

The challenges of unbounded treewidth in parameterised subgraph counting problems[☆]

Kitty Meeks^{*}

School of Mathematical Sciences, Queen Mary University of London, London E1 4NS, UK

ARTICLE INFO

Article history:

Received 18 August 2014

Received in revised form 8 May 2015

Accepted 21 June 2015

Available online 10 July 2015

Keywords:

Parameterised complexity

Counting complexity

Subgraph counting

FPTRAS

ABSTRACT

Parameterised subgraph counting problems are the most thoroughly studied topic in the theory of parameterised counting, and there has been significant recent progress in this area. Many of the existing tractability results for parameterised problems which involve finding or counting subgraphs with particular properties rely on bounding the treewidth of these subgraphs in some sense; here, we prove a number of hardness results for the situation in which this bounded treewidth condition does *not* hold, resulting in dichotomies for some special cases of the general subgraph counting problem. The paper also gives a thorough survey of known results on this subject and the methods used, as well as discussing the relationships both between multicolour and uncoloured versions of subgraph counting problems, and between exact counting, approximate counting and the corresponding decision problems.

© 2015 The Author. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The field of parameterised counting complexity, first introduced by Flum and Grohe in [24], has received considerable attention from the theoretical computer science community in recent years; by far the most thoroughly studied family of parameterised counting problems are so-called subgraph counting problems, informally those problems which can be phrased as follows:

Input: An n -vertex graph $G = (V, E)$, and $k \in \mathbb{N}$.

Parameter: k .

Question: How many (labelled) k -vertex subsets of V induce graphs with a given property?

It should be noted that, although this description refers to induced subgraphs, problems in which the goal is to count copies of a given graph H that are not necessarily induced can just as well be formulated in this manner: it suffices to define the property we are looking for to be “contains H as a subgraph”.

The aims of this paper are two-fold. Firstly, we give a thorough survey of existing results relating to these problems, the methods used, and the relationships between the complexities of different problem-variants. Secondly, we extend a method based on treewidth (previously used in the literature to demonstrate the intractability of various problems involving homomorphisms and coloured subgraphs) to give some general hardness results for exact and approximate subgraph counting problems, and also the related decision versions of these problems.

[☆] Research supported by EPSRC grant “Computational Counting” (EP/I011935/1).

^{*} Correspondence to: School of Mathematics and Statistics, University of Glasgow, 15 University Gardens, Glasgow G12 8QW, UK.

E-mail address: kitty.meeks@glasgow.ac.uk.

A formal model for subgraph counting problems was introduced in [31], and this is the model considered here. The full definition of this model, and its multicolour variant (in which the input graph is equipped with a vertex-colouring and we wish to count only those subgraphs in which each vertex has a distinct colour, as defined in [32]) is given in Section 2.3, along with a discussion of the relationship between this model and other specific problems.

Many problems falling within the scope of this model have previously been studied in the literature, and there are a number of results concerning the intractability of exact counting for specific problems (a detailed discussion of existing results is given in Section 3.1). Some general criteria which are sufficient to imply the intractability of a problem of this kind are also known (see, for example, [32]), and very recently a dichotomy result for an important subcase of the model was proved by Curticapean and Marx [16], in which nearly all problems are shown to be intractable from the point of view of parameterised complexity. The only tractable problems in this particular case are those for which, in order to determine whether a given k -vertex subset induces a graph with the desired property, it suffices to examine the edges incident with at most c vertices in this subset, where c is some fixed constant that does not depend on k ; it is clear that problems of this kind can be solved in polynomial time (moreover, the simple algorithm given in [16] can easily be extended to the more general family of problems considered in this paper).

The many intractability results for solving these parameterised counting problems exactly motivate the study of approximate counting in this setting. The first approximation algorithm for a parameterised counting problem was given by Arvind and Raman in [6], and their method has recently been extended to solve a larger family of subgraph counting problems [31]. Both of these positive approximation results apply to problems in which the minimal graphs that satisfy the property in question have bounded treewidth; a number of positive results for the decidability of subgraph problems also exploit treewidth in a similar manner.

In Section 4 we derive a number of new intractability results which apply in the case that minimal graphs which satisfy the property under consideration do not have treewidth bounded by any constant. Using a construction that exploits Robertson and Seymour's Excluded Grid Theorem [41],¹ we show that, in this particular setting:

1. exact counting is intractable,²
2. decision is intractable in the multicolour setting, and
3. again in the multicolour setting, approximate counting is intractable.

The last two intractability results, when combined with the existing algorithms, give rise to dichotomy results for restricted versions of the problems.

The remainder of the paper is organised as follows. In Section 2 we provide the necessary technical background, before giving a thorough analysis of existing results and methods in Section 3, as well as discussing the relationships between the various problem variants that have previously been studied. Section 4 then contains the proof of our new complexity results, and finally in Section 5 we review the limitations of the existing methods and suggest directions for future research.³

2. Preliminaries

In this section, we cover the key background material which will be required throughout the rest of the paper. We begin in Section 2.1 by summarising the main definitions used, and then review the key concepts we need from the theory of parameterised counting complexity in Section 2.2. In Section 2.3 we then give the formal definition of the general model we use for subgraph counting problems, and also discuss some variants and special cases of this general problem and their relationship to specific problems studied elsewhere.

2.1. Notation and definitions

In this section we introduce the key terms that will be used in the remainder of the paper.

Graphs, isomorphisms, subgraphs and minors

Throughout, all graphs are assumed to be simple, that is they do not have multiple edges or self-loops. The order of a graph G , written $|G|$, is the number of vertices in G . Given a graph $G = (V, E)$, and a subset $U \subseteq V$, we write $G[U]$ for the subgraph of G induced by the vertices of U . For any $k \in \mathbb{N}$, we write $[k]$ as shorthand for $\{1, \dots, k\}$, and denote by S_k the set of all permutations on $[k]$, that is, injective functions from $[k]$ to $[k]$. We write $V^{(k)}$ for the set of all subsets of V of size exactly k , and $V^{\underline{k}}$ for the set of k -tuples $(v_1, \dots, v_k) \in V^k$ such that v_1, \dots, v_k are all distinct.

Two graphs G and H are *isomorphic*, written $G \cong H$, if there exists a bijection $\theta : V(G) \rightarrow V(H)$ so that, for all $u, v \in V(G)$, we have $\theta(u)\theta(v) \in E(H)$ if and only if $uv \in E(G)$; θ is said to be an *isomorphism* from G to H . An *automorphism* on G is an

¹ This result has previously been used along the same lines in different settings by a number of authors; a discussion of the relationship of our methods to previous work is given in Section 3.2.5.

² This generalises a result of Curticapean and Marx [16], published since the writing of an initial version of this paper, which proves a special case of this result.

³ For a stand-alone survey, the reader may omit Section 4.

isomorphism from G to itself. We write $\text{aut}(G)$ for the number of automorphisms of G . A *subgraph* of G is a graph H which is isomorphic to a graph obtained from G by possibly deleting vertices and/or edges; a *spanning subgraph* is obtained by deleting edges only. H is a *proper subgraph* of G if H is a subgraph of G but is not isomorphic to G .

The notion of a subgraph is generalised by the definition of a *graph minor*. The graph H is said to be a *minor* of G if there is a mapping $m : V(H) \rightarrow \mathcal{P}(V(G))$, mapping vertices of H to sets of vertices from G , satisfying:

1. for all $u \in V(H)$, $G[m(u)]$ is connected,
2. for $u, v \in V(H)$ with $u \neq v$, $m(u) \cap m(v) = \emptyset$, and
3. for all $uv \in E(H)$, there exist $u' \in m(u)$ and $v' \in m(v)$ such that $u'v' \in E(G)$.

Such a mapping m is known as a *minor map* from H to G .

Coloured and labelled graphs

If G is coloured by some colouring $\omega : V \rightarrow [k]$ (not necessarily a proper colouring), we say that a subset $U \subseteq V$ is *colourful* (under ω) if, for every $i \in [k]$, there exists exactly one vertex $u \in U$ such that $\omega(u) = i$; note that this can only be achieved if $U \in V^{(k)}$.

We will be considering labelled graphs, where a labelled graph is a pair (H, π) such that H is a graph and $\pi : [|V(H)|] \rightarrow V(H)$ is a bijection. We denote by $\mathcal{L}(k)$ the set of all labelled subgraphs on k vertices.

For labelled graphs $(H, \pi), (H', \pi') \in \mathcal{L}(k)$, we say that (H, π) is a *labelled subgraph* of (H', π') if and only if, for every $uv \in E(H)$, we have $u'v' \in E(H')$, where $u' = \pi'(\pi^{-1}(u))$ and $v' = \pi'(\pi^{-1}(v))$.

Given a graph $G = (V, E)$ and a k -tuple of vertices (v_1, \dots, v_k) , we write $G[v_1, \dots, v_k]$ for the labelled graph (H, π) where $H = G[\{v_1, \dots, v_k\}]$ and $\pi(i) = v_i$ for each $i \in [k]$. For any set \mathcal{H} of labelled graphs with $(H, \pi) \in \mathcal{H}$, we set $\mathcal{H}^H = \{(H', \pi') \in \mathcal{H} : H' \cong H\}$.

Treewidth

A key concept in this paper is the *treewidth* of a graph. We say that (T, \mathcal{D}) is a *tree decomposition* of G if T is a tree and $\mathcal{D} = \{\mathcal{D}(t) : t \in V(T)\}$ is a collection of non-empty subsets of $V(G)$ (or *bags*), indexed by the nodes of T , satisfying:

1. $V(G) = \bigcup_{t \in V(T)} \mathcal{D}(t)$,
2. for every $e = uv \in E(G)$, there exists $t \in V(T)$ such that $u, v \in \mathcal{D}(t)$,
3. for every $v \in V(G)$, if $T(v)$ is defined to be the subgraph of T induced by nodes t with $v \in \mathcal{D}(t)$, then $T(v)$ is connected.

