Understanding the Complexity of Induced Subgraph Isomorphisms

Yijia Chen[1], Marc Thurley[2], and Mark Weyer[2]

[1] BASICS, Department of Computer Science, Shanghai Jiaotong University, 200030 Shanghai, China
[2] Institut für Informatik, Humboldt-Universität zu Berlin, 10099 Berlin, Germany

Abstract. We study left-hand side restrictions of the *induced subgraph isomorphism* problem: Fixing a class $\mathcal{C}$, for given graphs $G \in \mathcal{C}$ and arbitrary H we ask for induced subgraphs of H isomorphic to G.
For the homomorphism problem this kind of restriction has been studied by Grohe and Dalmau, Kolaitis and Vardi for the decision problem and by Dalmau and Jonsson for its counting variant.
We give a dichotomy result for both variants of the induced subgraph isomorphism problem. Under some assumption from parameterized complexity theory, these problems are solvable in polynomial time if and only if $\mathcal{C}$ contains no arbitrarily large graphs.
All classifications are given by means of parameterized complexity. The results are presented for arbitrary structures of bounded arity which implies, for example, analogous results for directed graphs.
Furthermore, we show that no such dichotomy is possible in the sense of classical complexity. That is, if P $\neq$ NP there are classes $\mathcal{C}$ such that the induced subgraph isomorphism problem on $\mathcal{C}$ is neither in P nor NP-complete. This argument may be of independent interest, because it is applicable to various parameterized problems.

1. Introduction

Given graphs G and H, the *induced subgraph isomorphism problem* asks for the existence of induced subgraphs of H isomorphic to G. A wide variety of graph theoretic problems can be formulated in this way, as is the case for the induced path or induced cycle problem. By the fact that the independent set problem is also an induced subgraph isomorphism problem, the latter is obviously NP-complete.

This inherent intractability is highly unsatisfactory, as it does not give any insight into the complexity of more restricted subproblems such as e.g. the induced path problem, necessitating separate investigation [3]. To uniformly study the complexity of such subproblems of the induced subgraph isomorphism problem, we therefore consider restrictions of this problem in the following way. Fixing a class $\mathcal{C}$ of graphs, we consider only inputs $G \in \mathcal{C}$ whereas H is still an arbitrary graph.

For the related homomorphism problem, the complexity of this kind of restrictions has been described in [5, 10] in terms of a dichotomy: If $\mathcal{C}$ has bounded treewidth up to homomorphic equivalence, the homomorphism problem is in polynomial time, otherwise it is intractable in the sense of parameterized complexity. A similar dichotomy from [4] states that for the counting version polynomial time is equivalent to bounded treewidth.

In this paper we settle the complexity of the restricted induced subgraph isomorphism problem by giving a further dichotomy. For both the decision and the counting variant, we show that the problem is computable in polynomial time if and only if there is an absolute bound on the size of the graphs in the class $\mathcal{C}$. Otherwise, the problem is intractable in terms of parameterized complexity.

We give two different proofs for the two versions of this problem. We prove the dichotomy for the decision problem based on the result from [10]. For the counting version, we give a more direct proof using an inclusion-exclusion style argument. This cannot be applied to the decision case.

Furthermore, all proofs will be given for arbitrary structures of bounded arity. Therefore, our results are not only applicable to graphs. They also extend to analogous results for the induced subgraph isomorphism problems on e.g. directed graphs, on coloured graphs, and on hypergraphs with bounded edge-size.

The fact that our hardness results rely on parameterized complexity theory raises the question of whether a similar dichotomy in terms of classical complexity could possibly be established. Using a Ladner-style argument (compare [11]) we show that this is not the case, unless P = NP (or FP = #P for the counting problem, resp.). More precisely, there are classes $\mathcal{C}$ such that the restricted induced subgraph isomorphism problem is neither in P nor NP-complete. This result is presented in a universal way which enables us to derive analogs for e.g. the homomorphism problem and the corresponding counting problems. Note that, for the homomorphism problem itself, a similar result has also been shown independently by [2].

Due to space limitations we have to defer some proofs to the full version of the paper.

