The complexity of counting homomorphisms seen from the other side

Víctor Dalmau
Departament de tecnologia
Universitat Pompeu Fabra
Estació de França
Passeig de la circumval.lació 8
08003, Barcelona, Spain
email: victor.dalmau@upf.edu

Peter Jonsson
Department of Computer and Information Science
Linköpings Universitet
SE-581 83 Linköping, Sweden
email: peter.jonsson@ida.liu.se

April 27, 2005

Abstract

For every class of relational structures $\mathcal{C}$, let HOM($\mathcal{C}$, _) be the problem of deciding whether a structure $\mathbf{A} \in \mathcal{C}$ has a homomorphism to a given arbitrary structure $\mathbf{B}$. Grohe has proved that, under a certain complexity-theoretic assumption, HOM($\mathcal{C}$, _) is solvable in polynomial time if and only if the cores of all structures in $\mathcal{C}$ have bounded tree-width. We prove (under a weaker complexity-theoretic assumption) that the corresponding counting problem #HOM($\mathcal{C}$, _) is solvable in polynomial time if and only if all structures in $\mathcal{C}$ have bounded tree-width. This answers an open question posed by Grohe.

Keywords: Computational complexity, counting, homomorphism, relational structure

1 Introduction

The homomorphism problem for relational structures is a fundamental algorithmic problem playing an important role in different areas of computer science including database theory and constraint satisfaction [7, 12, 13] For every two classes of relational structures $\mathcal{C}$,$\mathcal{D}$, let $\mathrm{HOM}(\mathcal{C},\mathcal{D})$ be the problem of deciding whether a structure $\mathbf{A} \in \mathcal{C}$ has a homomorphism to a given arbitrary structure $\mathbf{B} \in \mathcal{D}$. To simplify the notation, if either $\mathcal{C}$ or $\mathcal{D}$ is the class of all structures, we just use the placeholder '_'. The complexity of such problems has received a lot of attention – in particular, the problem when $\mathcal{C}$ is the class of all structures has been studied in depth. One of the most intensively studied problems (known as the *dichotomy problem*) is the following: is it true that the problem $\mathrm{HOM}(\_,\{\mathbf{B}\})$ is either polynomial or NP-complete for arbitrary relational structures $\mathbf{B}$? This question was recently settled for 3-element structures by Bulatov [3].

Grohe [10] has studied the opposite question: how to restrict $\mathcal{C}$ so that $\mathrm{HOM}(\mathcal{C},\_)$ is polynomial-time solvable? He proves the following:

Assume that FPT $\neq$ W[1]. Then for every recursively enumerable class $\mathcal{C}$ of structures of bounded arity, $\mathrm{HOM}(\mathcal{C},\_)$ *is polynomial-time solvable if and only if the core of every structure in $\mathcal{C}$ has tree-width at most w (for some fixed w).*

Here, FPT and W[1] are two parameterized complexity classes that are widely believed to be non-equal, and the *core* of a relational structure $\mathbf{A}$ is a substructure $\mathbf{A}' \subseteq \mathbf{A}$ such that there is a homomorphism from $\mathbf{A}$ to $\mathbf{A}'$ but to no proper substructure of $\mathbf{A}'$. Grohe's proof is based on the Excluded Grid Theorem [15] and it refines some of the ideas originally presented in Grohe et al. [11].

Let $\#\mathrm{HOM}(\mathcal{C},\mathcal{D})$ denote the problem of counting the number of homomorphisms from structures in $\mathcal{C}$ to structures in $\mathcal{D}$. The complexity of $\#\mathrm{HOM}(\_,\{\mathbf{B}\})$ has been thoroughly studied, cf. [4, 5, 6]. In comparison, much less is known about the complexity of $\#\mathrm{HOM}(\mathcal{C},\_)$. In [8], it is shown that if $\mathcal{C}$ is a class of structures of bounded tree-width, then $\#\mathrm{HOM}(\mathcal{C},\_)$ is solvable in polynomial time. One of the questions that was posed during the conference presentation of [10] was the following: is there a characterisation of those structures that allow for efficient counting of the number of homomorphisms, i.e. when is $\#\mathrm{HOM}(\mathcal{C},\_)$ solvable in polynomial time? In this article, we prove the following:

Assume that FPT $\neq$ #W[1]. Then for every recursively enumerable class $\mathcal{C}$ of structures of bounded arity, $\#\mathrm{HOM}(\mathcal{C}, \_)$ is polynomial-time solvable if and only if every structure in $\mathcal{C}$ has tree-width at most w (for some fixed w).