The *width* of the tree decomposition (T, \mathcal{D}) is defined to be $\max_{t \in V(T)} |\mathcal{D}(t)| - 1$, and the *treewidth* of G is the minimum width over all tree decompositions of G .

Embeddings

Given two graphs G and H , a *strong embedding* of H in G is an injective mapping $\theta : V(H) \rightarrow V(G)$ such that, for any $u, v \in V(H)$, $\theta(u)\theta(v) \in E(G)$ if and only if $uv \in E(H)$. We denote by $\#\text{StrEmb}(H, G)$ the number of strong embeddings of H in G . If \mathcal{H} is a class of labelled graphs on k vertices, we set

$$\begin{aligned} \#\text{StrEmb}(\mathcal{H}, G) = \\ |\{\theta : [k] \rightarrow V(G) : \theta \text{ is injective and } \exists (H, \pi) \in \mathcal{H} \text{ such that} \\ \theta(i)\theta(j) \in E(G) \iff \pi(i)\pi(j) \in E(H)\}|. \end{aligned}$$

If G is also equipped with a k -colouring ω , where $|V(H)| = k$, we write $\#\text{ColStrEmb}(H, G, \omega)$ for the number of strong embeddings of H in G such that the image of $V(H)$ is colourful under ω . Similarly, we set

$$\begin{aligned} \#\text{ColStrEmb}(\mathcal{H}, G, \omega) = \\ |\{\theta : [k] \rightarrow V(G) : \theta \text{ is injective, } \exists (H, \pi) \in \mathcal{H} \text{ such that} \\ \theta(i)\theta(j) \in E(G) \iff \pi(i)\pi(j) \in E(H), \\ \text{and } \theta([k]) \text{ is colourful under } \omega\}|. \end{aligned}$$

We can alternatively consider unlabelled embeddings of H in G . In this context we write $\#\text{SubInd}(H, G)$ for the number of subsets $U \in V(G)^{(|H|)}$ such that $G[U] \cong H$. Note that $\#\text{SubInd}(H, G) = \#\text{StrEmb}(H, G)/\text{aut}(H)$. If \mathcal{H} is a class of labelled graphs, we set

$$\#\text{SubInd}(\mathcal{H}, G) = |\{U \subseteq V(G) : \exists (H, \pi) \in \mathcal{H} \text{ such that } G[U] \cong H\}|.$$

Once again, we can also consider the case in which G is equipped with a k -colouring ω . In this case $\#\text{ColSubInd}(H, G, \omega)$ is the number of colourful subsets U such that $G[U] \cong H$, and

$$\begin{aligned} \#\text{ColSubInd}(\mathcal{H}, G, \omega) = |\{U \subseteq V(G) : \exists (H, \pi) \in \mathcal{H} \text{ such that } G[U] \cong H, \\ \text{and } U \text{ is colourful under } \omega\}|. \end{aligned}$$

Finally, we write $\#\text{Clique}_k(G)$ as shorthand for $\#\text{SubInd}(K_k, G)$, where K_k denotes a clique on k vertices.

2.2. Parameterised counting complexity

In this section, we introduce key notions from parameterised counting complexity, which we will use in the rest of the paper. Let Σ be a finite alphabet. A parameterised counting problem is a pair (Π, κ) , where $\Pi : \Sigma^* \rightarrow \mathbb{N}_0$ is a function and $\kappa : \Sigma^* \rightarrow \mathbb{N}$ is a parameterisation (a polynomial-time computable mapping).

Just as when considering the complexity of parameterised decision problems, an algorithm to solve a parameterised counting problem is considered to be efficient if, on input of size n with parameter k , its running time is bounded by $f(k)n^{O(1)}$, where f is any computable function of k . Problems (whether decision or counting problems) admitting such an *fpt-algorithm* belong to the class FPT and are said to be *fixed parameter tractable*.

To understand the intractability of parameterised counting problems, Flum and Grohe [24] introduce two kinds of reductions between such problems.

Definition. Let (Π, κ) and (Π', κ') be parameterised counting problems.

1. An fpt parsimonious reduction from (Π, κ) to (Π', κ') is an algorithm that computes, for every instance I of Π , an instance I' of Π' in time $f(\kappa(I)) \cdot |I|^c$ such that $\kappa'(I') \leq g(\kappa(I))$ and

$$\Pi(I) = \Pi'(I')$$

(for computable functions $f, g : \mathbb{N} \rightarrow \mathbb{N}$ and a constant $c \in \mathbb{N}$). In this case we write $(\Pi, \kappa) \leq_{\text{pars}}^{\text{fpt}} (\Pi', \kappa')$.

2. An fpt Turing reduction from (Π, κ) to (Π', κ') is an algorithm A with an oracle to Π' such that
 - (a) A computes Π ,
 - (b) A is an fpt-algorithm with respect to κ , and
 - (c) there is a computable function $g : \mathbb{N} \rightarrow \mathbb{N}$ such that for all oracle queries “ $\Pi'(y) = ?$ ” posed by A on input x we have $\kappa'(I') \leq g(\kappa(I))$.

In this case we write $(\Pi, \kappa) \leq_T^{\text{fpt}} (\Pi', \kappa')$.

Using these notions, Flum and Grohe introduce a hierarchy of parameterised counting complexity classes, $\#W[t]$, for $t \geq 1$; this is the analogue of the W -hierarchy for parameterised decision problems. In order to define this hierarchy, we need some more notions related to satisfiability problems.

The definition of levels of the hierarchy uses the following class of problems, where every first order formula ψ (with a free relation variable of arity s) gives rise to a specific problem in the class.

p-#WD $_{\psi}$

Input: A relational structure⁴ \mathcal{A} and $k \in \mathbb{N}$.

Parameter: k .

Question: How many relations $S \subseteq A^s$ of cardinality $|S| = k$ are such that $\mathcal{A} \models \psi(S)$ (where A is the universe of \mathcal{A})?

If Ψ is a class of first-order formulas, then **p-#WD- Ψ** is the class of all problems **p-#WD $_{\psi}$** where $\psi \in \Psi$. The classes of first-order formulas Σ_t and Π_t , for $t \geq 0$, are defined inductively. Both Σ_0 and Π_0 denote the class of quantifier-free formulas, while, for $t \geq 1$, Σ_t is the class of formulas

$$\exists x_1 \dots \exists x_i \psi,$$

where $\psi \in \Pi_{t-1}$, and Π_t is the class of formulas

$$\forall x_1 \dots \forall x_i \psi,$$

where $\psi \in \Sigma_{t-1}$. We are now ready to define the classes $\#W[t]$, for $t \geq 1$.

Definition ([24,25]). For $t \geq 1$, $\#W[t]$ is the class of all parameterised counting problems that are fpt parsimonious reducible to a problem in **p-#WD- Π_t** .

Unless $\text{FPT} = \text{W}[1]$, there does not exist an algorithm running in time $f(k)n^{O(1)}$ for any problem that is hard for the class $\#W[1]$ under either fpt parsimonious reductions or fpt Turing reductions. In the setting of this paper, a parameterised counting problem will be considered to be intractable if it is $\#W[1]$ -hard with respect to either form of reduction. A useful $\#W[1]$ -complete problem (shown to be hard in [24]) which we will use for reductions is **p-#CLIQUE**, defined formally as follows.

p-#CLIQUE

Input: A graph $G = (V, E)$, and $k \in \mathbb{N}$.

Parameter: k .

Question: How many k -vertex cliques are there in G ?

⁴ The relational structure \mathcal{A} of vocabulary τ of relations consists of the universe A together with an interpretation $R^{\mathcal{A}}$ of every relation R in τ . For further terminology we refer the reader to [25].

When considering approximation algorithms for parameterised counting problems, an “efficient” approximation scheme is an FPTRAS (fixed parameter tractable randomised approximation scheme), as introduced by Arvind and Raman [6]; this is the analogue in the parameterised setting of an FPRAS (fully polynomial randomised approximation scheme).

Definition. An FPTRAS for a parameterised counting problem Π with parameter k is a randomised approximation scheme that takes an instance I of Π (with $|I| = n$), and rational numbers $\epsilon > 0$ and $0 < \delta < 1$, and in time $f(k) \cdot g(n, 1/\epsilon, \log(1/\delta))$ (where f is any function, and g is a polynomial in $n, 1/\epsilon$ and $\log(1/\delta)$) outputs a rational number z such that

$$\mathbb{P}[(1 - \epsilon)\Pi(I) \leq z \leq (1 + \epsilon)\Pi(I)] \geq 1 - \delta.$$

Later in the paper, we will also consider the parameterised complexity of decision problems. This will involve the parameterised complexity class $W[1]$, the first level of the W -hierarchy (originally introduced by Downey and Fellows [18]), which is the decision equivalent of the $\#W$ -hierarchy. The levels of the W -hierarchy can be defined in a similar way to those of the $\#W$ -hierarchy, using the following problem.

WD_ψ

Input: A structure \mathcal{A} and $k \in \mathbb{N}$.

Parameter: k .

Question: Is there $S \subseteq A^s$ of cardinality $|S| = k$ such that $\mathcal{A} \models \psi(S)$ (where A is the universe of \mathcal{A})?

The classes $W[t]$, for $t \geq 1$, are then defined as follows.

Definition ([25]). For $t \geq 1$, $W[t]$ is the class of all parameterised problems that are fpt-reducible to $WD-\Pi_t$.

The definition of fpt-reductions for parameterised decision problems is analogous to that of fpt-parsimonious for counting problems given above; for the formal definition, and other notions from the theory of parameterised complexity, we refer the reader to [25]. When proving $W[1]$ -hardness of various problems later in the paper, we will use reductions from the problem \mathbf{p} -CLIQUE, shown to be $W[1]$ -complete in [18], and formally defined as follows:

\mathbf{p} -CLIQUE

Input: A graph $G = (V, E)$, and $k \in \mathbb{N}$.

