposition* of the graph.

Treewidth

Treewidth is a measure of how *tree-like* a structure is.

For a graph $G = (V, E)$, a *tree decomposition* of G is a relation $D \subset V \times T$ with a tree T such that:

- for each $v \in V$, the set $\{t \mid (v, t) \in D\}$ forms a connected subtree of T; and
- for each edge $(u, v) \in E$, there is a $t \in T$ such that $(u, t), (v, t) \in D$.

The *treewidth* of G is the least k such that there is a tree T and a tree decomposition $D \subset V \times T$ such that for each $t \in T$,

$$|\{v \in V \mid (v, t) \in D\}| \leq k + 1.$$

Treewidth

Looking at the decomposition *bottom-up*, a graph of treewidth k is obtained from graphs with at most $k+1$ nodes through a finite sequence of applications of the operation of taking *sums over sets* of at most k elements.

$$G_1 \oplus_X G_2$$

$$|X| \leq k$$

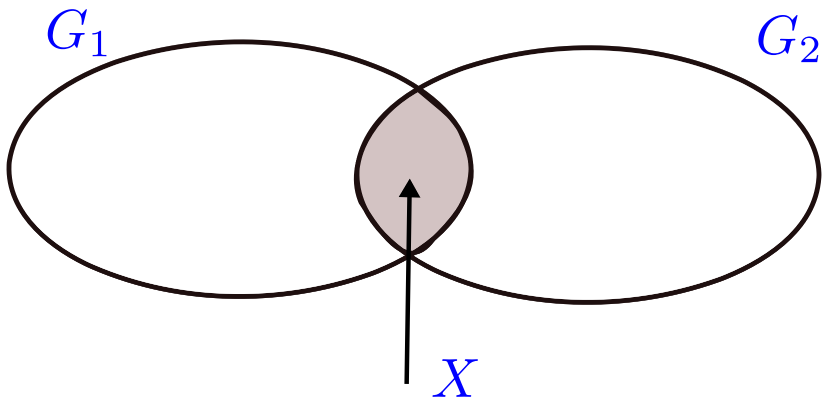

We let $\mathcal{T}_k$ denote the class of graphs G such that $\mathrm{tw}(G) \leq k$.

Treewidth

More formally,

Consider graphs with up to $k+1$ distinguished vertices $\mathbf{c} = c_0, \ldots, c_k$.

Define a *merge* operation $(G \oplus_{\mathbf{c}} H)$ that forms the union of G and H *disjointly apart from* $\mathbf{c}$.

Also define $\text{erase}_i(G)$ that erases the name c_i.

Then a graph G is in $\mathcal{T}_k$ if it can be formed from graphs with at most $k+1$ vertices through a sequence of such operations.

Examples

- Trees have treewidth 1.
- Cycles have treewidth 2.
- The clique K_k has treewidth $k-1$.
- The $m \times n$ grid has treewidth $\min(m, n)$.

Exercise

Graphs of Small Treewidth are Sparse

If $\mathrm{tw}(G) \leq k$ then G has at most $k \cdot |V(G)|$ edges.

This follows from the facts:

- if $\mathrm{tw}(G) \leq k$ then G contains a vertex with at most k neighbours;
- if $G \subset H$ then $\mathrm{tw}(G) \leq \mathrm{tw}(H)$.

Dynamic Programming

It has long been known that graphs of small treewidth admit efficient *dynamic programming* algorithms for intractable problems.

In general, these algorithms proceed bottom-up along a tree decomposition of G.

At any stage, a small set of vertices form the "*interface*" to the rest of the graph.

This allows a recursive decomposition of the problem.

Computing Treewidth

The problem of deciding, given a graph G and an integer k whether $\mathrm{tw}(G) \leq k$ is NP-complete.

But, it is fixed-parameter tractable with k as parameter.

This follows from a theorem of **Bodlaender** that there is an algorithm running in $O(2^{p(k)}n)$ time that given a graph $G \in \mathcal{T}_k$ computes a tree decomposition of G of width k.

Courcelle's Theorem

Theorem (Courcelle)

For any MSO (or MS_2) sentence φ and any k there is a linear time algorithm that decides, given $G \in \mathcal{T}_k$ whether $G \models \varphi$.

Given $G \in \mathcal{T}_k$ and φ, compute:

- from G a labelled tree T; and
- from φ a bottom-up tree automaton $\mathcal{A}$

such that $\mathcal{A}$ accepts T if, and only if, $G \models \varphi$.

The Labelled Tree

$C = \{c_0, \ldots, c_k\}$ a set of $k+1$ new constants.

(G, ρ)—expansion of G with $\rho : C \rightharpoonup V$, a partial map interpreting some of the constants in C.