#W[1] denotes the counting version of the complexity class W[1], and the classes FPT and #W[1] are conjectured to be non-equal [8] (in fact, if FPT = #W[1], then FPT = W[1].)

This article has the following structure: Section 2 contains a number of definitions and results on relational structures, graph theory and parameterized complexity. The proof of the main result can be found in Section 3 and Section 4 contains some conclusions and open questions.

2 Preliminaries

Most of the terminology presented in this section comes from [8] and [10]. In the next three subsections, we provide the necessary background material on relational structures, graph theory and parameterized complexity, respectively.

2.1 Relational structures

A *vocabulary* τ is a finite set of relation symbols of specified *arities*, denoted $ar(\cdot)$. The arity of τ is $\max\{ar(R) \mid R \in \tau\}$. A *$\tau$-structure* $\mathbf{A}$ consists of a finite set A (called the *universe* of $\mathbf{A}$) and for each relation symbol $R \in \tau$, a relation $R^A \subseteq A^{ar(R)}$. Throughout the article, we will use the same boldface and slanted capital letters to denote a structure and its universe, respectively.

We say that a class $\mathcal{C}$ of structures is of *bounded arity* if there is an r such that every structure in $\mathcal{C}$ is at most r-ary. A *substructure* of a τ-structure $\mathbf{A}$ is a τ-structure $\mathbf{B}$ with universe $B \subseteq A$ and relations $R^B \subseteq R^A$ for all $R \in \tau$. A substructure $\mathbf{B}$ is *induced* if for all $R \in \tau$, say, of arity r, $R^{\mathbf{B}} = R^{\mathbf{A}} \cap B^r$. We define the *size* $\|\mathbf{A}\|$ of the structure $\mathbf{A}$ as

$$\|A\| = |\tau| + |A| + \sum_{R \in \tau} |R^{\mathbf{A}}| \cdot |ar(R)|.$$

$\|\mathbf{A}\|$ is roughly the size of a reasonable encoding of $\mathbf{A}$.

Let $\mathbf{A}$ and $\mathbf{B}$ be τ-structures. We define $\mathbf{A} \cup \mathbf{B}$ to be the τ-structure with universe $A \cup B$ and such that for all $R \in \tau$, $R^{\mathbf{A} \cup \mathbf{B}} = R^{\mathbf{A}} \cup R^{\mathbf{B}}$

A *homomorphism* from an τ-structure $\mathbf{A}$ to a τ-structure $\mathbf{B}$ is a mapping $h : A \to B$ such that for all $R \in \tau$, say, of arity r, and all tuples $(a_1, \ldots, a_r) \in$

R^A, we have $h((a_1, \dots, a_r)) \in R^B$. For two classes $\mathcal{C}$ and $\mathcal{D}$ of structures, $\#\text{HOM}(\mathcal{C}, \mathcal{D})$ is the following problem:
INSTANCE: $\mathbf{A} \in \mathcal{C}$, $\mathbf{B} \in \mathcal{D}$.
OUTPUT: The number of homomorphism from $\mathbf{A}$ to $\mathbf{B}$.
and $\text{HOM}(\mathcal{C}, \mathcal{D})$ is the problem of deciding whether there exists a homomorphism or not. If $\mathcal{D}$ is the class of all finite structures, we write $\#\text{HOM}(\mathcal{C}, \_)$ instead of $\#\text{HOM}(\mathcal{C}, \mathcal{D})$.

2.2 Graph theory

Let E be a binary relation symbol. We view graphs as $\{E\}$-structures $\mathbf{G} = (G, E^{\mathbf{G}})$ and assume that they are undirected and loop-free. A graph $\mathbf{H} = (H, E^{\mathbf{H}})$ is a *minor* of a graph $\mathbf{G}$ if $\mathbf{H}$ is isomorphic to a graph that can be obtained from a subgraph of $\mathbf{G}$ by contracting edges. We define a *minor map* from $\mathbf{H}$ to $\mathbf{G}$ to be a mapping $M : H \to 2^G$ having the following properties:

1. for all $v \in H$, the set $M(v)$ is non-empty and connected in $\mathbf{G}$;
2. for all $v, w \in H$ with $v \neq w$, the sets $M(v)$ and $M(w)$ are disjoint; and
3. for all edges $\{v, w\} \in E^{\mathbf{H}}$, there are $v' \in M(v)$ and $w' \in M(w')$ such that $\{v', w'\} \in E^{\mathbf{G}}$.