Parameter: k .

Question: Does G contain a k -vertex clique?

We will also be considering the relationship between randomised and deterministic parameterised algorithms. Downey, Fellows and Regan [19] introduced a notion of randomised fpt reductions (rephrased here for consistency of terminology):

Definition. Let (Π, κ) and (Π', κ') be parameterised decision problems. A randomised (fpt, many-one) reduction from (Π, κ) to (Π', κ') is an algorithm that computes, for every instance I of Π , an instance I' of Π' in time $f(\kappa(I)) \cdot |I|^c$ such that $\kappa'(I') \leq g(\kappa(I))$ and

1. if I is a yes-instance for (Π, κ) then the probability that I' is a yes-instance for (Π', κ') is at least $\frac{1}{h(\kappa(I)) \cdot |I|^{c'}}$, and
 2. if I is a no-instance for (Π, κ) then I' is a no-instance for (Π', κ')
- (for computable functions $f, g, h : \mathbb{N} \rightarrow \mathbb{N}$ and constants $c, c' \in \mathbb{N}$).

We shall use this kind of reduction to analyse the relationship between algorithms for decision and approximate counting in Section 3.3.

2.3. The model

This paper is concerned with the formal model for subgraph counting problems introduced in [31]. We begin this section with the definition of the general problem and a discussion of some important special cases, as well as defining the multicolour version introduced in [32]. In Section 2.3.1 we then give a more detailed discussion of the relationship of this model to other problems that have previously been studied in the literature.

Let Φ be a family (ϕ_1, ϕ_2, \dots) of functions $\phi_k : \mathcal{L}(k) \rightarrow \{0, 1\}$, such that the function mapping $k \mapsto \phi_k$ is computable. For any k , we write \mathcal{H}_{ϕ_k} for the set $\{(H, \pi) \in \mathcal{L}(k) : \phi_k(H, \pi) = 1\}$, and set $\mathcal{H}_\Phi = \bigcup_{k \in \mathbb{N}} \mathcal{H}_{\phi_k}$. The general problem is then defined as follows.

\mathbf{p} -#INDUCED SUBGRAPH WITH PROPERTY(Φ) (\mathbf{p} -#ISWP(Φ))

Input: A graph $G = (V, E)$ and an integer k .

Parameter: k .

Question: What is $\#\text{StrEmb}(\mathcal{H}_{\phi_k}, G)$, that is, the cardinality of the set $\{(v_1, \dots, v_k) \in V^k : \phi_k(G[v_1, \dots, v_k]) = 1\}$?

Observe that this problem can equivalently regard this problem as that of counting induced labelled k -vertex subgraphs that belong to \mathcal{H}_Φ .

It was argued in [31] that this problem lies in $\#W[1]$:

Proposition 2.1. For any Φ , the problem $\mathbf{p}\text{-}\#\text{ISWP}(\Phi)$ belongs to $\#\text{W}[1]$.

Much of this paper will be concerned with the multicolour version of the general problem, introduced in [32] (and also known to belong to $\#\text{W}[1]$).

$\mathbf{p}\text{-}\#\text{MULTICOLOUR INDUCED SUBGRAPH WITH PROPERTY}(\Phi)$ ($\mathbf{p}\text{-}\#\text{MISWP}(\Phi)$)

Input: A graph $G = (V, E)$, an integer k and colouring $f : V \rightarrow [k]$.

Parameter: k .

Question: What is $\#\text{ColStrEmb}(\mathcal{H}_{\phi_k}, G, f)$, that is, the cardinality of the set $\{(v_1, \dots, v_k) \in V^k : \phi_k(G[v_1, \dots, v_k]) = 1 \text{ and } \{f(v_1), \dots, f(v_k)\} = [k]\}$?

Multicolour versions of specific subgraph counting problems were initially considered as a technical tool for deriving results about the uncoloured version (as in, for example, [6,32]); however, as we will see in Section 3.3 the multicolour versions turn out to have interesting properties when considered as separate problems in their own right.

In this paper we will also consider the decision versions of both the multicolour and uncoloured variants of the problem, which are defined as follows.

$\mathbf{p}\text{-INDUCED SUBGRAPH WITH PROPERTY}(\Phi)$ ($\mathbf{p}\text{-ISWP}(\Phi)$)

Input: A graph $G = (V, E)$, and an integer k .

Parameter: k .

Question: Is there any tuple $(v_1, \dots, v_k) \in V^k$ such that $\phi_k(G[v_1, \dots, v_k]) = 1$?

$\mathbf{p}\text{-MULTICOLOUR INDUCED SUBGRAPH WITH PROPERTY}(\Phi)$ ($\mathbf{p}\text{-MISWP}(\Phi)$)

Input: A graph $G = (V, E)$, an integer k and colouring $f : V \rightarrow [k]$.

Parameter: k .

Question: Is there any tuple $(v_1, \dots, v_k) \in V^k$ such that $\phi_k(G[v_1, \dots, v_k]) = 1$ and $\{f(v_1), \dots, f(v_k)\} = [k]$?

It is straightforward to verify that both these problems belong to the parameterised complexity class $\text{W}[1]$.

We now consider a number of special families of properties Φ which are of particular interest.

Symmetric properties

In Section 2.3.1, we will discuss the necessity of considering labelled graphs in order to represent a number of well-studied problems from the literature in this form. However, an interesting special case arises when we only consider “unlabelled” graph problems (such as $\mathbf{p}\text{-}\#\text{CLIQUE}$, or $\mathbf{p}\text{-}\#\text{CONNECTED INDUCED SUBGRAPH}$).

We say that the property Φ is *symmetric* if the value of $\phi_k(H, \pi)$ depends only on the graph H and not on the labelling π of the vertices. Thus, unlabelled graph problems correspond precisely to problems $\mathbf{p}\text{-}\#\text{ISWP}(\Phi)$ with Φ symmetric. The special case of the general problem for symmetric properties was also introduced in [31]:

$\mathbf{p}\text{-}\#\text{INDUCED UNLABELLED SUBGRAPH WITH PROPERTY}(\Phi)$

Input: A graph $G = (V, E)$ and $k \in \mathbb{N}$.

Parameter: k .

Question: What is $\#\text{SubInd}(\mathcal{H}_{\phi_k})$, that is, the cardinality of the set $\{(v_1, \dots, v_k) \in V^{(k)} : \phi_k(G[v_1, \dots, v_k]) = 1\}$?

Observe that, for any symmetric property Φ , the output of $\mathbf{p}\text{-}\#\text{ISWP}(\Phi)$ is exactly $k!$ times the output of $\mathbf{p}\text{-}\#\text{INDUCED UNLABELLED SUBGRAPH WITH PROPERTY}(\Phi)$, for any graph G and $k \in \mathbb{N}$.

Monotone and uniformly monotone properties

Much of the rest of the paper will be concerned with *monotone* properties. We say that Φ is monotone if, whenever $\phi_k(H, \pi) = 1$ and (H, π) is a labelled spanning subgraph of (H', π') , we also have $\phi_k(H', \pi') = 1$. It follows immediately from this definition that, if Φ is a monotone property and G is a subgraph of G' , then

$$\#\text{StrEmb}(\mathcal{H}_{\Phi}, G) \leq \#\text{StrEmb}(\mathcal{H}_{\Phi}, G').$$

Note that any monotone property is uniquely defined by the set of minimal labelled graphs (with respect to inclusion as labelled subgraphs) which satisfy the property, for each k .

For any property Φ (not necessarily monotone), we denote by $\min(\phi_k)$ the set of minimal elements satisfying ϕ_k for each $k \in \mathbb{N}$, and set $\min(\Phi) = \bigcup_{k \in \mathbb{N}} \min(\phi_k)$; $\min(\Phi)$ then uniquely specifies any *monotone* property, but there can be many distinct pairs of non-monotone properties $\Phi \neq \Phi'$ such that $\min(\Phi) = \min(\Phi')$.

The best studied example of a monotone property is $\mathbf{p}\text{-}\#\text{SUB}(\mathcal{C})$: this corresponds to the special case in which $|\min(\phi_k)| \leq 1$ for each k . The problem $\mathbf{p}\text{-}\#\text{SUB}(\mathcal{C})$ is defined for any class of graphs $\mathcal{C} = \{H_k : k \in I_{\mathcal{C}} \subseteq \mathbb{N}\}$, where each H_k has k vertices (so, for each $k \in \mathbb{N}$, there is at most one element of \mathcal{C} which has k vertices, and the indexing set $I_{\mathcal{C}} \subseteq \mathbb{N}$ contains every k such that some element of \mathcal{C} has exactly k vertices):

$\mathbf{p}\text{-}\#\text{SUB}(\mathcal{C})$

Input: A graph $G = (V, E)$ and $H \in \mathcal{C}$.




Fig. 1. Labeled subgraphs satisfying ϕ_3 , where Φ is not a uniformly monotone property: note that (H_4, π_4) belongs to $\min(\phi_k)$ (it is minimal as a labelled subgraph) but not $\min^*(\phi_k)$ (it is not minimal as an unlabelled graph).

Parameter: $k = |H|$.

Question: How many copies (not necessarily induced) of H are there in G ?