2. Preliminaries

Structures. We only consider relational vocabularies. Hence, a *vocabulary* is a set of relational symbols, each having an arity in $\mathbb{N}$. The arity of the symbol R is denoted $\mathrm{ar}(R)$. The arity of a vocabulary is the maximal arity of its symbols. Let τ be a vocabulary. A *structure* $\mathfrak{A}$ of vocabulary τ, or τ*-structure* for short, is a tuple $(A, (R^{\mathfrak{A}})_{R\in\tau})$, where the *universe* A of $\mathfrak{A}$ is some set and $R^{\mathfrak{A}} \subseteq A^{\mathrm{ar}(R)}$ for all $R \in \tau$. As we have done here, whenever we denote a structure by a German type letter, its universe is implicitly denoted by the corresponding Roman type letter. For algorithmic purposes, all vocabularies and universes are finite. The arity of a structure is the arity of its vocabulary.

Let $\mathfrak{A}$ and $\mathfrak{B}$ be structures of the same vocabulary τ. $\mathfrak{A}$ is a *substructure* of $\mathfrak{B}$, if $A \subseteq B$ and $R^{\mathfrak{A}} \subseteq R^{\mathfrak{B}}$ for all $R \in \tau$. $\mathfrak{A}$ is an *induced substructure* of $\mathfrak{B}$, if furthermore $R^{\mathfrak{A}} = R^{\mathfrak{B}} \cap A^{\mathrm{ar}(R)}$ for all $R \in \tau$. A *homomorphism* from $\mathfrak{A}$ to $\mathfrak{B}$ is a function $f : A \to B$ such that for all $R \in \tau$, we have $f(R^{\mathfrak{A}}) \subseteq R^{\mathfrak{B}}$. An *embedding* is a homomorphism that is injective. A *strong embedding* is an embedding f such that $f(R^{\mathfrak{A}}) = R^{\mathfrak{B}} \cap f(A)^{\mathrm{ar}(R)}$ for all $R \in \tau$. Note that (strong) embeddings coincide with isomorphisms to (induced) substructures.

For a structure $\mathfrak{A}$, say of vocabulary τ, the *Gaifman graph* $G(\mathfrak{A})$ of $\mathfrak{A}$ is the graph with vertex set A such that there is an edge between a and a' for $a \neq a'$ if and only if there is some $R \in \tau$, say of arity r, some $(a_1, \ldots, a_r) \in R^{\mathfrak{A}}$, and some $1 \leq i, j \leq r$ such that $a = a_i$ and $a' = a_j$.

Parameterized Complexity. Let Σ be a finite alphabet. A *parameterization of* Σ is a polynomial time computable mapping $\kappa : \Sigma^* \to \mathbb{N}$. A *parameterized decision problem* is a pair (P, κ) with $P \subseteq \Sigma^*$ and κ a parameterization. Similarly, a *parameterized counting problem* is a pair (F, κ) with $F : \Sigma^* \to \mathbb{N}$ and κ a parameterization. For problem instances $x \in \Sigma^*$, the value $\kappa(x)$ is called the *parameter* of x.

An algorithm is a *fixed-parameter algorithm*, if there is a computable function $f : \mathbb{N} \to \mathbb{N}$ and a constant $c \in \mathbb{N}$ such that for all $x \in \Sigma^*$ the algorithm stops after at most $f(\kappa(x)) \cdot |x|^c$ steps. A parameterized decision problem (P, κ) is *fixed-parameter tractable*, if there is a fixed-parameter algorithm which, for all $x \in \Sigma^*$, decides if $x \in P$. The class of all such problems is denoted by FPT. *Fixed-parameter tractable parameterized counting problems* are defined analogously, with fixed-parameter algorithms computing $F(x)$ and FFPT being the class of all of these problems.

In all well-behaved cases, our hardness results hold for the usual strongly uniform reductions. In general, however, we need nonuniform reductions for technical reasons. Therefore we define both versions, starting with the nonuniform variants. Note that FPT and FFPT are not closed under these, nor do we need them to be.

An FPT *many-one reduction* from a parameterized problem (P, κ) to a parameterized problem (P', κ') with $P \subseteq \Sigma^*$ and $P' \subseteq (\Sigma')^*$ is a family $(f_i : \Sigma^* \to (\Sigma')^*)_{i \in \mathbb{N}}$ with the following properties: There is a $c \in \mathbb{N}$ and some $h : \mathbb{N} \to \mathbb{N}$ such that on inputs of length n, every f_i is computable in time $h(i) \cdot n^c$. Furthermore there is some $g : \mathbb{N} \to \mathbb{N}$ such that for all $x \in \Sigma^*$ and $y := f_{\kappa(x)}(x)$, we have $x \in P$ if and only if $y \in P'$ and $\kappa'(y) \leq g(\kappa(x))$.