Let

- $\mathcal{B}_k$—the collection of (G, ρ) such that G has at most $k+1$ vertices.
- erase_i—an operation which takes (G, ρ) to (G, ρ'), where ρ' is as ρ but without c_i.
- a binary operation of union disjoint over C:

$$(G_1, \rho_1) \oplus_C (G_2, \rho_2)$$

Congruence

- Any $G \in \mathcal{T}_k$ is obtained from $\mathcal{B}_k$ by finitely many applications of the operations erase_i and $\oplus_C$.
- If $G_1, \rho_1 \equiv_m^{\mathsf{MSO}} G_2, \rho_2$, then

$$\mathsf{erase}_i(G_1, \rho_1) \equiv_m^{\mathsf{MSO}} \mathsf{erase}_i(G_2, \rho_2)$$

- If $G_1, \rho_1 \equiv_m^{\mathsf{MSO}} G_2, \rho_2$, and $H_1, \sigma_1 \equiv_m^{\mathsf{MSO}} H_2, \sigma_2$ then

$$(G_1, \rho_1) \oplus_C (H_1, \sigma_1) \equiv_m^{\mathsf{MSO}} (G_2, \rho_2) \oplus_C (H_2, \sigma_2)$$

Note: a special case of this is that $\equiv_m^{\mathsf{MSO}}$ is a congruence for *disjoint union* of graphs.

Satisfaction on $\mathcal{T}_k$

Any $G \in \mathcal{T}_k$ can be represented as a finite tree, with leaves labelled by elements of $\mathcal{B}_k$, internal nodes labelled by operations erase$_i$ and $\oplus_C$.

We can then compute the $\text{Type}^{\text{MSO}}(G)$ bottom-up.

This establishes the following:

> The satisfaction problem for MSO is decidable in time $f(l, k)n$, where
>
> - f is some computable function
> - l is the length of the input formula
> - k is the treewidth of the input structure
> - n is the size of the input structure.

The Method of Decompositions

Suppose $\mathcal{C}$ is a class of graphs such that there is a finite class $\mathcal{B}$ and 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 $\equiv_m^{(\text{MSO})}$ is an *effective* congruence with respect to to all operations $o \in$ Op (i.e., the $\equiv_m^{(\text{MSO})}$-type of $o(G_1, \ldots, G_s)$ can be computed from the $\equiv_m^{(\text{MSO})}$-types of $G_1, \ldots, G_s$).

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

Relaxations of the Method

1. Instead of requiring $\mathcal{B}$ be finite, it sufficers to require that *satisfaction is in FPT over* $\mathcal{B}$.
2. In place of $\equiv_m^{(\mathrm{MSO})}$, we can take any sequence of equivalence relations $\sim_m$ $(m \in \mathbb{N})$ satisfying
   - for every φ there is an m such that models of φ are closed under $\sim_m$; and
   - for all m, $\sim_m$ has finite index.

Note: letting $G \sim_m H$ if G, H cannot be distinguished by a formula of *length* m, does not yield a congruence with respect to disjoint union.

There is no elementary function e such that $G_1 \sim_{e(m)} H_1$ and $G_2 \sim_{e(m)} H_2$ implies $G_1 \oplus G_2 \sim_m H_1 \oplus H_2$.

(D.,Grohe, Kreutzer, Schweikardt)

Bounded Degree Graphs

In a graph $G = (V, E)$ the *degree* of a vertex $v \in V$ is the number of neighbours of v, i.e.

$$|\{u \in V \mid (u, v) \in E\}|.$$

We write $\delta(G)$ for the *smallest* degree of any vertex in G.

We write $\Delta(G)$ for the *largest* degree of any vertex in G.

$\mathcal{D}_k$—the class of graphs G with $\Delta(G) \leq k$.

Bounded Degree Graphs

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 $G \models \varphi$.

A proof is based on *locality* of first-order logic, which we look at next.

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

Construct, for any graph G, a graph G' such that $\Delta(G') \leq 5$ and G' is 3-colourable iff G is, and the map $G \mapsto G'$ is polynomial-time computable.

Gaifman's Locality 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 $N_r = \{y \mid \delta(x, y) < r\}$, i.e.

Each subformula $\exists y\theta$ is replaced by $\exists y(\delta(x, y) < r) \wedge \theta^r$

Each subformula $\forall y\theta$ is replaced by $\forall y(\delta(x, y) < r) \rightarrow \theta^r$

Gaifman's Locality Theorem

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.

Uses of Gaifman's Locality Theorem

Gaifman's theorem is often used to give simple proofs that some problems are *not expressible* in first-order logic.

