Bounded Tree-width and CSP-related Problems

Tommy Färnqvist* and Peter Jonsson**

Department of Computer and Information Science
Linköpings Universitet
S-581 83 Linköping, Sweden
{tomfa, petej}@ida.liu.se

Abstract. We study the complexity of structurally restricted homomorphism and constraint satisfaction problems. For every class of relational structures $\mathcal{C}$, let LHOM$(\mathcal{C}, _)$ be the problem of deciding whether a structure $\mathbf{A} \in \mathcal{C}$ has a homomorphism to a given arbitrary structure $\mathbf{B}$, when each element in $\mathbf{A}$ is only allowed a certain subset of elements of $\mathbf{B}$ as its image. We prove, under a certain complexity-theoretic assumption, that this *list homomorphism problem* is solvable in polynomial time if and only if all structures in $\mathcal{C}$ have bounded tree-width. The result is extended to the connected list homomorphism, edge list homomorphism, minimum cost homomorphism and maximum solution problems. We also show an inapproximability result for the minimum cost homomorphism problem.
Keywords: Computational complexity, constraint satisfaction, homomorphism, relational structure, inapproximability.

1 Introduction

A large class of problems in different areas of computer science can be viewed as constraint satisfaction problems [2, 7, 13, 15, 20, 23]. This includes problems in artificial intelligence, database theory, scheduling, frequency assignment, graph theory and satisfiability. The main model [13] considers constraint satisfaction problems with a fixed template determining the size of the domain and the set of allowed constraint types in an instance. Feder and Vardi [13] observed that constraint satisfaction problems can be described as homomorphism problems for relational structures. For an excellent introduction to and survey of the strongly related subject of *graph* homomorphisms, we refer to [17]. For every two classes of relational structures $\mathcal{C}$, $\mathcal{D}$, let 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 '_'. Grohe [15] has studied so called *structural* restrictions, i.e. the question of how to restrict $\mathcal{C}$, so that 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, 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).*

* Supported by the *National Graduate School in Computer Science* (CUGS), Sweden.

** Partially supported by the *Center for Industrial Information Technology* (CENIIT) under grant 04.01, and by the *Swedish Research Council* (VR) under grant 2006-4532.

FPT $\neq$ W[1] is a standard assumption from parameterised complexity theory that is widely believed to be true. A *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 no homomorphism from $\mathbf{A}'$ to a proper substructure of $\mathbf{A}'$. All cores of a structure $\mathbf{A}$ are isomorphic, so it is reasonable to speak of *the* core of $\mathbf{A}$.

In the *list homomorphism problem* [2, 6, 7, 9–12], $\text{LHOM}(\mathcal{C}, \mathcal{D})$, the goal is to decide whether there is a homomorphism from a structure $\mathbf{A} \in \mathcal{C}$ to a given structure $\mathbf{B} \in \mathcal{D}$, when each element in $\mathbf{A}$ is only allowed a certain subset of elements in the universe of $\mathbf{B}$ as its image. Such list homomorphisms generalise e.g. list colourings and have many natural applications. We show the following:

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

Incidentally, this complexity-theoretic classification coincides with that of Dalmau and Jonsson's in [3], where they study the problem of counting homomorphisms. Our result is then extended to the *connected* list homomorphism problem [6], where every list has to induce a connected substructure of the right hand side input structure and to the *edge* list homomorphism problem [8], where the lists contain tuples from the relations of the right hand side input structure that the tuples on the left hand side have to map to. We remark that our hardness results still apply when the classes of relational structures are restricted to graphs. We also extend the result to two optimisation problems. The *minimum cost* homomorphism problem was introduced by Gutin et al. in [16], where it was motivated by a real-world problem in defence logistics. Here, mapping an element from the left hand side to an element on the right hand side is afflicted with costs and the objective is to find a homomorphism of minimum cost. This problem includes as special cases the list homomorphism problem and the general optimum cost chromatic partition problem [24]. In the *maximum solution problem* [21], the right hand side elements are assumed to be a finite subset of the natural numbers and the objective is to find a homomorphism that has maximum possible total weight. In some sense, see [21], this is a generalisation of integer programming and captures e.g. the INDEPENDENT SET problem. When the right hand side is restricted to $\{0, 1\}$ this is the well-studied MAX ONES problem. The hard instances of the minimum cost homomorphism problem are also shown to be inapproximable as well. To our knowledge, this is the first inapproximability result for this problem.

The rest of this paper is organised as follows. Section 2 introduces the requisite background material and problem definitions for several variants of the homomorphism problem. Section 3 contains proofs of our intractability and inapproximability results. Finally, Section 4 concludes the paper and presents possible future work.