We would like to emphasize that a minor map maps nodes in $\mathbf{H}$ to sets of nodes of $\mathbf{G}$. We call a minor map M from $\mathbf{H}$ to $\mathbf{G}$ *onto* if $\bigcup_{v \in H} M(v) = G$. It is easy to see that there is a minor map from $\mathbf{H}$ to $\mathbf{G}$ if and only if $\mathbf{H}$ is a minor of $\mathbf{G}$. Moreover, if $\mathbf{H}$ is a minor of a connected graph $\mathbf{G}$, then we can always find a minor map from $\mathbf{H}$ onto $\mathbf{G}$.

A *tree-decomposition* of a graph $\mathbf{G}$ is a pair $(\mathbf{T}, \beta)$ where $\mathbf{T}$ is a tree and $\beta : T \to 2^G$ satisfies the following conditions:

1. for every $v \in G$, the set $\{t \in T \mid v \in \beta(t)\}$ is non-empty and connected in $\mathbf{T}$; and
2. for every $e \in E^{\mathbf{G}}$, there is a $t \in T$ such that $e \subseteq \beta(t)$.

The width of a tree-decomposition $(\mathbf{T}, \beta)$ is $\max\{|\beta(t)| \mid t \in T\} - 1$, and the *tree-width* $\omega(\mathbf{G})$ of a graph $\mathbf{G}$ is the minimum w such that $\mathbf{G}$ has a tree-decomposition of width w.

For $k, \ell \geq 1$, the $(k \times \ell)$-grid is the graph with vertex set $\{1, \dots, k\} \times \{1, \dots, \ell\}$ and an edge between (i, j) and (i', j') if and only if $|i - i'| + |j - j'| = 1$. It is known that the $(k \times k)$-grid has tree-width k, cf. Bodlaender [2, Sec.

13]. Robertson and Seymour have proved the following theorem which is known as the Excluded Grid Theorem.

Theorem 2.1 *[15] For every k there exists a $w(k)$ such that the $(k \times k)$-grid is a minor of every graph of tree-width at least $w(k)$.*

We will now generalise some of the graph-theoretic notions defined above to arbitrary relational structures. The *Gaifman graph* of a τ-structure $\mathbf{A}$ is the graph $\mathbf{G}(\mathbf{A})$ with vertex set A and an edge between a and b if $a \neq b$ and there is a relation symbol $R \in \tau$, say, of arity r, and a tuple $(a_1, \ldots, a_r) \in R^A$ such that $a, b \in \{a_1, \ldots, a_r\}$. Henceforth, we say that a subset $B \subseteq A$ is connected in a structure $\mathbf{A}$ if it is connected in $\mathbf{G}(\mathbf{A})$. A tree-decomposition of a τ-structure $\mathbf{A}$ is viewed as a tree-decomposition of $\mathbf{G}(\mathbf{A})$. A minor map from $\mathbf{A}$ to $\mathbf{B}$ is a mapping $M : A \to B$ that is a minor map from $\mathbf{G}(\mathbf{A})$ to $\mathbf{G}(\mathbf{B})$.

2.3 Parameterized complexity

Finally, we need some facts concerning parameterized complexity theory. Parameterized counting problems have been studied independently in a number of papers (cf.[1, 8, 14]) and our notation follows that of [8]. A *parameterized counting* problem is a function $F : \Sigma^* \times \mathbb{N} \to \mathbb{N}$, where Σ is a finite alphabet. If $(x, k) \in \Sigma^* \times \mathbb{N}$ is an instance of a parameterized counting problem, we call x the *input* and k the *parameter*. For example, the parameterized clique counting problem p-#CLIQUE, is the following problem:

INPUT: graph $\mathbf{G}$.
PARAMETER: $k \in \mathbb{N}$.
OUTPUT: The number of cliques in $\mathbf{G}$ of size k.

The problem we are mainly interested in is the counting homomorphism problem parameterized by the size of the left hand side input structure (p-#HOM($\mathcal{C}, \mathcal{D}$)):

INSTANCE: $\mathbf{A} \in \mathcal{C}$, $\mathbf{B} \in \mathcal{D}$.
PARAMETER: $||\mathbf{A}||$
OUTPUT: The number of homomorphisms from $\mathbf{A}$ to $\mathbf{B}$.

A parameterized counting problem $F : \Sigma^* \times \mathbb{N} \to \mathbb{N}$ is *fixed-parameter tractable* if there is a computable function $f : \mathbb{N} \to \mathbb{N}$, a constant $c \in \mathbb{N}$

and an algorithm that given $(x, k) \in \Sigma^* \times \mathbb{N}$ computes $F(x, k)$ in time $f(k) \cdot |x|^c$. FPT denotes the class of all fixed-parameter tractable counting problems. We note that all fixed-parameter tractable *decision* problems are also included in FPT.