The corresponding reductions for counting problems are defined similarly. Let (F, κ) and (F', κ') be parameterized counting problems with $F : \Sigma^* \to \mathbb{N}$ and $F' : (\Sigma')^* \to \mathbb{N}$. An FPT *parsimonious reduction* from (F, κ) to (F', κ') is a family $(f_i : \Sigma^* \to (\Sigma')^*)_{i \in \mathbb{N}}$ such that for a fixed $c \in \mathbb{N}$ and $h : \mathbb{N} \to \mathbb{N}$, on inputs of length n, every f_i is computable in time $h(i) \cdot n^c$. Furthermore there is some $g : \mathbb{N} \to \mathbb{N}$ such that for all $x \in \Sigma^*$ and $y := f_{\kappa(x)}(x)$ we have $F(x) = F'(y)$ and $\kappa'(y) \leq g(\kappa(x))$.

To develop our results, we need a second notion of reductions between counting problems. With (F, κ) and (F', κ') as above, an FPT *Turing reduction* from (F, κ) to (F', κ') is a family of functions $(f_i : \Sigma^* \to \mathbb{N})_{i \in \mathbb{N}}$ with the following properties. For all $x \in \Sigma^*$ we have $F(x) = f_{\kappa(x)}(x)$. Furthermore, there are $g, h : \mathbb{N} \to \mathbb{N}$ and $c \in \mathbb{N}$ such that on inputs on length n, every f_i can be computed in time $h(i) \cdot n^c$ by an algorithm with oracle access to F' such that every oracle query $F'(y)$ satisfies $\kappa'(y) \leq g(i)$.

The above reductions become *strongly uniform* if we further stipulate that the given families and the mappings $g, h : \mathbb{N} \to \mathbb{N}$ be computable.

Downey and Fellows [6] defined a hierarchy of complexity classes of decision problems FPT $\subseteq$ W[1] $\subseteq$ W[2] $\subseteq \ldots$, conjecturing that all of the given inclusions are strict. Analogs of this hierarchy in terms of counting problems have been proposed in [12] and, in slightly different form, in [8]. The differences between these definitions do not affect the aim of our paper. For our purposes, we rely on [8] which define a hierarchy FFPT $\subseteq$ #W[1] $\subseteq$ #W[2] $\subseteq \ldots$ with the same conjecture as before, that all inclusions are strict. We will be concerned only with the first level W[1] (#W[1], respectively) of these hierarchies.

Usually, these classes are defined such that they are closed under strongly uniform reductions. Nonuniform versions are immediate, however. For ease of presentation, we will denote both versions by W[1] respectively #W[1]. We will rely on two results, namely Theorem 7 explained below and the following.

Theorem 1 (Flum and Grohe [8]). *p*-#Clique *is* #W[1]*-complete under strongly uniform* FPT *parsimonious reductions, where p*-#Clique *is* *the classical* Clique *problem parameterized by the size of the clique we are looking for.*

As each strongly uniform reduction is nonuniform, this gives hardness for both variants of #W[1].

For a class $\mathcal{C}$ of structures, the problem Hom($\mathcal{C}$) asks, when given a structure $\mathfrak{A} \in \mathcal{C}$ and an arbitrary structure $\mathfrak{B}$, whether there is a homomorphism from $\mathfrak{A}$ to $\mathfrak{B}$. Similarly, we define restrictions of the problems Emb and StrEmb asking, whether a structure $\mathfrak{A}$ is isomorphic to a substructure, respectively an induced substructure, of a structure $\mathfrak{B}$. In some contexts these problems are called the *embedding* and *strong embedding* problem – hence the abbreviations Emb and StrEmb.

Further, #Hom, #Emb, #StrEmb are their counting analogs, which ask for the number of such homomorphisms or isomorphisms, and p-Hom, p-Emb, p-StrEmb, p-#Hom, p-#Emb, and p-#StrEmb are the parameterized versions, where the parameter is $|\mathfrak{A}|$. Note that the membership of all of these problems in W[1] (#W[1], respectively) is well-known. This follows, for example, from [7] and [8].

3. The Dichotomies

Throughout this paper we assume that $\mathcal{C}$ is a class of structures of bounded arity, i.e., there is a bound r_0 such that no structure in $\mathcal{C}$ has arity beyond r_0.

Furthermore, we say that $\mathcal{C}$ is *meagre*, if there is some $n_0 \in \mathbb{N}$ such that for all $\mathfrak{A} \in \mathcal{C}$ of arity at least 2 we have $|A| \leq n_0$.