This paper will in fact focus primarily on a sub-family of monotone properties, which satisfy an additional restriction. For any property Φ , let us write $\mathcal{H}_{\phi_k}^*$ for the set of unlabelled graphs which satisfy ϕ_k when equipped with an appropriate labelling. We then set $\min^*(\phi_k)$ to be the set of minimal elements of $\mathcal{H}_{\phi_k}^*$ (with respect to inclusion as unlabelled subgraphs), and write $\min^*(\Phi) = \bigcup_{k \in \mathbb{N}} \min^*(\phi_k)$. Note that, as illustrated in Fig. 1, it is possible to define a property Φ such that, for some $(H, \pi) \in \min(\phi_k)$, $H \notin \min^*(\phi_k)$. However, all the examples of monotone properties which have previously been studied in the literature (including **p**-#SUB(\mathcal{C}) and **p**-#CONNECTED INDUCED SUBGRAPH) have the property that all minimal labelled subgraphs satisfying Φ are also minimal when considered as unlabelled subgraphs (that is, $\min^*(\phi_k) = \{H : \exists \pi \text{ such that } (H, \pi) \in \min(\phi_k)\}$ for each k): indeed, it is hard to think of a natural subgraph counting problem that does not satisfy this condition. We call monotone properties that behave in this way, with minimal labelled and unlabelled subgraphs satisfying the property coinciding, *uniformly monotone* properties. Our dichotomy results in Section 4 apply to uniformly monotone properties; the key observation we exploit here is that, for a uniformly monotone property Φ , the collection of labelled graphs $\min(\Phi)$ has bounded treewidth if and only if the collection of unlabelled graphs $\min^*(\Phi)$ has bounded treewidth.

2.3.1. Relationship to other problems, and applications

In Section 3.1, we will see many examples of specific subgraph counting problems that have previously been studied in the literature. All of these specific problems can be formulated in terms of the model described above (up to, possibly, division by a constant factor, easily computable in time depending only on the parameter); in this section we describe how to do this for two specific examples (as previously discussed in [32]).

Our first and simplest example is **p**-#INDUCED SUBGRAPH ISOMORPHISM(\mathcal{C}), the problem of counting induced subgraphs of a graph G that are isomorphic to some graph $C \in \mathcal{C}$ (where \mathcal{C} is any recursively enumerable class of unlabelled graphs; to simplify the formulation, we shall assume that \mathcal{C} contains at most one graph on k vertices for each $k \in \mathbb{N}$). We then set

$$\phi_k(H, \pi) = \begin{cases} 1 & \text{if } H \in \mathcal{C} \\ 0 & \text{otherwise.} \end{cases}$$

Then, for any $\{v_1, \dots, v_k\} \in V^{(k)}$ such that $G[\{v_1, \dots, v_k\}] \cong C \in \mathcal{C}$, we will have $\phi_k(G[v_{\sigma(1)}, \dots, v_{\sigma(k)}]) = 1$ for every permutation $\sigma \in S_k$. Thus, the output of **p**-#ISWP(Φ) on the input (G, k) is equal to $k!$ times the desired value (the number of k -vertex induced subgraphs in G that belong to \mathcal{C}).

Our second example is **p**-#SUB(\mathcal{C}); unlike the previous example, in this case we need to exploit the fact that the model allows us to count *labelled* graphs. We begin with a concrete special case of **p**-#SUB(\mathcal{C}), namely **p**-#MATCHING (where \mathcal{C} is the set of all graphs consisting of a collection of disjoint edges). Here, I_e is the set of all even natural numbers and, for each $k \in I_e$, H_k is the graph consisting of $k/2$ disjoint edges. We can then set

$$\phi_k(H, \pi) = \begin{cases} 1 & \text{if } k \text{ is even and, for } 1 \leq i \leq k/2, \pi(2i - 1)\pi(2i) \in E(H) \\ 0 & \text{otherwise,} \end{cases}$$

and the output of **p**-#ISWP(Φ) on the input (G, k) will be equal to $(k/2)! \cdot 2^{k/2}$ times the output of **p**-#MATCHING (since $(k/2)! \cdot 2^{k/2}$ is the number of automorphisms of a $k/2$ -edge matching: there are $(k/2)!$ ways to map a set of $k/2$ edges to itself, and each edge can be mapped to any given edge in two different ways).

More generally, to count copies of graphs from \mathcal{C} we fix, for each $k \in I_e$, a labelling $\pi_k : [k] \rightarrow V(H_k)$, and set

$$\phi_k(H, \pi) = \begin{cases} 1 & \text{if } k \in I_e \text{ and, for every } uv \in E(H_k), \text{ we have } (\pi \circ \pi_k^{-1}(u))(\pi \circ \pi_k^{-1}(v)) \in E(H) \\ 0 & \text{otherwise.} \end{cases}$$

The output of **p**-#ISWP(Φ) on the input (G, k) is then equal to $\text{aut}(H_k)$ times the output of **p**-#SUB(\mathcal{C}).

Note that, if we were only able to count unlabelled subgraphs, we could not give different weight to induced unlabelled subgraphs containing different numbers of distinct copies of graphs from \mathcal{H} , and could only define properties corresponding to induced k -vertex subgraphs that contain at least r copies of H_k . For example, we could express the problem of counting the number of induced k -vertex subgraphs that contain at least one perfect matching using only unlabelled subgraphs, but to translate **p**-#MATCHING into a general framework of this kind we need to use labels.

In [16], the authors considered as a technical tool a related problem, \mathbf{p} -#PARTITIONEDSUB(\mathcal{C}), in which G and H_k are equipped with vertex-colourings $f_G : V(G) \rightarrow \{1, \dots, k\}$ and $f_H : V(H_k) \rightarrow \{1, \dots, k\}$ respectively, and the goal is to count colourful tuples $(v_1, \dots, v_k) \in V(G)^k$ which are *colour-preserving isomorphic* to H with colouring f_H (that is, $(f_G|_{\{v_1, \dots, v_k\}})^{-1} \circ f_H$ defines an isomorphism from H_k to $G[\{v_1, \dots, v_k\}]$). This problem is superficially similar to the multi-colour problem \mathbf{p} -#MISWP(Φ) defined above, and indeed for the special case of \mathbf{p} -#PARTITIONEDSUB(\mathcal{C}) they are equivalent with respect to fpt-Turing reductions; however, this relationship does not extend to the general problem.

Results concerning this very general model can be used in any application that involves counting small substructures in a large graph. Strategies that involve counting small substructures in large networks have previously been used for network security tools [27,45,46], and the specific example of counting the number of triangles for the analysis of social networks has been studied extensively [7,36,43,48]. The counting of more general “network motifs” has been used to analyse biological networks [38], and counting versions of the GRAPH MOTIF problem (which relates to metabolic networks) have recently been studied in the framework of parameterised subgraph counting [30,31].

3. The complexity landscape

We are concerned with determining the complexity of problems that can be expressed in the framework described in Section 2.3, and there are three related questions we can ask about the complexities of both the uncoloured and multicolour versions (giving a total of six problem variants to consider):

1. Is the decision version in FPT?
2. Is there an FPTRAS for the counting version?
3. Is the exact counting problem in FPT?

For a concrete example, consider the problem \mathbf{p} -#CONNECTED INDUCED SUBGRAPH, studied in [31]. In this case:

- exact counting is #W[1]-complete, in both the uncoloured and multicolour versions [31, Theorem 2.1],
- there is an FPTRAS for the problem, in both the uncoloured and multicolour versions [31, Corollary 3.2], and
- the decision problem, in both the uncoloured and multicoloured case, is in FPT (see Corollary 3.11); in fact the uncoloured version can be solved in polynomial time by a component exploration method.

In Section 3.1, we summarise existing results about the complexity of all six problem variants, for specific properties. In Section 3.2 we then discuss the different methods that have commonly been used to address these questions in the literature, before considering in Section 3.3 the relationships between the complexities of the different problem variants.

3.1. Existing results

The study of the complexity of counting subgraphs has a long history. Previous approaches have included the development of algorithms to solve the problem when certain restrictions are placed on either the input graph G (such as its arboricity [14]) or the graph H we wish to find (such as its flower number [47]). Exact algorithms for the problems of counting paths [9] and cycles [4] have been studied, and recently Williams and Williams [51] considered exact algorithms for the more general problem \mathbf{p} -#SUB(\mathcal{C}). In [5], Amini, Fomin and Saurabh explore the relationship between counting homomorphisms and isomorphisms to give exact algorithms for various cases of \mathbf{p} -#SUB(\mathcal{C}); in Section 3.2.5 we shall discuss in more detail the relationship between our results and previous methods used to determine the complexity of homomorphism problems.

In the seminal paper on parameterised counting complexity [24], Flum and Grohe demonstrated that the problems \mathbf{p} -#PATH and \mathbf{p} -#CYCLE (the problems of counting k -vertex paths and cycles respectively, both special cases of \mathbf{p} -#SUB(\mathcal{C})) are #W[1]-complete, in contrast to the fact that the decision versions of both of these problems belong to FPT (see Theorem 3.10); #W[1]-completeness of the problem of counting k -cliques (\mathbf{p} -#CLIQUE) was also proved in [24]. Further progress on the complexity of special cases of \mathbf{p} -#SUB(\mathcal{C}) did not come for nearly a decade, until Curticapean [15] resolved an open question from [24], proving the #W[1]-completeness of \mathbf{p} -#MATCHING (based on earlier joint work with Bläser [10] on the complexity of the weighted version of this problem), another problem whose decision version belongs to FPT (again by Theorem 3.10). Even more recently, Curticapean and Marx gave a simpler proof of this result, showing that it remains true even if input graph G is bipartite, as part of their more general dichotomy result for \mathbf{p} -#SUB(\mathcal{C}): their main result shows that \mathbf{p} -#SUB(\mathcal{C}) is in FPT (and is indeed solvable in polynomial time) if all graphs in \mathcal{C} admit a vertex cover of size at most c (where c is some fixed constant, independent of the number of vertices in the graph), and otherwise this problem is #W[1]-complete.