2 Preliminaries

Most of the terminology presented in this section comes from [3, 14–16]. In the next three subsections, we provide the necessary background material on relational structures and graph theory, homomorphism problems and parameterised complexity, respectively.

2.1 Relational Structures and Graph Theory

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 τ*-structure* $\mathbf{A}$ consists of a finite set A (called the *universe* of **A**) and for each relation symbol $R \in \tau$, a relation $R^A \subseteq A^{ar(R)}$. 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 $\|\mathbf{A}\| = |\tau| + |A| + \sum_{R \in \tau} |R^{\mathbf{A}}| \cdot |ar(R)|$. $\|\mathbf{A}\|$ 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 $R \in \tau$, $R^{\mathbf{A} \cup \mathbf{B}} = R^{\mathbf{A}} \cup R^{\mathbf{B}}$

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

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

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

A *tree-decomposition* of a graph **G** is a pair $(\mathbf{T}, \beta)$ where **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 **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 **G** is the minimum w such that **G** has a tree-decomposition of width w.

For $k, \ell \geq 1$, the $(k \times \ell)$-grid is the graph with vertex set $\{1, \ldots, k\} \times \{1, \ldots, \ell\}$ and an edge between (i, j) and (i', j') if and only if $|i - i'| + |j - j'| = 1$. It is not hard to see that the $(k \times k)$-grid has tree-width k. Robertson and Seymour have proved the following theorem which is known as the Excluded Grid Theorem:

Theorem 1. *[25] 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 **A** is the graph **G(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 **A** if it is connected in **G(A)**. A tree-decomposition of a τ-structure **A** is viewed as a tree-decomposition of **G(A)**. A minor map from **A** to **B** is a mapping $\mu : A \to 2^B$ that is a minor map from **G(A)** to **G(B)**.

2.2 Homomorphism Problems

A *homomorphism* from a τ-structure **A** to a τ-structure **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 $(h(a_1, \ldots, h(a_r)) \in R^B$.

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: "yes" if a homomorphism from **A** to **B** exists, "no" if no homomorphism from **A** to **B** exists.

If $\mathcal{D}$ is the class of all finite structures, we write $\text{HOM}(\mathcal{C}, \_)$ instead of $\text{HOM}(\mathcal{C}, \mathcal{D})$.

In the *list homomorphism problem*, each element of the left hand side input structure is given together with a set, called a *list*, of possible images in the right hand side input structure. This problem has been well studied with regard to restrictions to the right hand side input structure, see e.g. [2, 6, 7, 9–12] for some results. We denote it $\text{LHOM}(\mathcal{C}, \mathcal{D})$:

INSTANCE: $\mathbf{A} \in \mathcal{C}$, $\mathbf{B} \in \mathcal{D}$, $L_a \subseteq B$ for each $a \in A$.
OUTPUT: "yes" if a homomorphism h from **A** to **B** such that $h(a) \in L_a$ for each $a \in A$ exists, "no" otherwise.

By restricting $\text{LHOM}(\mathcal{C}, \mathcal{D})$ to those inputs in which each list L_a induces a connected subgraph of the Gaifman graph $\mathbf{G}(\mathbf{B})$ of $\mathbf{B}$, we get the *connected list homomorphism problem*, $\text{CLHOM}(\mathcal{C}, \mathcal{D})$, introduced for graphs in [6]:

INSTANCE: $\mathbf{A} \in \mathcal{C}$, $\mathbf{B} \in \mathcal{D}$, $L_a \subseteq B$ for each $a \in A$, such that each L_a induces a connected substructure in $\mathbf{B}$.
OUTPUT: "yes" if a homomorphism h from **A** to **B** such that $h(a) \in L_a$ for each $a \in A$ exists, "no" otherwise.

Feder and Hell introduce the *edge list homomorphism problem* for undirected graphs in [8]. Here we generalise this to arbitrary relational structures and let $\text{ELHOM}(\mathcal{C}, \mathcal{D})$ be the following problem:

INSTANCE: $\mathbf{A} \in \mathcal{C}$, $\mathbf{B} \in \mathcal{D}$, lists of tuples from the relations of $\mathbf{B}$ for each tuple of the relations in $\mathbf{A}$.
OUTPUT: "yes" if a homomorphism h from **A** to **B** such that each tuple in the relations of **A** maps to a tuple in the corresponding list of tuples from **B** exists, "no" otherwise.