Theorem 2 (p-StrEmb($\cdot$) Dichotomy). *Let $\mathcal{C}$ be a class of structures of bounded arity.*

If $\mathcal{C}$ is meagre, then StrEmb($\mathcal{C}$) $\in$ P*. Otherwise,* p-StrEmb($\mathcal{C}$) *is complete for* W[1] *using nonuniform* FPT *many-one reductions.*

If $\mathcal{C}$ is recursively enumerable, then W[1]*-completeness holds even for strongly uniform* FPT *many-one reductions.*

Theorem 3 (p-#StrEmb($\cdot$) Dichotomy). *Let $\mathcal{C}$ be a class of structures of bounded arity.*

If $\mathcal{C}$ is meagre, then #StrEmb($\mathcal{C}$) $\in$ FP*. Otherwise,* p-#StrEmb($\mathcal{C}$) *is complete for* #W[1] *using nonuniform* FPT *Turing reductions.*

If $\mathcal{C}$ is recursively enumerable, then #W[1]*-completeness holds even for strongly uniform* FPT *Turing reductions.*

For membership in P or FP, we view our problems as promise problems, unless $\mathcal{C}$ happens to be decidable in polynomial time. When we consider classes of graphs instead of classes of arbitrary structures, then arity 2 is guaranteed. Hence meagreness just means bounded size and the theorems read as follows:

Corollary 4. *Let $\mathcal{C}$ be a class of graphs.*

If the graphs in $\mathcal{C}$ have bounded size, then $\mathrm{StrEmb}(\mathcal{C}) \in \mathrm{P}$*. Otherwise,* p-$\mathrm{StrEmb}(\mathcal{C})$ *is complete for* $\mathrm{W}[1]$ *under FPT many-one reductions.*

The analogue holds for the counting problem.

The first parts of Theorem 2 and Theorem 3 are easy:

Lemma 5. *If $\mathcal{C}$ is meagre, then* $\#\mathrm{StrEmb}(\mathcal{C}) \in \mathrm{FP}$ *and* $\mathrm{StrEmb}(\mathcal{C}) \in \mathrm{P}$*.*

Hence, in the following we may assume that $\mathcal{C}$ contains arbitrarily large structures of arity at least 2.

Roughly speaking, we want to find structures in $\mathcal{C}$ which exhibit large cliques. We can do this only up to taking complements. So, for a τ-structure $\mathfrak{A}$, define the complement $\mathfrak{A}^{\mathrm{comp}}$ of $\mathfrak{A}$ as the following τ-structure: The universe is again A, and for each relational symbol $R \in \tau$, say of arity r, we have $R^{\mathfrak{A}^{\mathrm{comp}}} = A^r \setminus R^{\mathfrak{A}}$. For a class $\mathcal{C}$ of structures, let $\mathcal{C}^{\mathrm{comp}} := \{\mathfrak{A}^{\mathrm{comp}} \mid \mathfrak{A} \in \mathcal{C}\}$ be the class of complements of structures in $\mathcal{C}$. Note that this is not the complement of $\mathcal{C}$.

The following lemma is immediate.

Lemma 6. p-$\mathrm{StrEmb}(\mathcal{C}) \equiv^{\mathrm{FPT}} p$-$\mathrm{StrEmb}(\mathcal{C}^{\mathrm{comp}})$ *by parsimonious reductions.*

For a structure $\mathfrak{A}$ and a symbol R in its vocabulary, say of arity $r \geq 2$, let

$$D(\mathfrak{A}, R) := (A, \{(a, b) \in A^2 \mid a \neq b, (a, \ldots, a, b) \in R^{\mathfrak{A}}\})$$

be the digraph associated with $\mathfrak{A}$ and R. For any given $k \in \mathbb{N}$ and sufficiently large A, Ramsey's Theorem guarantees that $D(\mathfrak{A}, R)$ contains a clique or a tournament or an independent set of size k. Then, at least one of $D(\mathfrak{A}, R)$ and $D(\mathfrak{A}^{\mathrm{comp}}, R)$ contains a clique or a tournament of size k. Hence, for at least one of $\mathcal{C}$ and $\mathcal{C}^{\mathrm{comp}}$ we can find arbitrarily large cliques or tournaments in the digraphs associated with its structures. Using Lemma 6, we can assume without loss of generality that this is the case for $\mathcal{C}$. Then, in particular, the Gaifman graphs of structures in $\mathcal{C}$ contain arbitrarily large cliques.