Let $F : \Sigma^* \times \mathbb{N} \to \mathbb{N}$ and $G : \Sigma^* \times \mathbb{N} \to \mathbb{N}$ be parameterized counting problems. A *parameterized parsimonious reduction* from F to G is an algorithm that computes for every instance (x, k) of F an instance (y, ℓ) of G in time $f(k) \cdot |x|^c$ such that $\ell \leq g(k)$ and $F(x, k) = G(y, \ell)$ (for computable functions $f, g : \mathbb{N} \to \mathbb{N}$ and a constant $c \in \mathbb{N}$). *Hardness* and *completeness* are defined in the usual way. Flum & Grohe [8] defined a hierarchy $\#W[1] \subseteq \#W[2] \subseteq \ldots$ of parameterized counting complexity classes, and they conjecture that this hierarchy is strict and that FPT is strictly contained in #W[1]. They have also proved that p-#CLIQUE is #W[1]-complete under parameterized parsimonious reductions.

3 Main result

We are now ready to prove the main result. Subsection 3.1 contains a construction of certain structures that the proof relies on, and Subsection 3.2 contains the actual proof. The construction in Subsection 3.1 shows that given an arbitrary connected τ-structure $\mathbf{A}$ such that there exists a minor map from the $(k \times K)$-grid onto $\mathbf{A}$, one can compute a structure $\mathbf{B}$ which allows only a certain number of restricted (in a sense to be defined later) homomorphisms from $\mathbf{A}$ to $\mathbf{B}$. The proof presented in Subsection 3.2 is an fpt-reduction from p-#CLIQUE to p-#HOM$(\mathcal{C}, _)$ that uses $\mathbf{B}$ as one of its building blocks.

3.1 The structure B

Let $\mathbf{A}$ be a connected τ-structure. Let $k \geq 2$, $K = \binom{k}{2}$, and $M : \{1, \ldots, k\} \times \{1, \ldots, K\} \to 2^A$ a minor map from the $(k \times K)$-grid onto $\mathbf{A}$. Let us assume that we have fixed some bijection ϱ between $\{1, \ldots, K\}$ and the set of all unordered pairs of elements of $\{1, \ldots, k\}$. For improved readability, we sloppily write $i \in p$ instead of $i \in \varrho(p)$.

Let the $\{E\}$-structure $\mathbf{G}$ be a graph. We present the τ-structure $\mathbf{B} = \mathbf{B}(\mathbf{A}, M, \mathbf{G})$ as defined by Grohe [10]; he shows that there is a homomorphism from $\mathbf{A}$ to $\mathbf{B}$ if and only if $\mathbf{G}$ contains a k-clique. Let $\mathbf{id}$ denote the identity function on the set A. We show that the number of homomorphisms h from $\mathbf{A}$ to $\mathbf{B}$ such that $\Pi \circ h = \mathbf{id}$ equals the number of k-cliques in $\mathbf{G}$

multiplied by $k!$ (where Π will be defined below). The universe B of $\mathbf{B}$ is given by:

$$\begin{aligned}\{(v,e,i,p,a)| \quad & v \in G, e \in E^{\mathbf{G}}, \\ & 1 \leq i \leq k, 1 \leq p \leq K \text{ s.t. } (v \in e \Longleftrightarrow i \in p), \\ & a \in M(i,p)\}\end{aligned}$$

We define the function $\Pi : B \to A$ by letting

$$\Pi(v,e,i,p,a) = a$$

As usual, we extend Π and Π^{-1} to tuples by defining it componentwise.

For every relation $R \in \tau$, say, of arity r, and for all tuples $(a_1, \ldots, a_r) \in R^A$, we add to R^B all tuples $(b_1, \ldots, b_r) \in \Pi^{-1}(a_1, \ldots, a_r)$ satisfying the following two constraints for all $b, b' \in \{b_1, \ldots, b_r\}$:

(C1) if $b = (v,e,i,p,a)$ and $b' = (v',e',i,p',a')$, then $v = v'$; and

(C2) if $b = (v,e,i,p,a)$ and $b' = (v',e',i',p,a')$, then $e = e'$.

It follows immediately from the definition of $\mathbf{B}$ that Π defines an homomorphism from $\mathbf{B}$ to $\mathbf{A}$. We now prove a lemma concerning the number of homomorphisms from $\mathbf{A}$ to $\mathbf{B}$ by refining Grohe's [10] original proof.