In [16] an optimisation problem is introduced, where every *graph* homomorphism is associated with a cost. We generalise this framework to arbitrary relational structures. If each element $a \in A$ is associated with, positive integral, costs $c_b(a)$, $b \in B$, then the cost of a homomorphism h is $\sum_{a \in A} c_{h(a)}(a)$ and the *minimum cost homomorphism problem*, $\text{MINHOM}(\mathcal{C}, \mathcal{D})$, is the following problem:

INSTANCE: $\mathbf{A} \in \mathcal{C}$, $\mathbf{B} \in \mathcal{D}$, positive integer costs $c_b(a)$, where $a \in A$ and $b \in B$.
OUTPUT: The cost of a minimum cost homomorphism from **A** to **B**, "no" if no homomorphism from **A** to **B** exists.

If we let the universes B of the right hand side input structures of HOM$(\mathcal{C}, \mathcal{D})$ be finite subsets of the natural numbers equipped with the usual total order $<$, the *maximum solution problem* [21], MAX SOL$(\mathcal{C}, \mathcal{D})$, is the following problem:

INSTANCE: $\mathbf{A} \in \mathcal{C}$, $\mathbf{B} \in \mathcal{D}$, weight function $\omega : A \to \mathbb{N}$
OUTPUT: The maximum of $\sum_{a \in A} \omega(a) \cdot h(a)$ for any homomorphism h from **A** to **B**, "no" if no homomorphism from **A** to **B** exists.

We note that MAX SOL is an extension of the MAX ONES problem and, as in [22], where Khanna et al. classify the approximability of MAX ONES with respect to restrictions to the right hand side input structure, we restrict our attention to instances of MAX SOL satisfying the following restriction: if a, a' occur in the same tuple $(a_1, \ldots, a_r)$ in some relation in **A**, then $a \neq a'$ must hold. We say that a structure having this property is *replication free*.

2.3 Parameterised Complexity

Finally, we need some facts concerning parameterised complexity theory. Here we relax the classical notion of tractability, polynomial time computability, by admitting algorithms whose running time is exponential in some *parameter* of the problem instance that can be expected to be small in the typical application.

A parameterisation of a problem $P \subseteq \Sigma^*$ is a polynomial time computable mapping $\kappa : \Sigma^* \to \mathbb{N}$. If $(x, k) \in \Sigma^* \times \mathbb{N}$ is an instance of a parameterised decision problem, we call x the input and k the parameter. For example, the parameterised clique problem p-CLIQUE, is the following problem:

INPUT: graph **G**.
PARAMETER: $k \in \mathbb{N}$.
OUTPUT: "Yes" if **G** has a clique of size k, "no" otherwise.

A parameterised problem (P, κ) over Σ 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 the solution in time $f(k) \cdot |x|^c$. FPT denotes the class of all fixed-parameter tractable parameterised problems.

An *fpt-reduction* from a parameterised problem (P, κ) over Σ to a parameterised problem (P', κ') over Σ' is a mapping $R : \Sigma^* \to (\Sigma')^*$ such that for all $x \in \Sigma^*$ we have $R(x) \in P'$, R is computable in time $f(\kappa(x)) \cdot |x|^c$ and $\kappa'(R(x)) \leq g(\kappa(x))$ (for computable functions $f, g : \mathbb{N} \to \mathbb{N}$ and a constant c).

Hardness and completeness of parameterised problems for a parameterised complexity class are defined in the usual way. Downey and Fellows [4] defined a hierarchy W[1] $\subseteq$ W[2] $\subseteq \cdots$ of parameterised complexity classes. They conjecture that this hierarchy is strict and that FPT is strictly contained in W[1]. p-CLIQUE is shown to be W[1]-complete under fpt-reductions in [5]. This theorem is used in our hardness proofs.

The problems we are interested in are the homomorphism problems defined in Subsection 2.2 parameterised by the size of the left hand side input structure. E.g.

we have the following definition of the parameterised list homomorphism problem, p-LHOM$(\mathcal{C}, \mathcal{D})$:

INPUT: $\mathbf{A} \in \mathcal{C}$, $\mathbf{B} \in \mathcal{D}$, $L_a \subseteq B$ for each $a \in A$.
PARAMETER: $||\mathbf{A}||$.
OUTPUT: "yes" if a homomorphism h from $\mathbf{A}$ to $\mathbf{B}$ such that $h(a) \in L_a$ for each $a \in A$ exists, "no" otherwise.

The parameterised versions of the other problems in Subsection 2.2 are defined analogously and with the same parameter.