From this point on, the proofs for Theorem 2 and Theorem 3 diverge.

3.1. Hardness of Deciding. First to the proof of Theorem 2. We use $\leq^{\mathrm{FPT}}$ to denote nonuniform FPT many-one reducibility in the general case. If $\mathcal{C}$ is recursively enumerable, we instead intend it to denote strongly uniform FPT many-one reducibility.

We base the hardness part of Theorem 2 on the following result:

Theorem 7 (Grohe [10]). *If $\mathcal{C}$ is a class of structures with* cores *of unbounded* treewidth*, then p-*Hom*$(\mathcal{C})$ is hard for* W[1]*.*

Some explanations are in order. Hardness uses, just as we need it, nonuniform FPT many-one reductions, which are strongly uniform in case $\mathcal{C}$ is recursively enumerable. As to the notions of cores and treewidth, let us omit the definitions and just state the two facts we actually need:

1. The core of a structure $\mathfrak{A}$ is some particular homomorphic image of $\mathfrak{A}$ in $\mathfrak{A}$.
2. If $G(\mathfrak{A})$ contains a clique of size k, then the treewidth of $\mathfrak{A}$ is at least $k-1$.

A relation $R \subseteq A^r$ is *antireflexive*, if for all $a \in A$ we have $(a, \ldots, a) \notin R$. A structure is *antireflexive*, if all its relations are. For a given structure $\mathfrak{A}$, the *antireflexive part* $\mathfrak{A}^{\text{antiref}}$ of $\mathfrak{A}$ is obtained from $\mathfrak{A}$ by deleting all tuples of the form $(a, \ldots, a)$ from all relations of $\mathfrak{A}$. Further, let $\mathcal{C}^{\text{antiref}} := \{\mathfrak{A}^{\text{antiref}} \mid \mathfrak{A} \in \mathcal{C}\}$. If $\mathcal{C}$ is recursively enumerable, then so is $\mathcal{C}^{\text{antiref}}$.

Lemma 8. *p-Hom$(\mathcal{C}^{\text{antiref}}) \leq^{\text{FPT}} p$-StrEmb$(\mathcal{C})$.*

Proof: Assume given an input $(\mathfrak{A}', \mathfrak{B})$ to the reduction. Let $\mathfrak{A}$ be such, that $\mathfrak{A}' = \mathfrak{A}^{\text{antiref}}$. In case $\mathcal{C}$ is recursively enumerable, such an $\mathfrak{A}$ can be found effectively, otherwise there is no need for effectiveness because we use nonuniform reductions.

Define the structure $\mathfrak{C}$ as the following variant of $\mathfrak{A} \otimes \mathfrak{B}$: The universe is $C = A \times B$ and for each symbol R, say of arity r, let

$$\begin{aligned} R^{\mathfrak{C}} := &\{((a_1, b_1), \ldots, (a_r, b_r)) \mid (a_1, \ldots, a_r) \in R^{\mathfrak{A}}, (b_1, \ldots, b_r) \in R^{\mathfrak{B}}\} \\ &\cup \{((a, b), \ldots, (a, b)) \mid (a, \ldots, a) \in R^{\mathfrak{A}}, b \in B\}. \end{aligned}$$

Now if $f : A \to B$ is a homomorphism from $\mathfrak{A}'$ to $\mathfrak{B}$, then $g : A \to C$ defined by $g(a) = (a, f(a))$ is a strong embedding of $\mathfrak{A}$ in $\mathfrak{C}$. Conversely, if $g : A \to C$ is a strong embedding of $\mathfrak{A}$ in $\mathfrak{C}$, then the projection of g to B is a homomorphism from $\mathfrak{A}'$ to $\mathfrak{B}$.

Hence, the reduction outputs $(\mathfrak{A}, \mathfrak{C})$. □

Proof of Theorem 2: Lemma 5 already covered the lower part of the dichotomy. For the upper part, membership in W[1] is widely known (see e.g. [9].