Lemma 3.1 *The number of homomorphisms h from $\mathbf{A}$ to $\mathbf{B}$ such that $\Pi \circ h = \mathbf{id}$ equals the number of k-cliques in $\mathbf{G}$ multiplied by $k!$.*

Proof: Let $v_1, \ldots, v_k$ be a vertex set of a k-clique in G. Recall that ϱ is a bijection between $\{1, \ldots, K\}$ and the set of unordered pairs of elements of $\{1, \ldots, k\}$. We define the mapping $h_{v_1,\ldots,v_k} : A \to B$ from the universe A of $\mathbf{A}$ to the universe B of $\mathbf{B}$ as

$$h_{v_1,\ldots,v_k}(a) = (v_i, \{v_j, v_l\}, i, p, a)$$

for $1 \leq i \leq k$, $1 \leq p \leq K$, $a \in M(i,p)$ and $\{j,l\} = \varrho(p)$. It is immediate that $h_{v_1,\ldots,v_k}$ defines an homomorphism from $\mathbf{A}$ to $\mathbf{B}$. Furthermore, $\Pi \circ h_{v_1,\ldots,v_k} = \mathbf{id}$.

We shall now show that if h is an homomorphism from $\mathbf{A}$ to $\mathbf{B}$ such that $\Pi \circ h = \mathbf{id}$ then $h = h_{v_1,\ldots,v_k}$ for some vertex set $v_1, \ldots, v_k$ of a k-clique in G. Let h be an homomorphism from $\mathbf{A}$ to $\mathbf{B}$ with $\Pi \circ h = \mathbf{id}$. This means that for $1 \leq i \leq k$, $1 \leq p \leq K$, and $a \in M(i,p)$ we have

$$h(a) = (v_a, e_a, i, p, a)$$

for a vertex $v_a \in G$ and an edge $e_a \in E^{\mathbf{G}}$ such that $(v_a \in e_a \Longleftrightarrow i \in p)$.

Claim 1: For $1 \leq i \leq k$, $1 \leq p \leq K$, and $a, a' \in M(i,p)$ we have $v_a = v_{a'}$ and $e_a = e_{a'}$.

Proof: Since $M(i,p)$ is connected in $\mathbf{A}$, it suffices to prove the claim for a, a' such that there is an edge between a and a' in the Gaifman graph of $\mathbf{A}$. So, let $R \in \tau$ and $\overline{a} = (a_1, \ldots, a_{ar(R)}) \in R^{\mathbf{A}}$ such that $a, a' \in \{a_1, \ldots, a_{ar(R)}\}$. Since h is an homomorphism we have $h(\overline{a}) \in R^{\mathbf{B}}$. Thus by conditions (C1) and (C2), we must have $v_a = v_{a'}$ and $e_a = e_{a'}$. This proves Claim 1.

Claim 2: For $1 \leq i, i' \leq k$, $1 \leq p \leq K$, and $a \in M(i,p)$, $a' \in M(i',p)$ we have $e_a = e_{a'}$.

Proof: By a simple inductive argument in which Claim 1 is the case base, it suffices to prove Claim 2 for $i' = i + 1$.

Since M is a minor map from the $(k \times K)$-grid to $\mathbf{A}$ and there is an edge between (i,p) and (i',p) in the grid, there must some edge (a^1, a^2) in the Gaifman graph $\mathbf{G}(\mathbf{A})$ where $a^1 \in M(i,p)$ and $a^2 \in M(i',p)$. Thus there exists a relation symbol $R \in \tau$ and a tuple $\overline{a} \in R^{\mathbf{A}}$ such that both a^1 and a^2 occur in $\overline{a}$. Since h is an homomorphism we have that $h(\overline{a}) \in R^{\mathbf{B}}$. Thus, by condition (C2) we have $e_{a^1} = e_{a^2}$ and by Claim 1, we have $e_a = e_{a^1}$ and $e_{a'} = e_{a^2}$. This completes the proof of Claim 2.

Claim 3: For $1 \leq i \leq k$, $1 \leq p, p' \leq K$, and $a \in M(i,p)$, $a' \in M(i,p')$ we have $v_a = v_{a'}$.

Proof: Analogously to the proof of Claim 2 using condition (C1) instead of (C2)

Together the three claims imply that there are vertices $v_1, \ldots, v_k$ in G and edges $e_1, \ldots, e_K \in E^{\mathbf{G}}$ such that for all $1 \leq i \leq k$, $1 \leq p \leq K$ and $a \in M(i,p)$ we have $h(a) = (v_i, e_p, i, p, a)$. Furthermore, since $h(a) \in B$, we have that $e_p = \{v_j, v_l\}$ where $\{j, l\} = \varrho(p)$.