3 Main Results

We are now ready to prove the main results. First, we make the observation that when our homomorphism problems are restricted to classes of structures that have bounded tree-width, standard techniques using tree-decompositions, cf [17, 19], may be employed to solve the problems in question in polynomial time. Then we see that what is left to do to get a classification of our problems, with regard to structural restrictions, is to prove hardness for classes of structures with unbounded tree-width. The proofs need a bit of preparation, that is taken care of in Subsection 3.1. Subsection 3.2 contains the actual proofs.

3.1 The Structure $\mathbb{B}$

Let $\mathbf{A}$ be a connected τ-structure. Let $k \geq 2$, $K = \binom{k}{2}$, and $\mu : \{1, \ldots, k\} \times \{1, \ldots, K\} \to 2^A$ 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 write $i \in p$ instead of $i \in \varrho(p)$.

Let the $\{E\}$-structure $\mathbf{G}$ be a graph. We now concentrate on the τ-structure $\mathbf{B} = \mathbf{B}(\mathbf{A}, \mu, \mathbf{G})$, as defined by Grohe [15]. The universe B of $\mathbf{B}$ is given by:

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

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 component-wise.

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 $e = e'$.

In the remainder of this paper, we will focus on homomorphisms from $\mathbf{A}$ to $\mathbf{B}$ such that each $a \in A$ is mapped to an element $b \in B$ that was "generated" by a, i.e. $b \in \Pi^{-1}(a)$. We will denote this by saying that for a homomorphism $h : A \to B$,

$h(a) = (\_,\_,\_,\_,a)$ for each $a \in A$, where the placeholders '_' are used to indicate that the values in question are arbitrary, as long as the element is a member of B. To proceed we need the following fact:

Lemma 2. *The graph $\mathbf{G}$ contains a k-clique if and only if there exists a homomorphism h from $\mathbf{A}$ to $\mathbf{B}$ such that $h(a) = (\_,\_,\_,\_,a)$ for all $a \in A$.*

Proof. In the proof of Lemma 3.1 in [3] it is shown that the graph $\mathbf{G}$ contains a k-clique if and only if there exists a homomorphism h from $\mathbf{A}$ to $\mathbf{B}$ satisfying $\Pi \circ h = \mathbf{id}$, where $\mathbf{id}$ is the identity function on the set A. Now, if h is a homomorphism from $\mathbf{A}$ to $\mathbf{B}$ such that $h(a) = (\_,\_,\_,\_,a)$ for all $a \in A$, h obviously satisfies $\Pi \circ h = \mathbf{id}$ and vice versa. □

3.2 Hardness Results

The problem p-LHOM$(\mathcal{C},_)$ is trivially in FPT when LHOM$(\mathcal{C},_)$ is in FP, and we know that LHOM$(\mathcal{C},_)$ is solvable in polynomial time if the structures in $\mathcal{C}$ have bounded tree-width. What is left to prove, to achieve the result announced in Section 1, is that if p-LHOM$(\mathcal{C},_)$ is in FPT, then the structures in $\mathcal{C}$ have bounded tree-width. We do this by assuming that p-LHOM$(\mathcal{C},_)$ is in FPT even when $\mathcal{C}$ has unbounded tree-width and showing that this implies p-CLIQUE is in FPT, in contradiction with the fact that it is W$[1]$-complete. This is accomplished by exhibiting an fpt-reduction from p-CLIQUE to p-LHOM$(\mathcal{C},_)$, where the result in the previous subsection is applied. As the same reasoning applies to the four other problems under study, this proof is then adapted and extended to fit our different problem variations. However, due to space constraints, some proofs are omitted from this paper.

Lemma 3. *Let $\mathcal{C}$ be a recursively enumerable class of structures of bounded arity that does not have bounded tree-width. If either p-LHOM$(\mathcal{C},_)$, p-CLHOM$(\mathcal{C},_)$, p-ELHOM$(\mathcal{C},_)$ or p-MINHOM$(\mathcal{C},_)$ is in FPT, then FPT $=$ W$[1]$.*

Proof. 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 μ from the $(k \times K)$-grid to $\mathbf{A}$. Let $\mathbf{A_1}, \ldots, \mathbf{A_m}$ be a decomposition of $\mathbf{A}$ into its connected components. We can assume, without loss of generality, that the $(k \times K)$-grid is a minor of (the Gaifman graph of) $\mathbf{A_1}$ and that the minor map μ is onto $\mathbf{A_1}$.