To illustrate the undefinability of *connectivity* consider (for any r, q) the following two graphs chosen big enough so that any basic local sentence of *radius* r and *quantifier rank* q cannot distinguish them.

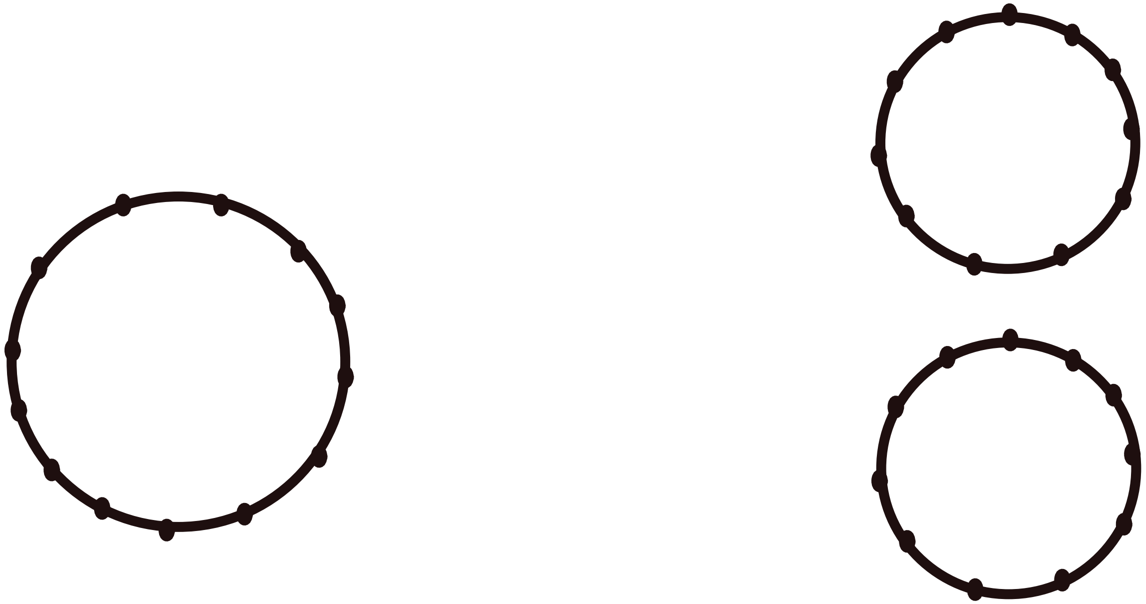

Planarity

A figure illustrating that *planarity* is not first-order definable.

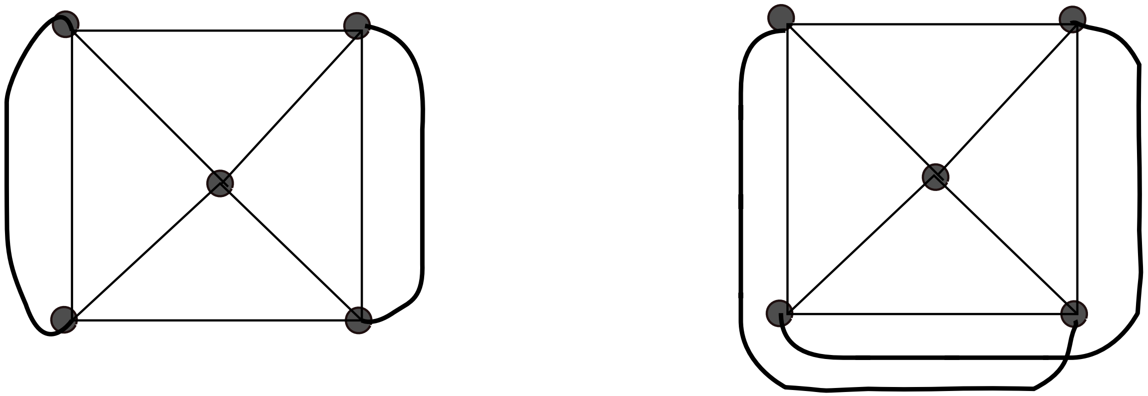

Exercise: use locality to show that 2-colourability and 3-colourability are not first-order definable.

Proof of Gaifman's Theorem

Write $G \sim_q^r H$ to indicate that G and H agree on all *basic, local* sentences with *radius* r and *quantifier rank* q.

It suffices to show that there are functions r and q such that

$$G \sim_{q(p)}^{r(p)} H \quad \text{implies} \quad G \equiv_p H$$

$r = 7^p$ suffices. The value of q will emerge from the proof.