All of the results in the previous paragraph relate to monotone (indeed, uniformly monotone) properties. Previously, Arvind and Raman [6] had demonstrated the existence of an FPTRAS for the problem \mathbf{p} -#SUB(\mathcal{C}) whenever \mathcal{C} is a class of graphs of bounded treewidth (implying that, among others, \mathbf{p} -#PATH, \mathbf{p} -#CYCLE and \mathbf{p} -#MATCHING are efficiently approximable). This result was recently extended by Jerrum and Meeks [31] to give an FPTRAS for \mathbf{p} -#ISWP(Φ) whenever Φ is a monotone property and $\min(\Phi)$ is a class of graphs having bounded treewidth. This latter result includes problems such as that of counting the number of connected k -vertex induced subgraphs, \mathbf{p} -#CONNECTED INDUCED SUBGRAPH, although solving this problem exactly was also shown in [31] to be #W[1]-complete.

Properties that are not monotone have also been studied in the literature. Chen, Thurley and Weyer demonstrated that $\mathbf{p}\text{-}\#\text{INDUCED SUBGRAPH ISOMORPHISM}(\mathcal{C})$ is $\#\text{W}[1]$ -complete whenever \mathcal{C} contains arbitrarily large graphs; if this condition does not hold, then the problem is clearly solvable in polynomial time. They further demonstrated that the same criterion distinguishes between the polynomial-time solvable and $\text{W}[1]$ -complete cases of the corresponding decision problem. Some general $\#\text{W}[1]$ -hardness results for $\mathbf{p}\text{-}\#\text{ISWP}(\Phi)$ in the case that Φ is not monotone were also given by Jerrum and Meeks in [32], and very recently the same authors showed that $\mathbf{p}\text{-}\#\text{EVEN SUBGRAPH}$ (the problem of counting induced k -vertex subgraphs with an even number of edges) is $\#\text{W}[1]$ -complete, although the problem does admit an FPTRAS, and the decision version is polynomial-time solvable.

3.2. Methods

In this section we discuss some of the methods that have previously been used in the literature to determine the complexity of subgraph counting problems (and the related decision problems), focusing on those that have been used to address a number of different problems.

3.2.1. Interpolation and matrix inversion

When giving reductions to exact counting problems, certain additional techniques become available that are not applicable when considering decision or approximate counting problems; probably the most thoroughly exploited technique of this kind is polynomial interpolation, which has been widely used to prove the $\#\text{P}$ -completeness of counting problems in the setting of classical counting complexity (see, for example, [20,49,50,52]). In essence, the reduction method involves defining a polynomial whose coefficients are in some way related to the solution of a known computationally hard problem, and which we can evaluate at different points by making suitable calls to an oracle for the problem under consideration; this makes it possible to recover the coefficients of the polynomial in question, and thus to solve the known hard problem. The final step in such a reduction typically involves demonstrating that the resulting system of linear equations has a unique solution, or equivalently proving that a matrix is non-singular.

Techniques along these lines have also been used to obtain a number of results about the $\#\text{W}[1]$ -completeness of parameterised counting problems, including $\mathbf{p}\text{-}\#\text{WEIGHTEDMATCHING}$ [10] and $\mathbf{p}\text{-}\#\text{MATCHING}$ [16]. In other parameterised subgraph counting situations, reductions have involved demonstrating the invertibility of a matrix that does not arise from polynomial interpolation, as for example in proving the hardness of $\mathbf{p}\text{-}\#\text{CYCLE}$ [24], $\mathbf{p}\text{-}\#\text{CONNECTED INDUCED SUBGRAPH}$ [31] and $\mathbf{p}\text{-}\#\text{EVEN SUBGRAPH}$ [33]. The strategies used to demonstrate that the resulting matrix is non-singular vary greatly in style and difficulty, and have drawn on diverse areas of mathematics including non-commutative algebra [16] and lattice theory [31,33].

3.2.2. Ramsey theory

Ideas from Ramsey Theory have played a starring role in many results related to the parameterised complexity of counting subgraphs; here we discuss their use in proving $\#\text{W}[1]$ -completeness, but we will see in Section 3.2.3 that these ideas have also been used to give positive results for decision and approximate counting.

Many of the hardness results exploit the simplest form of Ramsey's theorem; in these applications it is typically only required that the Ramsey number $R(k)$ is a computable function of k , and so the following bound, which follows immediately from a result of Erdős and Szekeres [21], suffices:

Theorem 3.1. *Let $k \in \mathbb{N}$. Then there exists $R(k) < 2^{2k}$ such that any graph on $n \geq R(k)$ vertices contains either a clique or independent set on k vertices.*

The strategy in most of the reductions using Ramsey theory is to show that graphs satisfying the property Φ , which we can presumably count using our oracle, must (for a large enough value of the parameter) contain arbitrarily large copies of substructures from some class (such as cliques) that we know is hard to count; with the construction of appropriate gadgets, the oracle can then be used to count (for example) cliques in the original graph. This strategy was first exploited in the parameterised counting setting by Chen, Thurley and Weyer [13] to address the problem $\mathbf{p}\text{-}\#\text{INDUCED SUBGRAPH ISOMORPHISM}(\mathcal{C})$. For more complicated properties, which are satisfied by more than one k -vertex graph for any k (as, for example, in [32]), much of the work in such reductions involves constructing gadgets that ensure there is no “confusion” between different subgraphs that satisfy our property, so that there is indeed a correspondence between subgraphs satisfying Φ (perhaps also with certain additional properties) and the structures we wish to count in the original graph; this seems to be more easily achievable when the properties in question are in some sense “sparse” (as in [13,32]).

Such strategies sometimes give a more natural reduction to the multicolour version of the problem, which implies the hardness of the uncoloured version in the exact counting setting (as explained in Section 3.3); reductions relying on the multicolour version as an intermediate step do not then transfer to the settings of decision or approximate counting.

Ramsey-based methods have been applied with particular success to subgraph counting problems involving hereditary classes of graphs (that is, classes that are closed under taking induced subgraphs). Khot and Raman [35] investigated the

parameterised complexity of the decision problem \mathbf{p} -ISWP($\Phi_{\mathcal{C}}$), where \mathcal{C} is any hereditary graph class and $\phi_k(H, \pi) = 1$ if and only if $H \in \mathcal{C}$ with $|H| = k$. They proved the following result:

Theorem 3.2 ([35]). *If the hereditary class Π contains all empty graphs and all complete graphs, or excludes some empty graphs and some complete graphs, then \mathbf{p} -ISWP (Φ_{Π}) is in FPT; otherwise, the problem is W[1]-complete.*

Both parts of this result were proved using ideas from Ramsey theory. Very recently, Curticapean and Marx [16] also used Ramsey-theoretic ideas to prove their dichotomy for \mathbf{p} -#SUB(\mathcal{C}) in the special case that \mathcal{C} is a hereditary class, proving in the process the Ramsey-style result that any graph containing a sufficiently large matching must contain either a clique on k vertices, a complete bipartite graph on $k + k$ vertices, or an induced matching with k edges.

In some cases, an extension of the basic Ramsey theorem is required. Indeed, the argument in [32] makes use of a straightforward corollary of Theorem 3.1, which gives a lower bound on the number of k -vertex subsets in a large graph which must induce either a clique or an independent set.

Corollary 3.3 ([32]). *Let $G = (V, E)$ be an n -vertex graph, where $n \geq 2^{2k}$. Then the number of k -vertex subsets $U \subset V$ such that U induces either a clique or independent set in G is at least*

$$\frac{(2^{2k} - k)!}{(2^{2k})!} \frac{n!}{(n - k)!}.$$

3.2.3. Random sampling

Perhaps the simplest way of devising an FPTRAS for a subgraph counting problem is to use a random sampling technique: the idea is to estimate the total number of k -tuples of vertices that induce graphs satisfying ϕ_k by picking some number of such tuples uniformly at random and counting the number of selected tuples that are “good” (i.e. satisfy ϕ_k). However, for such a method to work, we need to know that the total number of “good” k -tuples in the input graph satisfies certain conditions. For example, suppose that G contains only one good k -tuple: in order to have a reasonable chance of finding this tuple (and thus distinguishing between this situation and the case in which G contains no good k -tuple) we would have to sample $O(n^k)$ tuples, requiring more time than is allowed by the definition of the FPTRAS.

We now give a general sufficient condition on the number of good k -tuples in an underlying graph which will ensure that we can obtain a good approximation with a random sampling method. The techniques are considered folklore in the classical random counting community, and so only a sketch proof is given.

Lemma 3.4. *Let $G = (V, E)$ be a graph on n vertices and ϕ_k a mapping from labelled k -vertex graphs to $\{0, 1\}$, and set N to be the number of k -tuples of vertices $(v_1, \dots, v_k) \in V^k$ satisfying $\phi_k(G[v_1, \dots, v_k]) = 1$. Suppose that there exists a polynomial $q(n)$ and a computable function $g(k)$ such that either $N = 0$ or $N \geq \frac{1}{g(k)q(n)} \frac{n!}{(n-k)!}$. Then, for every $\epsilon > 0$ and $\delta \in (0, 1)$ there is an explicit randomised algorithm which outputs an integer α , such that*

$$\mathbb{P}[|\alpha - N| \leq \epsilon \cdot N] \geq 1 - \delta,$$

and runs in time at most $g(k)\tilde{q}(n, \epsilon^{-1}, \log(\delta^{-1}))$, where g is a computable function and \tilde{q} is a polynomial.

Proof (Sketch). We obtain an approximation to N using a simple random sampling algorithm. At each step, a k -tuple (v_1, \dots, v_k) of vertices is chosen uniformly at random among all elements of V^k ; we then determine whether $\phi_k(G[v_1, \dots, v_k]) = 1$. This sampling and checking step can clearly be performed in polynomial time. To obtain a good estimate of the total number of k -tuples satisfying ϕ_k , we repeat this sampling process t times (for some value of t to be determined), and compute the proportion p of our sampled tuples which satisfy ϕ_k . We then output as our approximation $p \frac{n!}{(n-k)!}$. Note that in the case that $N = 0$ we are certain to output 0, as required.