Let $\mathbf{B} = (\mathbf{A}, \mu, \mathbf{G})$ constructed as above. By Lemma 2, we know that in order to decide if there exists a k-clique in $\mathbf{G}$ we only need to check if there is a homomorphism h from $\mathbf{A_1}$ to $\mathbf{B}$ such that h maps every $a \in A_1$ to some $(\_,\_,\_,\_,a) \in B$, since such an h exists if and only if $\mathbf{G}$ has a k-clique. We would like to differentiate $\mathbf{B}$, so that only homomorphisms mapping $a \in A_1$ to $(\_,\_,\_,\_,a) \in B$ are allowed. Fortunately, the list homomorphism problem lets us enforce precisely such a differentiation of $\mathbf{B}$.

To do this construct $\mathbf{B'}$ as $\mathbf{B} \cup \mathbf{A_2} \cup \ldots \cup \mathbf{A_m}$ and lists $L_a \subseteq B'$, $a \in A$ defined by:

$$L_a = \begin{cases} \{b \mid b \in B \text{ and } b = (\_,\_,\_,\_,a)\} \text{ if } a \in A_1 \\ \{b \mid b \in B' \setminus B, b = a\} \text{ otherwise} \end{cases}$$

This way, we will always be able to find a homomorphism from $\mathbf{A} \setminus \mathbf{A_1}$ to $\mathbf{B}' \setminus \mathbf{B}$: it is just a matter of selecting the only element b available in L_a for each $a \in A \setminus A_1$. Since $b = a$ in each case this obviously results in a homomorphism from $\mathbf{A} \setminus \mathbf{A_1}$ to $\mathbf{B}' \setminus \mathbf{B}$.

It is also clear that the only possible homomorphisms h from $\mathbf{A_1}$ to $\mathbf{B}$ (and hence also the only possible homomorphisms from $\mathbf{A}$ to $\mathbf{B}'$), under our lists, are the ones obeying the condition that h maps each $a \in A_1$ to some $(\_,\_,\_,\_,a) \in B$, due to the definition of the lists for elements $a \in A_1$.

Thus, the conclusion is that if $\mathbf{G}$ contains a k-clique, then we will be able to find a homomorphism from $\mathbf{A}$ to $\mathbf{B}'$, since then a homomorphism h from $\mathbf{A}_1$ to $\mathbf{B}$, obeying $h(a) = (\_,\_,\_,\_,a)$ for each $a \in A_1$, exists (by Lemma 2). If $\mathbf{G}$ has no k-clique, then we will not be able to find any homomorphism from $\mathbf{A}$ to $\mathbf{B}'$.

The construction of $\mathbf{A}$ only depends on k and is polynomial-time because $\mathcal{C}$ is recursively enumerable. Computing the minor map μ may require exponential time in the size of $\mathbf{A}$, but this is still bounded in terms of k. The size of an r-ary relation $R^{\mathbf{B}}$ is at most $|\Pi^{-1}(A^r)| \leq (|V^G| \cdot |E^G| \cdot |A|)^r$. This is polynomial in $||\mathbf{A}||$ and $||\mathbf{G}||$ since the arity of $\mathcal{C}$ is bounded. It follows that the size of $\mathbf{B}$ and $\mathbf{B}'$ is polynomially bounded in terms of $||\mathbf{A}||$ and $||\mathbf{G}||$ and so, $\mathbf{B}'$ can be computed in polynomial time. The lists L_a for $a \in A \setminus A_1$ are easy to compute and only hold one element each. While generating $\mathbf{B}$ it is easy to construct the lists L_a for $a \in A_1$ and the size of these lists are linear in the size of B. This shows that the reduction from $\mathbf{G}, k$ to $\mathbf{A}$, $\mathbf{B}'$, L_a is an fpt-reduction.

To be able to prove hardness for CLHOM we have to modify the structure $\mathbf{B}$ a bit; by adding some dummy elements to $\mathbf{B}$ we make our lists of elements in $\mathbf{B}$ induce connected substructures of $\mathbf{B}$. The result for ELHOM follows from a straightforward adaption of the proof for LHOM. Finally, the hardness result for MINHOM follows from transforming the instance of LHOM, in the proof of Lemma 3, to an instance of MINHOM by assigning $c_b(a) = 1$ if $b \in L_a$ and $c_b(a) = 2$ otherwise. □

An immediate consequence of the above is that the problem of *counting* list homomorphisms [18] is hard when $\mathcal{C}$ does not have bounded tree-width.