We have just proved that if h is an homomorphism from $\mathbf{A}$ to $\mathbf{B}$ such that $\Pi \circ h = \mathbf{id}$ then h must be of the form $h_{v_1,\ldots,v_k}$ for some vertex set $v_1, \ldots, v_k$ of a k-clique in $\mathbf{G}$. Since for every vertex set $v_1, \ldots, v_k$ of a k-clique in $\mathbf{G}$ we have that $h_{v_1,\ldots,v_k}$ defines an homomorphism we have a complete description of the homomorphisms from $\mathbf{A}$ to $\mathbf{B}$ satisfying $\Pi \circ h = \mathbf{id}$: a mapping h from the universe A of $\mathbf{A}$ to the universe B of $\mathbf{B}$ is an homomorphism satisfying $\Pi \circ h = \mathbf{id}$ if and only if $h = h_{v_1,\ldots,v_k}$ for some vertex set $v_1, \ldots, v_k$ of a k-clique in $\mathbf{G}$. Consequently the number of such mappings equals the number of k-cliques in $\mathbf{G}$ multiplied by $k!$. ∎

3.2 The proof

This subsection contains the proof of the main result. The proof is basically an fpt-reduction from p-#CLIQUE to p-#HOM($\mathcal{C}$, _). The reduction exploits the fact that if $\mathcal{C}$ do not have bounded tree-width, then it contains some structure $\mathbf{A}$ such that the $(k \times K)$-grid is a minor of the Gaifman graph of $\mathbf{A}$ (by the Excluded Grid Theorem). This makes the results in the previous subsection applicable and we show that the number of cliques can be efficiently computed by using the inclusion-exclusion principle.

Theorem 3.2 *Assume that FPT $\neq$ #W[1]. Then for every recursively enumerable class $\mathcal{C}$ of structures of bounded arity the following three statements are equivalent:*

1. #HOM($\mathcal{C}$, _) *is in FP.*
2. p-#HOM($\mathcal{C}$, _) *is in FPT.*
3. *There is a $\omega \geq 1$ such that every structure in $\mathcal{C}$ has tree-width at most ω.*

Proof: The implication (1)⇒(2) is trivial. Implication (3)⇒(1) is Proposition 7 of [8]. It remains to prove that (2)⇒(3). We shall show that if (2) is true and (3) false then p-#CLIQUE is in FPT, in contradiction with the fact that it is #W[1]-complete. We prove this by exhibiting an fpt-reduction from p-#CLIQUE to p-#HOM($\mathcal{C}$, _).

Let $(\mathbf{G}, k)$ be an instance of p-#CLIQUE. By the Excluded Grid Theorem, there is some structure $\mathbf{A}$ in $\mathcal{C}$ such that the $(k \times K)$-grid is a minor of the Gaifman graph of $\mathbf{A}$. We enumerate the recursively enumerable class $\mathcal{C}$ until we find such an $\mathbf{A} = \mathbf{A}(k)$. Then we compute a minor map M from the $(k \times K)$-grid to $\mathbf{A}$ and we compute the number of automorphisms I of $\mathbf{A}$. Let $\mathbf{A_1}, \ldots, \mathbf{A_m}$ be a decomposition of $\mathbf{A}$ into its connected components. We can assume wlog. that the $(k \times K)$-grid is a minor of (the Gaifman graph of) $\mathbf{A_1}$ and that the minor map M is onto $\mathbf{A_1}$.

Let $\mathbf{B} = \mathbf{B}(\mathbf{G}, M, \mathbf{A})$ constructed as above. In order to compute the number of k-cliques of $\mathbf{G}$ we only need to compute the number N of homomorphisms h from $\mathbf{A_1}$ to $\mathbf{B}$ such that $\Pi \circ h = \mathbf{id}$, and divide this number by

$k!$ by Lemma 3.1. To do this construct $\mathbf{C}$ as $\mathbf{B} \cup \mathbf{A_2} \cup \ldots \cup \mathbf{A_m}$. Then the function $\Gamma : C \to A$ defined by

$$\Gamma(c) = \begin{cases} \Pi(c) & \text{if } c \in B \\ c & \text{otherwise} \end{cases}$$

defines an homomorphism from $\mathbf{C}$ to $\mathbf{A}$. Furthermore the number N of homomorphisms h from $\mathbf{A_1}$ to $\mathbf{B}$ such that $\Pi \circ h = \mathbf{id}$ is identical to the number of homomorphisms h from $\mathbf{A}$ to $\mathbf{C}$ such that $\Gamma \circ h = \mathbf{id}$. Furthermore, in order to compute N we only need to compute the number N' of homomorphisms h from $\mathbf{A}$ to $\mathbf{C}$ such that $\Gamma \circ h(A) = A$ and divide it by I.