The value of t must be chosen to be large enough that $\mathbb{P}[|p \frac{n!}{(n-k)!} - N| \leq \epsilon \cdot N] \geq 1 - \delta$. However, it is straightforward to verify (using, for example, a Chernoff bound) that the number of trials required is not too large (for example, $4 \log(\delta^{-1})g(k)q(n)\epsilon^{-2}$ will do). \square

In many cases, random sampling methods give rise to an FPTRAS for properties Φ for which the decision version is trivial (at least for sufficiently large input graphs G). As an example, consider the following (symmetric) problem:

\mathbf{p} -#CLIQUE OR INDEPENDENT SET

Input: A graph $G = (V, E)$ and $k \in \mathbb{N}$.

Parameter: k .

Question: How many k -vertex subsets of G induce either a clique or an independent set?

It was observed by Arvind and Raman [6] that the decision version of this problem is trivially in FPT: for G large enough, the answer is always YES by Theorem 3.1 (and if the number of vertices in G is less than 2^{2k} , the problem can be solved by brute force search). The approximability of the corresponding counting problem also follows easily from Lemma 3.4 together with Corollary 3.3:

Corollary 3.5. *There exists an FPTRAS for \mathbf{p} -#CLIQUE OR INDEPENDENT SET.*

Similarly, the proof of the positive direction of [Theorem 3.2](#) given in [35] invoked Ramsey theory to prove that the answer to the decision problem is trivially YES when the input graph is sufficiently large; thus, exploiting [Lemma 3.4](#) again, we easily obtain the following result.

Corollary 3.6. *If the hereditary class Π contains all empty graphs and all complete graphs, or excludes some empty graphs and some complete graphs, then \mathbf{p} -#ISWP(Φ_Π) admits an FPTRAS.*

In both of these examples, it was known that any sufficiently large graph must contain at least one good k -tuple; however, the formulation of [Lemma 3.4](#) allows for a slightly more complicated situation, in which there is also the possibility of the input graph containing no good k -tuple. This is the case when considering the problem \mathbf{p} -#EVEN SUBGRAPH in [33]: Ramsey theory is exploited to show that if there is at least one induced k -vertex subgraph with an even number of edges then there must in fact be a large number of such subgraphs. In this situation the decision problem is not entirely trivial, but nevertheless admits an fpt-algorithm, as it is shown that (for a sufficiently large graph G), if there is no k -vertex induced subgraph with an even number of edges, then G must have a very specific and easily recognisable structure.

These methods are typically more successful in devising approximation algorithms for uncoloured problems than for the multicolour version, primarily because it seems to be much more difficult to guarantee the existence of (a large number of) specific structure(s) when an additional condition on the colouring of the structures is also imposed. As an example, we consider the multicolour version of the problem \mathbf{p} -CLIQUE OR INDEPENDENT SET discussed above.

Proposition 3.7. *\mathbf{p} -MULTICOLOUR CLIQUE OR INDEPENDENT SET is $W[1]$ -complete.*

Proof. We proceed by means of a reduction from \mathbf{p} -MULTICOLOUR CLIQUE, shown to be $W[1]$ -complete in [23]. Let G with colouring f be the input to an instance of \mathbf{p} -MULTICOLOUR CLIQUE; without loss of generality we may assume that f colours the vertices of G with colours $[k]$. We now construct G' from G by adding a new vertex v adjacent to every vertex of G ; we extend the colouring f , to give a colouring f' of the vertices of G' , by setting $f'(v) = k + 1$. Any colourful subset of vertices in G' with colouring f' must therefore contain v , but as v is adjacent to all other vertices it cannot be part of any $k + 1$ -independent set. Thus, there is a colourful clique or independent set in G' (with colouring f') if and only if there is a colourful clique in G (with colouring f). \square

3.2.4. Colour-coding

An extremely important technique in problems that involve finding or counting small subgraphs is that of *color-coding*, first introduced by Alon, Yuster and Zwick [3]. This technique allows us to transform algorithms for the multicolour versions of the problems into algorithms for the uncoloured case, in the setting of decision or approximate counting.

The simplest form of this technique gives randomised algorithms. If the vertices of G are coloured independently and uniformly at random with k colours, the probability that any given k -vertex subset is colourful with respect to this random colouring depends only on k . Thus, by repeating this random colouring process sufficiently many times (which will be bounded by a function of k), the probability that a given k -vertex subset is colourful with respect to *at least one* of the random colourings can be made arbitrarily high. Thus, if the multicolour decision problem is in FPT, we can apply this known algorithm to each of the randomly coloured graphs in turn, and return the answer YES to the uncoloured decision problem if and only if we find a multicolour yes-instance with respect to at least one of the colourings.

In fact, this process can be derandomised, using *k -perfect families of hash functions*: these are families of functions $\mathcal{F}_{n,k}$ from $[n]$ to $[k]$ such that, for every $U \in [n]^{(k)}$, there exists some $f \in \mathcal{F}_{n,k}$ such that U is colourful with respect to f . Bounds on the size of such families and the time required to construct them are known:

Theorem 3.8 ([3], *Refining Constructions in [44,42,26]*). *For all $n, k \in \mathbb{N}$, there is a k -perfect family $\mathcal{F}_{n,k}$ of hash functions from $[n]$ to $[k]$ of cardinality $2^{O(k)} \cdot \log n$, and (given n and k) the family $\mathcal{F}_{n,k}$ can be computed in time $2^{O(k)} \cdot n \cdot \log^2 n$.*

Thus if, instead of using random colourings, we begin by constructing a k -perfect family of hash functions and then apply an FPT algorithm for the multicolour version with each of these colourings, we obtain a deterministic FPT algorithm for the uncoloured problem. This observation gives the following result (implicit in [3]):

Proposition 3.9. *For any family Φ , if \mathbf{p} -MISWP(Φ) is in FPT then so is \mathbf{p} -ISWP(Φ).*

This technique was used to prove the following result in [3]; inclusion in FPT had previously been proved by Plehn and Voigt [40] using different techniques, but the color-coding method gives an algorithm with improved running time.

Theorem 3.10. *Let \mathcal{C} be a class of graphs having treewidth at most t . Then \mathbf{p} -SUB(\mathcal{C}) can be solved in time $2^{O(k)} |G|^{\ell+1} \log |G|$.*

This result covers, for example, the decision problems \mathbf{p} -PATH, \mathbf{p} -CYCLE and \mathbf{p} -MATCHING. As an intermediate step in the proof of [Theorem 3.10](#), it is also shown that the problem of deciding whether there exists an *colourful* copy of a graph from \mathcal{C} is also in FPT; by applying this result for the multicolour case to each element of $\min(\Phi)$ in turn, we obtain the following more general corollary.

Corollary 3.11. *Let Φ be a monotone property such that $\min(\Phi)$ is a class of graphs of bounded treewidth. Then \mathbf{p} -ISWP(Φ) and \mathbf{p} -MISWP(Φ) are in FPT.*

When considering approximate counting rather than decision, a more careful probability analysis reveals that the same strategy of random colourings (combined with repeated calls to an appropriate approximation algorithm for the multicolour version of the problem) can be used to give an FPTRAS for the uncoloured counting problem; Arvind and Raman exploited this technique in [6], implicitly using the following result.

Proposition 3.12. *For any family Φ , if \mathbf{p} -#MISWP(Φ) admits an FPTRAS, then so does \mathbf{p} -#ISWP(Φ).*

Indeed, Arvind and Raman provided in [6] a general framework for proving the existence of an FPTRAS for a parameterised counting problem, translating the classic Karp–Luby result on approximate counting [34] into the parameterised setting. Specifically, they used this framework to prove the following result.

Theorem 3.13 ([6]). *Let \mathcal{C} be a class of graphs of bounded treewidth. Then there is an FPTRAS for \mathbf{p} -#SUB(\mathcal{C}).*

This result was generalised in [31] to cover all monotone properties whose minimal elements have bounded treewidth.

Theorem 3.14 ([31]). *Let Φ be a monotone property such that $\min(\Phi)$ is a class of graphs of bounded treewidth. Then there exists an FPTRAS for \mathbf{p} -#ISWP(Φ); the same is true for \mathbf{p} -#MISWP(Φ).*

Once again, the random colouring process can be derandomised, this time using so-called (ϵ, k) -balanced families of hash functions from $[n]$ to $[k]$: such a family $\mathcal{F}_{n,k}$ has the property that there exists some constant T so that, for any $U \in [n]^{(k)}$, the number of functions $f \in \mathcal{F}_{n,k}$ such that U is colourful with respect to f is between $(1 - \epsilon)T$ and $(1 + \epsilon)T$. Constructions of such families were given in [1,2]. However, it was also shown in [1] that the minimum size of a family of perfectly k -balanced hash functions from $[n]$ to $[k]$ (that is, a $(0, k)$ -balanced family) must be at least $c(k)n^{\lfloor k/2 \rfloor}$ (where $c(k)$ is a positive constant depending only on k), meaning that we cannot obtain an exact fpt-algorithm for \mathbf{p} -#ISWP(Φ) by this method even if there exists an exact fpt-algorithm for \mathbf{p} -#MISWP(Φ). In fact, it is clear that this method must fail in the exact counting case unless $\text{FPT} = \#\text{W}[1]$, as it would give fpt-algorithms for \mathbf{p} -#PATH and \mathbf{p} -#MATCHING.

3.2.5. The excluded grid theorem

We have already seen that a number of positive results for decision and approximate counting rely on methods for finding copies of graphs whose treewidth is bounded. It is natural to ask whether the use of tree decompositions to give efficient algorithms for such problems is necessary, or merely convenient: some partial answers to this question have been given by exploiting the celebrated Excluded Grid Theorem due to Robertson and Seymour [41], and we use this result to explore the question more thoroughly in Section 4.