In the last reduction in the proof of Lemma 3, from p-CLIQUE to p-MINHOM$(\mathcal{C}, _)$, a gap that can be utilised to show the following (For further details regarding approximability we refer to [1].) is produced:

Proposition 4. *Let $\mathcal{C}$ be a recursively enumerable class of structures that does not have bounded tree-width. If* MINHOM$(\mathcal{C}, _)$ *is approximable within $2^{p(|A|)}$, (where p is a fixed polynomial), for every structure $\mathbf{A} \in \mathcal{C}$, then FPT = W[1].*

Before we continue dealing with our hardness results, a remark about our chosen proof method is in place. Why do we need to use the structure $\mathbf{B}$ at all, could we not just reduce LHOM$(\mathcal{C}, _)$ to HOM$(\mathcal{C}', _)$, for some suitable class $\mathcal{C}'$, i.e. for $\mathbf{A} \in \mathcal{C}$, let $\mathbf{A}' \in \mathcal{C}'$ be the expansion of $\mathbf{A}$ having a relation R_a for each $a \in A$ such that $R_a^{A'} = \{(a)\}$, and go from there? This way, an instance $(\mathbf{A}, \mathbf{B})$ of LHOM reduces to $(\mathbf{A}', \mathbf{B}')$, where $\mathbf{B}'$ has $R_a^{B'} = L_a$ and $R^{B'} = R^B$ for all other relations R. The error in this line of reasoning, is that the structure $\mathbf{A}$ might not allow unary relations on all its members. To illustrate this point, think of the problems HOM$(\mathcal{C}, _)$ for a class $\mathcal{C}$ of structures with unbounded tree-width. Using the method of adding unary relations to the structures in $\mathcal{C}$, described above, we can now modify the proof of Lemma 3 to become a hardness proof for $\text{HOM}(\mathcal{C}, \_)$! This is, of course, contradictory to Grohe's result. (If $\mathcal{C}$ is *restricted* to classes of structures that have all unary singleton relations, the cores of the structures can not be smaller than the structures themselves and the tree-widths of the structures and their respective cores coincide.)

In the hardness proof for MAX SOL, we are able to exploit the fact that we have to impose some total order on the elements in B; by letting elements of the form $(\_, \_, \_, \_, a)$, for some $a \in A$, have essentially the same values and inter-spacing these clusters with large gaps, the positive and negative instances of p-CLIQUE are separated.

Lemma 5. *Let $\mathcal{C}$ be a recursively enumerable class of replication free structures of bounded arity that does not have bounded tree-width. If p-MAX SOL$(\mathcal{C}, _)$ is in FPT, then FPT = W$[1]$.*

Proof. We start out as in the proof of Lemma 3 and construct $\mathbf{B'}$ as $\mathbf{B} \cup \mathbf{A_2} \cup \ldots \cup \mathbf{A_m}$. To proceed, we have to impose some total order on the elements in B'. Fix the natural order $<$ on $\mathbb{N}$. The intuition is to let elements in B on the form $(\_, \_, \_, \_, a)$, for some $a \in A_1$, have essentially the same values in B'. If these small intervals where the $(\_, \_, \_, \_, a) \in B$ reside, for each a, are inter-spaced by large gaps and the weights assigned to $a \in A_1$ are chosen accordingly we might be able to separate the positive and negative instances of p-CLIQUE.

Let $\sigma = \max_{a \in A_1} |\Pi^{-1}(a)|$, the maximum number of elements in B "generated" by an element in A_1. Clearly, σ is bounded in terms of k and $||\mathbf{G}||$.

Let $B' \setminus B = \{1, \ldots, d\}$. Also, let $w(a) = 0$ when $a \in A \setminus A_1$. Furthermore, take an $a \in A_1$, let $w(a) = d + 1$ and let each $b \in \Pi^{-1}(a)$ have a distinct value in $[d + 1, d + \sigma]$. The next $a \in A_1$ gets $w(a) = d + \Delta + 1$ while the associated $b \in \Pi^{-1}(a)$ get distinct values in $[d + \Delta + 1, d + \Delta + \sigma]$. We continue this process until A_1 is exhausted and end up with the arrangement in Figure 1.