For hardness, let $\mathcal{C}$ contain arbitrarily large structures of arity at least 2. We have already seen that without loss of generality, the digraphs associated with structures from $\mathcal{C}$ contain arbitrarily large cliques or tournaments. For ease of presentation, let us assume only the latter. Say, $\mathfrak{A} \in \mathcal{C}$ and R satisfy that $D(\mathfrak{A}, R)$ contains a tournament of size k. Then $D(\mathfrak{A}^{\text{antiref}}, R)$ still contains the tournament. Every homomorphic image of this tournament into an antireflexive structure is necessarily injective. As $\mathfrak{A}^{\text{antiref}}$ itself is antireflexive, it follows for the core $\mathfrak{A}'$ of $\mathfrak{A}^{\text{antiref}}$, that $D(\mathfrak{A}', R)$ contains a tournament as a subdigraph. Then, $G(\mathfrak{A}')$ contains a clique of size k, so $\mathfrak{A}'$ has treewidth at least $k-1$. As k is arbitrary, the cores of structures from $\mathcal{C}^{\text{antiref}}$ have unbounded treewidth.

Using Theorem 7, we conclude that p-Hom($\mathcal{C}^{\text{antiref}}$) is W[1]-hard. Lemma 8 then implies the W[1]-hardness of p-StrEmb($\mathcal{C}$). □

3.2. Hardness of Counting. We now turn to the proof for the counting problems. Generally, we use $\leq^{\text{FPT-T}}$ to denote nonuniform FPT Turing reducibility. If $\mathcal{C}$ is recursively enumerable, we instead intend it to denote strongly uniform FPT Turing reducibility.

Proof of Theorem 3: Membership of p-#StrEmb($\mathcal{C}$) in #W[1] is well-known [8] and if $\mathcal{C}$ is meagre, then Lemma 5 implies membership in FP.

So we may assume that $\mathcal{C}$ contains arbitrarily large structures of arity at least 2. By the above considerations, we can accordingly assume that the Gaifman graphs of structures in $\mathcal{C}$ contain arbitrarily large cliques. We show hardness, by giving an FPT Turing reduction from p-#Clique to p-#StrEmb($\mathcal{C}$). Let $G = (V, E)$ be a graph and $k \in \mathbb{N}$. First we find a structure $\mathfrak{A} \in \mathcal{C}$ such that $G(\mathfrak{A})$ contains a k-clique. We assume $A = [k']$ with $k' \geq k$ and $G(\mathfrak{A}) \restriction [k]$ is a k-clique. Let τ be the vocabulary of $\mathfrak{A}$. We define a τ-structure $\mathfrak{B} = \mathfrak{B}(\mathfrak{A}, G, k)$ with universe

$$B := \big(V \times [k]\big) \,\dot\cup\, [k+1, k'].$$

To define the relations of $\mathfrak{B}$, we need two projections $\pi_1 : B \to V \,\dot\cup\, \{\bot\}$ and $\pi_2 : B \to A$ defined by

$$\pi_1(b) := \begin{cases} u, & \text{if } b = (u,i) \text{ for some} \\ & \quad u \in V \text{ and } i \in [k] \\ \bot, & \text{if } b \in [k+1, k'], \end{cases} \quad \pi_2(b) := \begin{cases} i, & \text{if } b = (u,i) \text{ for some} \\ & \quad u \in V \text{ and } i \in [k] \\ b, & \text{if } b \in [k+1, k']. \end{cases}$$

Now for every $R \in \tau$ with arity r we let

$$R^{\mathfrak{B}} := \big\{(b_1, \dots, b_r) \in B^r \mid (\pi_2(b_1), \dots, \pi_2(b_r)) \in R^{\mathfrak{A}} \text{ and } \{\pi_1(b_1), \dots, \pi_1(b_r)\} \setminus \{\bot\} \text{ is a clique in } G\big\}. \tag{1}$$

By our bounded arity assumption, we always have $r \leq r_0$ here, hence $|\mathfrak{B}|$ is polynomial in $|\mathfrak{A}|$ and k.

Let h be a strong embedding from $\mathfrak{A}$ to $\mathfrak{B}$. We call h *good* if

$$\pi_2(h(A)) = [k'] \tag{2}$$

Note that this implies that π_2 is bijective on $h(A)$. Then we can establish:

Claim 1. For every good h, if we let

$$\big\{(v_i, i)\big\} := h(A) \cap \big(V \times \{i\}\big)$$

for every $i \in [k]$, then the set $\{v_1, \dots, v_k\}$ is a k-clique in G. ⊣

The proof of the next claim is straightforward.