Proof of Gaifman's Theorem

The aim is to prove that if $G \sim_q^r H$, then in the p-round Ehrenfeucht game on G and H, *Duplicator* can inductively maintain the following condition with $m = p - l$ rounds left to play:

$$\bigcup_{i \leq l} N^G_{r(m)}(a_i) \equiv_{q(m)} \bigcup_{i \leq k} N^H_{r(m)}(b_i)$$

where $N^G_r(a)$ denotes the *subgraph* of G induced by the vertices whose distance from a is at most r.

Proof of Gaifman's Theorem

Suppose *w.l.o.g.* that *Spoiler*, in round $l + 1$ plays on a in G.

We describe the response of *Duplicator*.

We distinguish three cases:

1. for some $i \leq k$, $dist(a, a_i) \leq 2r(m-1)$;
2. for all i, $dist(a, a_i) > 2r(m-1)$ and for some i, $dist(a, a_i) \leq 6r(m-1)$; and
3. for all i, $dist(a, a_i) > 6r(m-1)$.

Proof of Gaifman's Theorem

Case 1:

$q(m)$ is bigger than the quantifier rank of the sentence:

$$\exists x(\delta(x, a_i) \leq 2r(m-1) \wedge \theta(\mathbf{x}, x))$$

where $\theta(\mathbf{x}, x)$ is the formula that characterises $\text{Type}_{q(m-1)}(N, \mathbf{a}a)$ for N the graph

$$N = N^G_{r(m-1)}(a) \cup \bigcup_{i \leq l} N^G_{r(m-1)}(a_i)$$

Proof of Gaifman's Theorem

Case 2:

$q(m)$ is bigger than the quantifier rank of the sentence:

$$\exists x(\bigwedge_i (\delta(x, a_i) > 2r(m-1)) \wedge \bigvee_i (\delta(x, a_i) \leq 6r(m-1)) \wedge \theta(x))$$

where $\theta(x)$ is the formula that characterises $\text{Type}_{q(m-1)}(N, a)$ where

$$N = N^G_{r(m-1)}(a)$$

Proof of Gaifman's Theorem

Case 3

Let s be maximal such that $\bigcup_{i \leq l} N^G_{2r(m-1)}(a_i)$ contains s elements, pairwise distance $4r(m-1)$ apart, each satisfying $\theta(x)$.

Note $s \leq l$.

$q(m)$ is big enough so that the value of s is the same in $\bigcup_{i \leq l} N^H_{2r(m-1)}(b_i)$

Since $dist(a, a_i) > 6r(m-1)$ for all i,

$$G \models \exists x_1 \cdots \exists x_{s+1} \left(\bigwedge_{i \neq j} \delta(x_i, x_j) > 4r(m-1) \wedge \bigwedge_i \theta(x_i) \right)$$

Using Gaifman's Theorem

We now want to use Gaifman's theorem to establish Seese's theorem:

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 $G \models \varphi$.

By Gaifman's theorem, it suffices to prove the above for *basic local sentences*.

Seese's Theorem

How do we evaluate a basic local sentence
$\exists x_1 \cdots \exists x_s \left(\bigwedge_{i \neq j} \delta(x_i, x_j) > 2r \wedge \bigwedge_i \psi^r(x_i) \right)$ in a graph $G \in \mathcal{D}_k$?

For each $v \in G$, determine whether

$$N_r(a) \models \psi[a].$$

Since the size of $N_r(a)$ is bounded, this takes linear time.

Label a red if so. We now want to know whether there exists a $2r$-*scattered* set of red vertices of size s.

Finding a Scattered Set

(Frick and Grohe) describe a method to do this efficiently.

Choose red vertices from G in some order, removing the $2r$-neighbourhood of each chosen vertex.

$a_1 \in G$,
$a_2 \in G \setminus N_{2r}(a_1)$,
$a_3 \in G \setminus (N_{2r}(a_1) \cup N_{2r}(a_2)), \ldots$

If the process continues for s steps, we have found a $2r$-scattered set of size s.

Otherwise, for some $u < s$ we have found $a_1, \ldots, a_u$ such that all red vertices are contained in

$$N_{2r}(a_1, \ldots, a_u)$$

This is a graph of bounded size and the property of containing a $2r$-scattered set of *red* vertices of size s can be stated in FO.

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.

Local Tree-Width

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 $\mathcal{C} \subseteq$ 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 $t(r) = 3r$.

Bounded Local Tree-Width

Theorem (Frick-Grohe)

For any class $\mathcal{C}$ of bounded local tree-width and any $\varphi \in \mathsf{FO}$, there is a *quadratic* time algorithm that decides, given $\mathbb{A} \in \mathcal{C}$, whether $\mathbb{A} \models \varphi$.

The proof is a direct application of the method of locality.