We are interested in homomorphisms h between $\mathbf{A_1}$ and $\mathbf{B}$, such that each $a \in A_1$ maps to some $(\_, \_, \_, \_, a) \in B$, i.e. where the $a \in A_1$ with highest weight get mapped to some $(\_, \_, \_, \_, a) \in B$ in the highest interval of values, the $a \in A_1$ with second highest weight get mapped to some $(\_, \_, \_, \_, a) \in B$ in the second highest interval of values and so on. Such an h will receive a measure m_{id} with

$$(d+1)^2 + (d+\Delta+1)^2 + \ldots + (d + (|A_1| - 1)\Delta + 1)^2 \leq m_{id} \leq$$

$$\leq (d+1)(d+\sigma) + (d+\Delta+1)(d+\Delta+\sigma) + \ldots +$$

$$+ (d + (|A_1| - 1)\Delta + 1)(d + (|A_1| - 1)\Delta + \sigma).$$

It is easy to extend h to a homomorphism h' from $\mathbf{A}$ to $\mathbf{B'}$ (by mapping each $a \in A \setminus A_1$ to the $b \in B' \setminus B$ with $b = a$) and the measure for h' will still be m_{id}.

What false positives could we get? Recall that for each relation $R \in \tau$ and for all tuples $(a_1, \ldots, a_r) \in R^{\mathbf{A}_1}$, we add tuples $(b_1, \ldots, b_r) \in \Pi^{-1}(a_1, \ldots, a_r)$ satisfying certain conditions to $R^{\mathbf{B}}$ and that, in this case, $\mathbf{A}_1$ is replication free. This means that $\mathbf{B}$ is constructed so that any homomorphism h from $\mathbf{A}_1$ to $\mathbf{B}$ must have the property that the image of h contains at most one element from each interval, $[d+n\Delta+1, d+n\Delta+\sigma]$, of values in B.

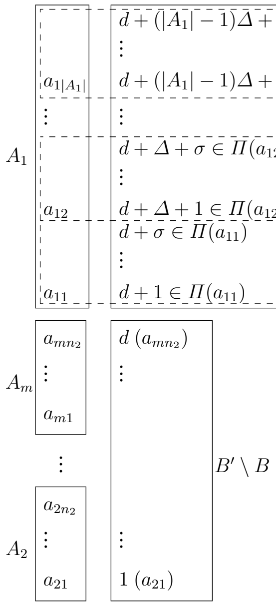

Fig. 1. The total order imposed on B'.

That leaves the possibility that some intervals of values have been permuted in some way, i.e. at least a pair of elements in A_1 have been mapped to somewhere in "each others" intervals. It can be shown by induction that the maximum measure of such a homomorphism occurs when the two elements in A_1 that have lowest weight have swapped intervals, i.e. we have $h(a_{11}) = (\_, \_, \_, \_, a_{12})$ and $h(a_{12}) = (\_, \_, \_, \_, a_{11})$ in Figure 1, and the maximum value of each interval is picked as image. This measure matches the maximum possible m_{id} except for the two first summands.

The difference, denoted δ, between the lowest possible m_{id} and the measure of such a homomorphism is

$$\delta = \sum_{n=1}^{|A_1|} (d + (n-1)\varDelta + 1)^2 - (d+1)(d + \varDelta + \sigma) - (d + \varDelta + 1)(d + \sigma) -$$

$$- \sum_{n=3}^{|A_1|} (d + (n-1)\varDelta + 1)(d + (n-1)\varDelta + \sigma),$$

which is the same as (omitting the calculations) δ being equal to

$$\varDelta^2 + \left(|A_1|^2 - |A_1| - \sigma|A_1|^2 + \sigma|A_1|\right)\varDelta/2 + |A_1| + d|A_1| - d\sigma|A_1| - \sigma|A_1|.$$

If we choose $\varDelta$ large enough, for example $\varDelta = d^2\sigma^2|A_1|^2$, the difference δ will be positive and hence, we can say that if we find a homomorphism with measure m_{id}, $\mathbf{G}$ has a k-clique and that if the maximum measure of any homomorphism from $\mathbf{A}$ to $\mathbf{B}'$ is strictly less than the smallest possible m_{id}, $\mathbf{G}$ contains no k-clique. $\square$

4 Conclusions and Open Questions

We have utilised the structure $\mathbf{B}$ defined by Grohe to classify a number of homomorphism problems by computational complexity with regard to structural restrictions, under the assumption that FPT $\neq$ W$[1]$. It is interesting to note that while the variants of the homomorphism problem we have treated have their boundary between tractability and intractability at bounded tree-width of the left hand side input structure, the original HOM$(\mathcal{C}, _)$ problem exhibits the same boundary at bounded tree-width for the core of the structures in $\mathcal{C}$. It would be interesting to characterise exactly what properties make the computational complexity of our problems different from that of the "regular" homomorphism problem.

Of course it would be nice to classify further homomorphism problems. E.g. the *retraction problem*, also known as the *one-or-all list homomorphism problem*, see [6], would be an interesting subject. Here, inputs of the list homomorphism problem are restricted to each list containing only a single element or the entire universe of the right hand side input structure.