Claim 2. Let $\bar{u} := (u_1, \dots, u_k) \in V^k$ such that $\{u_1, \dots, u_k\}$ is a k-clique in G. Then the mapping $h_{\bar{u}} : A \to B$ with

$$h_{\bar{u}}(i) := \begin{cases} (u_i, i), & \text{if } i \in [k], \\ i, & \text{if } i \in [k+1, k'] \end{cases}$$

is a good strong embedding from $\mathfrak{A}$ to $\mathfrak{B}$. ⊣

Let η be the number of good strong embeddings from $\mathfrak{A}$ to $\mathfrak{B}$, $\alpha := \big|\mathrm{Aut}(\mathfrak{A})\big|$ the number of automorphisms of $\mathfrak{A}$, and κ the number of k-cliques in G. Then:

Claim 3.

$$\kappa = \frac{\eta}{\alpha \cdot k!}$$

⊣

Theorem 3 now follows, if we can show how to compute the number η of good strong embeddings from $\mathfrak{A}$ to $\mathfrak{B}$. This is done using the principle of inclusion and exclusion:

For a τ-structure $\mathfrak{B}$ and a set $X \subseteq B$ let $\mathfrak{B}[X]$ denote the *induced* substructure of $\mathfrak{B}$ defined by X, i.e. $\mathfrak{B}[X] = (X, (R^{\mathfrak{B}} \cap X^{ar(R)})_{R \in \tau})$. Define for every set $I \subseteq [k']$ the structure $\mathfrak{B}_I := \mathfrak{B}[\pi_2^{-1}(I)]$. Let $\mathrm{StrEmb}(\mathfrak{A}, \mathfrak{B}_I)$ denote the set of strong embeddings from $\mathfrak{A}$ to $\mathfrak{B}_I$ and let b_I be the value returned by the p-#StrEmb($\mathcal{C}$) oracle on input $(\mathfrak{A}, \mathfrak{B}_I)$, i. e. $b_I =$

$|\mathrm{StrEmb}(\mathfrak{A}, \mathfrak{B}_I)|$. Furthermore, define C_I as the set of strong embeddings $f : \mathfrak{A} \to \mathfrak{B}$ satisfying $\pi_2(f(A)) = I$. Let $c_I := |C_I|$. The definition of C_I immediately implies that $C_{[k']}$ is the set of all good strong embeddings of $\mathfrak{A}$ into $\mathfrak{B}$.

Obviously, $\mathrm{StrEmb}(\mathfrak{A}, \mathfrak{B}_I) = \dot{\bigcup}_{I' \subseteq I} C_{I'}$ for all $I \subseteq [k']$. Hence

$$c_I = b_I - \sum_{I' \subsetneq I} c_{I'}.$$

Then, by recursion on $|I|$, we can compute the $2^{k'}$ values c_I from our knowledge of the values b_I. As k' is bounded in terms of the parameter k, we can compute all of the $2^{k'}$ values by $2^{k'}$ oracle calls within the time bounds of an FPT Turing reduction. □

4. The Nondichotomies

For a parameterized (decision, counting, or otherwise) problem (Q, κ) and $A \subseteq \mathbb{N}$, the *restriction* of (Q, κ) to A, denoted $(Q, \kappa) \restriction A$, is the classical problem Q, restricted to inputs x such that $\kappa(x) \in A$. We consider the case that A is decidable in polynomial time ($A \in \mathrm{P}$), when numbers are encoded in unary.

Let $\leq$ denote the reducibility for Q, e.g. polynomial time many-one reducibility for decision problems and polynomial time Turing reducibility for counting problems. If $A_1, A_2 \subseteq \mathbb{N}$ are in P and $A_1 \subseteq A_2$, and if Q is nontrivial, then, clearly, $(Q, \kappa) \restriction A_1 \leq (Q, \kappa) \restriction A_2$. Hence the lattice of polynomial time decidable subsets of $\mathbb{N}$ induces a partial order of degrees witnessed by restrictions of (Q, κ). We now establish a dense linear suborder.

Theorem 9. *Let (Q, κ) be a parameterized problem. Assume that Q is not solvable in polynomial time, but that (Q, κ) is solvable in* XP *time, i. e. on input x in time $|x|^{g(\kappa(x))}$ for some function $g : \mathbb{N} \to \mathbb{N}$.*

Then there is a dense linear order $\mathcal{O}$ of polynomial time decidable subsets of $\mathbb{N}$ such that for all $A_1, A_2 \in \mathcal{O}$ with $A_1 \subsetneq A_2$ we have $(Q, \kappa) \restriction A_2 \not\leq (Q, \kappa) \restriction A_1$.

The proof follows the lines of Ladner's classical argument.