Theorem 3.15 (Excluded Grid Theorem [41]). *There is a computable function $w : \mathbb{N} \rightarrow \mathbb{N}$ such that the $(k \times k)$ grid is a minor of every graph of treewidth at least $w(k)$.*

This result has been used to prove the $\text{W}[1]$ -hardness of problems involving homomorphisms by Grohe, Schwenck and Segoufin [29] and by Grohe [28]; the latter considers the following problem, where \mathcal{C} is a fixed class of recursively enumerable graphs⁵:

p-HOM($\mathcal{C}, -$)

Input: A graph $G = (V, E)$ and $H \in \mathcal{C}$.

Parameter: $k = |H|$.

Question: Is there a homomorphism from H to G ?

A homomorphism from H to G is a mapping $\theta : V(H) \rightarrow V(G)$ such that, for all $uv \in E(H)$, we have $\theta(u)\theta(v) \in E(G)$; the key difference from an isomorphism is that the mapping θ need not be injective in this case. The *core* of a graph H is the subgraph H' of H such that there is a homomorphism from H to H' , but there is no homomorphism from H' to a proper subgraph of H' (if more than one subgraph H' of H satisfies this definition, it is known that the two subgraphs must be isomorphic). A graph H has treewidth at most t modulo homomorphic equivalence if the core of H has treewidth at most t .

In this setting, Grohe proved in [28] that \mathbf{p} -HOM($\mathcal{C}, -$) is $\text{W}[1]$ -hard whenever \mathcal{C} is a class of recursively enumerable graphs that do not have bounded treewidth modulo homomorphic equivalence. Since then, a number of authors have built on this result, including [17,12,13,22].

Broadly speaking, the key idea in all of these constructions is to use a $k \times \binom{k}{2}$ grid (whose existence as a minor in one of the graphs in the class under consideration follows from Theorem 3.15) to encode the existence of a clique in some other graph G ; to achieve this, a new graph G' is constructed (which will be the input to the problem to which the reduction is

⁵ In fact [28] addressed the more general problem of homomorphisms between arbitrary relational structures, but we restrict attention to the special case of graph homomorphisms here to avoid introducing additional notation.

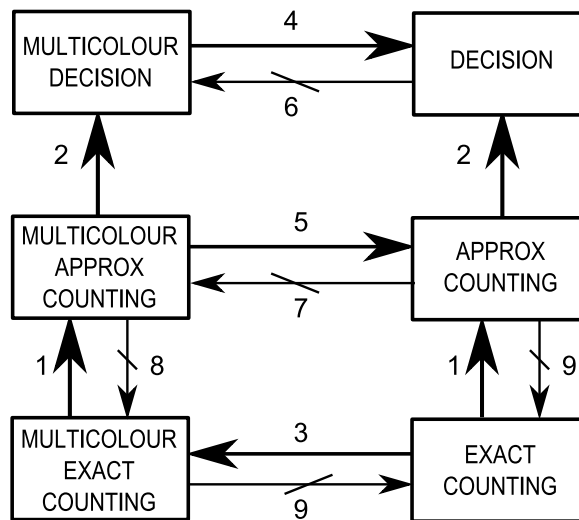


Fig. 2. Summary of the complexity implications between the problems considered 1. Trivial. 2. Proposition 3.16 (randomised decision algorithm). 3. Proposition 3.18 (inclusion–exclusion). 4. Proposition 3.9 (color-coding). 5. Proposition 3.12 (color-coding). 6. Counterexample: \mathbf{p} -CLIQUE OR INDEPENDENT SET ([6] and Proposition 3.7). 7. Counterexample: \mathbf{p} -#CLIQUE OR INDEPENDENT SET (Corollary 3.5 and Proposition 3.7, together with Corollary 3.17). 8. Counterexample: \mathbf{p} -#CONNECTED INDUCED SUBGRAPH [31]. 9. Counterexample: \mathbf{p} -#SUB(\mathcal{C}) for any class \mathcal{C} of bounded treewidth but unbounded vertex-cover number, e.g. \mathbf{p} -#PATH ([6] and [16]).

being made). In G' , there is a vertex-class corresponding to each vertex of a $k \times \binom{k}{2}$ grid, with vertices in each class indexed by elements of $V(G) \times E(G)$; conditions are placed on edges between the vertex classes to ensure that any copy of a $k \times \binom{k}{2}$ grid in G' which includes one vertex from each vertex class must correspond to a clique in G .

The important point to emphasise in this description is the fact that we only want to identify grids which include one vertex from each vertex class. Different authors have used varying strategies to ensure that the grids we identify do indeed include one vertex from each vertex class: in [28], the condition that we are considering a graph whose core has unbounded treewidth means we cannot miss any of the vertex classes, whereas in [29] the authors considered a *coloured* version of the homomorphism problem, for which unbounded treewidth alone is enough to imply $W[1]$ -hardness. Dalmau and Jonsson [17] consider the counting version of the homomorphism problem, \mathbf{p} -#HOM(\mathcal{C} , $-$), and show that fixed parameter tractability coincides precisely with classes \mathcal{C} of bounded treewidth, using an inclusion–exclusion method (only possible in the exact counting setting) to count precisely the number of grids that use every vertex class. Färnqvist and Jonsson studied a variant, \mathbf{p} -LHOM(\mathcal{C} , $-$) in which restrictions are placed on the set of vertices in G to which each vertex in H may be mapped by the homomorphism (resulting in a problem that generalises the well-known LISTCOLOURING for graphs), which again allows grids that do not use every vertex class to be avoided, so the condition for $W[1]$ -hardness is again that \mathcal{C} has bounded treewidth.

Very recently, Curticapean and Marx also used these ideas to prove # $W[1]$ -completeness of \mathbf{p} -#SUB(\mathcal{C}) when \mathcal{C} is a class of graphs of unbounded treewidth; this is a special case of our result on exact counting given in Theorem 4.6. The proof of their result involves proving hardness of the problem \mathbf{p} -#PARTITIONEDSUB(\mathcal{C}) (analogous to a counting version of \mathbf{p} -LHOM(\mathcal{C} , $-$) in the discussion above) as an intermediate step; this problem is then reduced to \mathbf{p} -#SUB(\mathcal{C}).

In Section 4, we will also exploit colourings to allow us to identify those grids that include one vertex from each vertex class. Due to a general result in Section 3.3 (Proposition 3.18) we can reduce the multicolour version to the uncoloured version when considering exact counting, but this is not possible for approximate counting or the decision version, so our results are only valid for the multicolour problem in these settings.

3.3. Relationships between results

We have already seen in Section 3.2 certain relationships between the complexities of the different problem variants; in this section we derive some further implications between the complexities of the different variants of subgraph counting problems, and then summarise all the known relationships in Fig. 2.

Relationships between exact counting, approximate counting and decision

We begin by noting the trivial fact that, if the counting version of a problem is in FPT, then it admits an FPTRAS and the corresponding decision problem is also in FPT. Similarly, we now prove that, if a counting problem admits an FPTRAS, then there is a *randomised* algorithm for the corresponding decision problem.

Proposition 3.16. Let Φ be a family (ϕ_1, ϕ_2, \dots) of functions $\phi_k : \mathcal{L}(k) \rightarrow \{0, 1\}$, such that the function mapping $k \mapsto \phi_k$ is computable, and suppose that there exists an FPTRAS for $\mathbf{p}\text{-}\#\text{ISWP}(\Phi)$. Then, given a graph G on n vertices and $k \in \mathbb{N}$, there is a randomised algorithm, running in time $f(k) \cdot n^{O(1)}$, which

1. if (G, k) is a yes-instance for $\mathbf{p}\text{-ISWP}(\Phi)$, returns YES with probability greater than $1/2$, and
2. if (G, k) is a no-instance for $\mathbf{p}\text{-ISWP}(\Phi)$, returns NO.

The same is true for the corresponding multicolour problems.

Proof. Suppose that there exists an FPTRAS for $\mathbf{p}\text{-}\#\text{ISWP}(\Phi)$. Running this randomised approximation algorithm with $\epsilon = 1/2$ and $\delta = \frac{1}{2n+1}$ immediately gives a randomised fpt-algorithm A for the decision problem with two-sided error probability strictly smaller than $\frac{1}{2n}$: if the answer to the decision problem is NO then the approximation algorithm must output “0” with probability at least $1 - \frac{1}{2n+1}$, whereas if the answer to the decision problem is YES then the approximation algorithm must output a value of at least $1/2$ with probability at least $1 - \frac{1}{2n+1}$. It remains to show that this two-sided error can be reduced to a one-sided error, without increasing the probability of error too much.

We do this using standard techniques, similar to those employed to show that certain problems (so-called “self-reducible” problems) in the class BPP are in fact contained in RP; the idea is that before outputting the answer YES we will search for a witness to the fact that the input is indeed a yes-instance. We begin by running our randomised decision algorithm A ; if the output of A is NO then we return NO. If, on the other hand, the output of A is YES, then we begin our search for a witness.

To do this, we consider each vertex in turn, and either delete it or add it to a list of “necessary” vertices. We begin by choosing an arbitrary vertex v of G , which we (temporarily) delete, and then run A on the resulting graph. If the output of A on this new graph is YES, we permanently delete v and then repeat the process (but always ensuring that we do not delete a vertex we have previously marked as “necessary”); if the output is NO then we reinstate the vertex v , marking it as “necessary”, and repeat. This process terminates when either:

1. $k + 1$ vertices have been marked as necessary, in which case we return NO, or
2. there are exactly k vertices remaining, in which case we check (in time depending only on k) whether ϕ_k is true on any tuple formed of the remaining vertices; if so we return YES, and otherwise we return NO.