This fact requires a bit of proof: Let $\mathcal{N}$ be the set of all homomorphisms h from $\mathbf{A}$ to $\mathbf{C}$ such that $\Gamma \circ h = \mathbf{id}$, let $\mathcal{N}'$ be the set of all homomorphisms h from $\mathbf{A}$ to $\mathbf{C}$ such that $\Gamma \circ h(A) = A$, and let $\mathcal{I}$ be the set of all isomorphisms of $\mathbf{A}$. We shall show that

$$\mathcal{N}' = \{f \circ g | f \in \mathcal{N}, g \in \mathcal{I}\}$$

The $\supseteq$ part is trivial. For the $\subseteq$ part, let h be any function in $\mathcal{N}'$. Consequently, $g = \Gamma \circ h$ is an isomorphism of $\mathbf{A}$. Henceforth, g^{-1} is also an isomorphism and $h \circ g^{-1}$ is also an homomorphism from $\mathbf{A}$ to $\mathbf{C}$. Notice that $\Gamma \circ h \circ g^{-1}(a) = a$ for all $a \in A$ and hence $h \circ g^{-1}$ belongs to $\mathcal{N}$. Furthermore, $h = h \circ g^{-1} \circ g$ and we are done. Finally, it is easy to see that for every $f, f' \in \mathcal{N}$, and every $g, g' \in \mathcal{I}$, if $f \neq f'$ or $g \neq g'$, then $f \circ g \neq f' \circ g'$, and hence $|\mathcal{N}'| = |\mathcal{N}| \cdot |\mathcal{I}|$.

It only remains to compute the value of N'. Let S be any subset of A. We denote by M_S the number of homomophisms h from $\mathbf{A}$ to $\mathbf{C}$ such that $\Gamma \circ h(A) \subseteq S$. Clearly, M_S is equal to the number of homomorphisms from $\mathbf{A}$ to $\mathbf{C}_{|\Gamma^{-1}(S)}$ where $\mathbf{C}_{|\Gamma^{-1}(S)}$ is the induced substructure of $\mathbf{C}$ with universe $\Gamma^{-1}(S)$. Consequently, for every S, M_S is computable in time $f_1(k) \cdot \|\mathbf{G}\|^{c_1}$ for some $f_1 : N \to N$ and some constant $c_1 \in N$. Finally, the value of N' can be computed by the inclusion-exclusion principle as

$$N' = M_{\{a_1,\ldots,a_m\}} - M_{\{a_1,\ldots,a_{m-1}\}} - M_{\{a_1,\ldots,a_{m-1},a_m\}} \cdots - M_{\{a_2,\ldots,a_{m-1}\}} + M_{\{a_1,\ldots,a_{m-2}\}} + \cdots$$

where $A = \{a_1, \ldots, a_m\}$. Since the number of subsets of A only depends on k, we can conclude that there exists some $f_2 : N \to N$ and some constant $c_2 \in N$ such that the value of $|\mathcal{M}|$ can be computed in time $f_2(k) \cdot \|\mathbf{G}\|^{c_2}$. ∎

4 Conclusions

We have proved that $\#\mathrm{HOM}(\mathcal{C}, \_)$ is solvable in polynomial time if and only if all structures in $\mathcal{C}$ have bounded tree-width, under the following three assumptions:

1. FPT $\neq$ #W[1];
2. $\mathcal{C}$ is of bounded arity; and
3. $\mathcal{C}$ is recursively enumerable.

An interesting question is whether all of these assumptions are necessary for the result to hold or not. First we note that lifting the assumption FPT $\neq$ #W[1] is probably very hard since this would separate P from #P. Lifting the assumption that $\mathcal{C}$ is of bounded arity is not possible since there exist $\mathcal{C}$ of unbounded arity and unbounded tree-width such that $\#\mathrm{HOM}(\mathcal{C}, \_)$ is polynomial: for instance, if $\mathcal{C}$ only contains structures such that their underlying hypergraphs are hypertrees, then $\#\mathrm{HOM}(\mathcal{C}, \_)$ can be solved in polynomial time by a straightforward dynamic programming algorithm. Many other examples of $\mathcal{C}$ of unbounded arity such that $\mathrm{HOM}(\mathcal{C}, \_)$ can be solved in polynomial-time are known, cf. Gottlob et al. [9]. We conjecture that $\#\mathrm{HOM}(\mathcal{C}, \_)$ is also polynomial-time solvable for many of these classes. Proving an analogue of Theorem 3.2 for structures of unbounded arity thus seems to be a quite challenging problem. We remark that the complexity of the decision problem $\mathrm{HOM}(\mathcal{C}, \_)$ when $\mathcal{C}$ is allowed to be of unbounded arity is still unknown.