Corollary 10. *If* $\mathrm{P} \neq \mathrm{NP}$*, respectively* $\mathrm{P} \neq \#\mathrm{P}$*, then the complexities of problems of the form* $\mathrm{StrEmb}(\mathcal{C})$*, respectively* $\#\mathrm{StrEmb}(\mathcal{C})$*, with $\mathcal{C}$ being some polynomial time decidable class of graphs, contain a dense linear order between* P *and* NP*, respectively between* P *and* #P*.*

The same holds for the homomorphism and embedding problems, and for structures instead of graphs.

The order's denseness implies that there is no finite classification of the unparameterized complexities of problems of the form StrEmb($\mathcal{C}$). Contrast this to our dichotomies.

As noted in the introduction, [2] contains an independently obtained proof of the fact that there is a polynomial time decidable class $\mathcal{C}$ of structures such that Hom($\mathcal{C}$) is neither in P nor complete for NP.

5. Conclusion and Open Problems

We give dichotomy results for the complexity of the restricted induced subgraph isomorphism problem. The upper parts of both our dichotomies are parameterized hardness, while the lower parts are classical tractability. Strong evidence is given that classifications of these problems cannot be given by classical complexity theory alone.

We were not able to classify the restricted subgraph isomorphism problem p-Emb($\mathcal{C}$). The decision problem is known to be fixed-parameter tractable if $\mathcal{C}$ is of bounded tree-width [1]. If $\mathcal{C}$ is of unbounded treewidth modulo homomorphic equivalence, then the problem is easily seen to be W[1]-hard. A natural example of the remaining cases is $\mathcal{C} = \{K_{k,k} \mid k \in \mathbb{N}\}$ for which p-Emb($\mathcal{C}$) coincides with the complete bipartite subgraph problem ([9], p. 355).

A classification of the counting problem p-#Emb($\mathcal{C}$) is wide open as well. Clearly, this problem is hard if $\mathcal{C}$ is of unbounded treewidth. However, treewidth is not the measure of choice here, as p-#Emb($\mathcal{C}$) is hard even if $\mathcal{C}$ is the class of all paths [8]. A natural example of the unknown cases is the parameterized problem of counting matchings [8].

Acknowledgement. Lemma 8 and its proof were suggested by an anonymous referee. The authors' original proof of Theorem 2 was much more complicated. Further thanks are due to Jörg Flum, Martin Grohe and Berit Grußien for comments on earlier versions of this paper. The first author also acknowledges the support of the National Nature Science Foundation of China (60673049).

References

1. Noga Alon, Raphael Yuster, and Uri Zwick. Color-coding. *J. ACM*, 42(4):844–856, 1995.

2. Manuel Bodirski and Martin Grohe. Non-dichotomies in Constraint Satisfaction Complexity. To appear in *Proceedings of the 35th International Colloquium on Automata, Languages and Programming (ICALP'08, Track B)*, 2008.
3. Yijia Chen and Jörg Flum. On parameterized path and chordless path problems. In *Proceedings of the Twenty-Second Annual IEEE Conference on Computational Complexity*, pages 250–263. IEEE Computer Society, 2007.
4. Víctor Dalmau and Peter Jonsson. The complexity of counting homomorphisms seen from the other side. *Theor. Comput. Sci.*, 329(1-3):315–323, 2004.
5. Víctor Dalmau, Phokion G. Kolaitis, and Moshe Y. Vardi. Constraint satisfaction, bounded treewidth, and finite-variable logics. In *Proceedings of the 8th International Conference on Principles and Practice of Constraint Programming, CP 2002*, pages 310–326, 2002.
6. Rodney G. Downey and Michael R. Fellows. Fixed-parameter tractability and completeness I: Basic results. *SIAM J. Comput.*, 24(4):873–921, 1995.
7. Jörg Flum and Martin Grohe. Fixed-parameter tractability, definability, and model-checking. *SIAM J. Comput.*, 31(1):113–145, 2001.
8. Jörg Flum and Martin Grohe. The parameterized complexity of counting problems. *SIAM J. Comput.*, 33(4):892–922, 2004.
9. Jörg Flum and Martin Grohe. *Parameterized Complexity Theory.* Springer, 2006.
10. Martin Grohe. The complexity of homomorphism and constraint satisfaction problems seen from the other side. *J. ACM*, 54(1), 2007.
11. Richard E. Ladner. On the structure of polynomial time reducibility. *J. ACM*, 22(1):155–171, 1975.
12. Catherine McCartin. Parameterized counting problems. *Ann. Pure Appl. Logic*, 138(1-3):147–182, 2006.