In the reduction from p-CLIQUE to p-MINHOM$(\mathcal{C}, _)$ a gap that can be used to show inapproximability properties of the intractable instances is produced. A gap is also produced in the MAX SOL case, but it is not exploitable in the same way. Is it possible to change the reduction somewhat to achieve a gap large enough for proving inapproximability?

A further observation is that the structure $\mathbf{B}$, so far, only has been applied when classifying homomorphism problems: is it possible to modify the structure $\mathbf{B}$, or the analysis of it, so that hardness proofs for problems where the solution is not necessarily a homomorphism, e.g. MAX CSP, becomes plausible?

Acknowledgements

We appreciate the comments of an anonymous reviewer who helped improve the presentation of this paper.

References

1. G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann and A. Marchetti Spaccamela, Complexity and approximation: Combinatorial optimization problems and their approximability properties, Springer (1999).
2. A. A. Bulatov, Tractable conservative constraint satisfaction problems, ACM Transactions on Computational Logic, to appear.
3. V. Dalmau and P. Jonsson, The complexity of counting homomorphisms seen from the other side, Theoretical Computer Science 329(1-3) (2004) 315-323.
4. R. G. Downey and M. R. Fellows, Fixed-parameter tractability and completeness I: Basic results, SIAM Journal on Computing, 24 (1995) 873–921.
5. R. G. Downey and M. R. Fellows, Fixed-parameter tractability and completeness II: On completeness for W[1], Theoretical Computer Science, 141 (1995) 109–131.
6. T. Feder and P. Hell, List homomorphisms to reflexive graphs, J. Comb. Theory Series B 72 (1998) 236–250

7. T. Feder and P. Hell, Full constraint satisfaction problems, SIAM Journal on Computing, 36 (2006) 230–246.
8. T. Feder and P. Hell, Edge list homomorphisms, manuscript.
9. T. Feder and P. Hell and J. Huang, List homomorphisms and circular arc graphs, Combinatorica 19 (1999) 487–505.
10. T. Feder and P. Hell and J. Huang, Bi-arc graphs and the complexity of list homomorphisms, J. Graph Theory 42 (1999) 61–80.
11. T. Feder and P. Hell and J. Huang, List homomorphisms to reflexive digraphs, manuscript 2005.
12. T. Feder and P. Hell and J. Huang, List homomorphisms of graphs with bounded degree, to appear in Discrete Math.
13. 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.
14. 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.
15. M. Grohe, The complexity of homomorphism and constraint satisfaction problems seen from the other side, Journal of the ACM 54(1) (2007) 1–24.
16. G. Gutin, A. Rafiey, A. Yeo and M. Tso, Level of repair analysis and minimum cost homomorphisms of graphs, Discrete Appl. Math. 154 (2006), 881-889.
17. P. Hell, Algorithmic aspects of graph homomorphisms, in: Survey in combinatorics 2003, London Math. Society Lecture Note Series, Cambridge University Press 307 (2003), 239–276.
18. P. Hell and J. Nešetřil, Counting list homomorphisms for graphs with bounded degree, in: Graphs, morphisms and statistical physics, DIMACS Series in Discrete Mathematics and Theoretical Computer Science 63 (2004) 105–112.
19. I. Hicks, A. Koster and E. Kolotoğlu, Branch and tree decomposition techniques for discrete optimization, in: TutORials 2005, INFORMS TutORials in Operations Research Series (2005), 1–29.
20. P. Jeavons, On the algebraic structure of combinatorial problems, Theoretical Computer Science 200(1–2) (1998) 185–204.
21. P. Jonsson and G. Nordh, Generalised integer programming based on logically defined relations, in: Proceedings of the 31st International Symposium on Mathematical Foundations of Computer Science (2006) 549–560.
22. S. Khanna, M. Sudan, L. Trevisan and D. P. Williamson, The approximability of constraint satisfaction problems, SIAM Journal on Computing 30(6) (2001) 1863–1920.
23. P. G. Kolaitis and M. Vardi, Conjunctive-query containment and constraint satisfaction, Journal of Computer and System Sciences 61 (2000) 302–332.
24. L. G. Kroon, A. Sen, H. Deng and A. Roy, The optimal cost chromatic partition problem for trees and interval graphs, in: Graph-Theoretic Concepts in Computer Science (1997) 279–292.
25. N. Robertson and P. Seymour, Graph minors V: Excluding a planar graph, Journal of Combinatorial Theory, ser. B, 62 (1994) 323–348.