Finally, the set of structures $\mathcal{C}$ do not have to be recursively enumerable if we are willing to use the following stronger complexity-theoretic assumption: non-uniform-FPT $\neq$ non-uniform-#W[1]. Here, a parameterized counting problem $F : \Sigma^* \times \mathbb{N} \rightarrow \mathbb{N}$ is *non-uniformly fixed-parameter tractable* (and a member of the class non-uniform-FPT) if there is a constant $c \in \mathbb{N}$ such that for every $k \in \mathbb{N}$ there is an algorithm that given $x \in \Sigma^*$, decides if $(x, k) \in P$ in time $O(|x|^c)$. Similarly, one can define the notion of non-uniform parsimonious fpt-reductions and define non-uniform-#W[1] as the class of all parameterized counting problems that can be reduced to p-#CLIQUE by non-uniform parsimonious fpt-reductions. Under this stronger assumption, the recursive enumerability of the class $\mathcal{C}$ can be lifted since it is only used to guarantee that the reduction from p-#CLIQUE to p-$\#\mathrm{HOM}(\mathcal{C}, \_)$ in the proof of Theorem 3.2 is a proper fpt-reduction as defined in Section 2.3.

Acknowledgements

Víctor Dalmau is supported by the *Spanish MCyT* under grants TIC2002-04019-C03 and TIC2002-04470-C03. Peter Jonsson is partially supported by the *Center for Industrial Information Technology* (CENIIT) under grant 04.01, and by the *Swedish Research Council* (VR) under grants 221-2000-361 and 621-2003-3421.

References

[1] V. Arvind and V. Raman, Approximation algorithms for some parameterized counting problems, in: Proceedings of the 13th Annual International Symposium on Algorithms and Computation (2002) 453–464.

[2] H. L. Bodlaender, A partial k-arboretum of graphs with bounded treewidth, Theoretical Computer Science 209(1–2) (1998) 1–45.

[3] A. Bulatov, A dichotomy theorem for constraints on a three-element set, in: Proceedings of the 43rd IEEE Symposium on Foundations of Computer Science (2002) 649–658.

[4] A. A. Bulatov and V. Dalmau, Towards a dichotomy theorem for the counting constraint satisfaction problem, in: Proceedings of the 44th IEEE Symposium of Foundations of Computer Science (2003) 562–573.

[5] N. Creignou and M. Hermann, Complexity of generalized satisfiability counting problems, Information and Computation 125 (1996) 1–12.

[6] M. E. Dyer and C. S. Greenhill, The complexity of counting graph homomorphisms, Random Structures and Algorithms, 17 (2000) 260–289.

[7] T. Feder and M. Y. Vardi, The computational structure of monotone monadic SNP and constraint satisfaction: a study through datalog and group theory, SIAM Journal on Computing 28(1) (1998) 57–104.

[8] J. Flum and M. Grohe, The parameterized complexity of counting problems, in: Proceedings of the 43rd IEEE Symposium on Foundations of Computer Science (2002) 538–547.

[9] G. Gottlob, N. Leone, and F. Scarcello, A comparison of structural CSP decomposition methods, Artificial Intelligence 124(2) (2000) 243–282.

[10] M. Grohe, The complexity of homomorphism and constraint satisfaction problems seen from the other side, in: Proceedings of the 44th IEEE Symposium on Foundations of Computer Science (2003) 552–561.

[11] M. Grohe, T. Schentwick, and L. Segoufin, When is the evaluation of conjunctive queries tractable?, in: Proceedings of the 33rd ACM Symposium on Theory of Computing (2001) 657–666.

[12] P. Jeavons, On the algebraic structure of combinatorial problems, Theoretical Computer Science 200(1–2) (1998) 185–204.

[13] P. G. Kolaitis and M. Vardi, Conjunctive-query containment and constraint satisfaction, Journal of Computer and System Sciences 61 (2000) 302–332.

[14] C. McCartin, Parameterized counting problems, in: Proceedings of the 27th International Symposium on Mathematical Foundations of Computer Science (2002) 556–567.

[15] N. Robertson and P. Seymour., Graph minors V: Excluding a planar graph, Journal of Combinatorial Theory, ser. B, 62 (1994) 323–348.