It is clear that this witness search process is fixed parameter tractable, and moreover that it will only return YES if we do indeed have a yes-instance. To bound the probability of returning NO if we actually have a yes-instance, observe that we call the algorithm A at most n times (where n is the number of vertices in G), so by the union bound the probability that the algorithm will return the incorrect answer on at least one of these calls is at most $\frac{n}{2n+1} < \frac{1}{2}$; if A returns the correct answer on all of the calls then we are sure to return YES on any yes-instance. This gives the required algorithm in the uncoloured case.

The corresponding result for the multicolour versions of the problems follows by exactly the same argument. \square

Subject to the assumption that $W[1]$ is not equal to FPT under randomised parameterised reductions (as described in Section 2.2), this gives the following immediate corollary.

Corollary 3.17. Assume that $W[1]$ is not equal to FPT under randomised parameterised reductions, and let Φ be a family (ϕ_1, ϕ_2, \dots) of functions $\phi_k : \mathcal{L}(k) \rightarrow \{0, 1\}$, such that the function mapping $k \mapsto \phi_k$ is computable. Then, if $\mathbf{p}\text{-ISWP}(\Phi)$ is $W[1]$ -complete, there is no FPTRAS for $\mathbf{p}\text{-}\#\text{ISWP}(\Phi)$. Similarly, if $\mathbf{p}\text{-MISWP}(\Phi)$ is $W[1]$ -complete, there is no FPTRAS for $\mathbf{p}\text{-}\#\text{MISWP}(\Phi)$.

Recently, more refined models of parameterised random complexity have been developed by Montoya and Müller [39] and Chauhan and Rao [11], in which the use of randomness is restricted. However, given that the definition of an FPTRAS does not place any restrictions on the nature of the randomised algorithm, the existence of an FPTRAS for a problem does not have immediate implications in this new framework. Possible future research directions to address this issue are discussed in Section 5.

Relationships between multicolour and uncoloured problems

The relationships between the complexities of multicolour and uncoloured versions of the problems are somewhat more complicated. For the case of exact counting, the following result was proved in [31], using an inclusion–exclusion method reminiscent of that used in [17] (and discussed in Section 3.2.5).

Proposition 3.18. For any family Φ , we have $\mathbf{p}\text{-}\#\text{MISWP}(\Phi) \leq_{\text{fpt}}^{\text{fpt}} \mathbf{p}\text{-}\#\text{ISWP}(\Phi)$.

However, we saw in Section 3.2.4 (Propositions 3.9 and 3.12) that this relationship between the multicolour and uncoloured variants for any given property Φ is reversed in the cases of approximate counting and decision.

All of these relationships between the complexities of the six different problem variants are summarised in Fig. 2: an arrow from problem A to problem B implies that (subject to standard assumptions in parameterised complexity theory), if problem A is efficiently solvable (either it is in FPT, or it admits an FPTRAS, as appropriate), then there must also be an efficient algorithm for problem B . A crossed out arrow from A to B indicates that there is some property Φ for which variant A has an efficient algorithm but variant B does not.

4. New results

In this section, we extend methods based on the Excluded Grid Theorem, discussed in Section 3.2.5, to give new hardness results for several of the problem variants. Specifically, we prove that the following results hold whenever $\min^*(\Phi)$ is a class of graphs having unbounded treewidth.

1. Both $\mathbf{p}\text{-}\#\text{ISWP}(\Phi)$ and $\mathbf{p}\text{-}\#\text{MISWP}(\Phi)$ are $\#\text{W}[1]$ -complete (Theorem 4.6).
2. The multicolour decision problem, $\mathbf{p}\text{-}\text{MISWP}(\Phi)$ is $\text{W}[1]$ -complete (Theorem 4.7).
3. Assuming that $\text{W}[1]$ is not equal to FPT under randomised parameterised reductions, there is no FPTRAS for $\mathbf{p}\text{-}\#\text{MISWP}(\Phi)$ (Corollary 4.9).

Recall that inclusion of these problems in the appropriate classes was discussed in Section 2.3, so it suffices for the first two results to prove the appropriate hardness results.

Our hardness results apply to any property Φ satisfying the condition that $\min^*(\Phi)$ has unbounded treewidth. We will later restrict our attention to uniformly monotone properties, allowing us to give dichotomy results for decision and approximate counting in the multicolour setting, when the hardness results are combined with existing positive results discussed above; the dichotomies are stated explicitly in Section 4.2.

All of these results are based on a common construction, described and analysed in Section 4.1, which exploits the Excluded Grid Theorem (Theorem 3.15). Making use of the properties of the construction established in Section 4.1, the proofs of the complexity results given in Section 4.2 are reasonably straightforward.

Before embarking on the proof of these results, we consider briefly some examples of properties to which they are applicable, namely properties whose sets of minimal satisfying elements do not have bounded treewidth. Some obvious examples come from properties that in some way depend on the treewidth of the graph: for example, if we set $\phi_k(H, \pi) = 1$ if and only if H has treewidth at least $\theta(k)$, where θ is any function such that $\theta(k) \rightarrow \infty$ as $k \rightarrow \infty$, then it is clear that there is no constant that bounds the treewidth of elements of $\min(\Phi)$. The same is true if, for example, we set $\phi_k(H, \pi) = 1$ if and only if the chromatic number of H is at least $\theta(k)$ (where again $\theta(k) \rightarrow \infty$ as $k \rightarrow \infty$), since any graph with treewidth at most t has chromatic number at most $t + 1$ (a greedy colouring with $t + 1$ colours can be obtained by considering the bags of the indexing tree in a post-order traversal, with respect to some arbitrarily chosen root).

It is also clear that our results apply when ϕ_k is true only if H contains a complete bipartite spanning subgraph. This category of problems includes the case in which $\phi_k(H) = 1$ if and only if k is even and H contains $K_{\frac{k}{2}, \frac{k}{2}}$ as a subgraph; the parameterised complexity of this problem from the point of view of decision (in the uncoloured case) was a long-standing open problem, recently resolved by Lin [37].

Slightly less obvious examples of properties Φ for which $\min(\Phi)$ does not have bounded treewidth arise when ϕ_k is true on a k -vertex subgraph if and only if certain connectivity conditions are satisfied: while $\mathbf{p}\text{-}\#\text{CONNECTED INDUCED SUBGRAPH}$, studied in [31], is an example of a property for which all minimal elements do have bounded treewidth (since the minimal elements satisfying ϕ_k are precisely all trees on k vertices), if we instead demand stronger connectivity conditions this is no longer the case. The graph obtained by subdividing every edge of the $k \times k$ grid exactly once (that is, replacing every edge of the grid with a new vertex adjacent to both endpoints of the original edge) has treewidth at least k , as it contains the $k \times k$ grid as a minor, and is also a minimally 2-connected graph, in the sense that it is 2-connected but deleting any edge will result in a graph that is no longer 2-connected (as it will contain a vertex of degree one). Similarly, the graph illustrated in Fig. 3 has treewidth at least k and is 3-connected, but deleting any edge will result in a graph that is no longer 3-connected (as such a subgraph will contain at least one vertex of degree at most 2). Both of these problems were already shown to be hard from the point of view of decision in the uncoloured case in [8], in which the authors proved that the problem of deciding whether a graph contains an r -connected induced subgraph on k vertices is $\text{W}[1]$ -complete.

Combining any of these properties with an additional condition on, say, the number of edges in the subgraph leads to non-monotone properties to which our results apply: a concrete example would be a property Φ such that $\phi_k(H, \pi)$ is true if and only if H is 2-connected and has an even number of edges.

4.1. Graphs with grid minors

In this section we present a construction based on a graph which contains a large grid minor; we define the construction in Section 4.1.1, then analyse its key properties in Section 4.1.2. This construction will form the basis of the complexity results mentioned above and proved in Section 4.2. The goal of the construction is to define, given an arbitrary graph G and a graph H containing a large grid as a minor, a new coloured graph G_H such that the number of colourful copies of H in G_H and the number of k -cliques in G are closely related.

Before presenting this construction, we make a simple observation about situations in which we will be able to apply this construction.

Proposition 4.1. *Let Φ be a family (ϕ_1, ϕ_2, \dots) of functions $\phi_k : \mathcal{L}(k) \rightarrow \{0, 1\}$, such that the function mapping $k \mapsto \phi_k$ is computable. Suppose that $\min^*(\Phi)$ is a class of graphs having unbounded treewidth. Then, for every $j \in \mathbb{N}$ there exists $k \in \mathbb{N}$ such that some $H \in \min^*(\phi_k)$ contains the $(j \times \binom{j}{2})$ grid as a minor; moreover, the value of this k is bounded by some computable function of j .*




Fig. 3. A minimally 3-connected graph with treewidth at least k .

Proof. By the Excluded Grid Theorem (Theorem 3.15), there exists a computable function w such that every graph with treewidth at least $w\binom{j}{2}$ contains as a minor the $\binom{j}{2} \times \binom{j}{2}$ grid, and hence the $j \times \binom{j}{2}$ grid. The fact that $\min^*(\Phi)$ has unbounded treewidth means that there is a computable function $g : \mathbb{N} \rightarrow \mathbb{N}$ such that, for any $w\binom{j}{2} \in \mathbb{N}$, there exists $k \leq g(w\binom{j}{2})$ such that some $H \in \min^*(\phi_k)$ has treewidth at least $w\binom{j}{2}$ (observe that g can be computed by considering ϕ_1, ϕ_2, \dots in turn). Thus, $H \in \min^*(\phi_k)$ contains the $j \times \binom{j}{2}$ grid as a minor, as required. \square

4.1.1. Construction

In this section we describe the construction that will be used to prove our complexity results in Section 4.2.

Let G be any graph, and $k \in \mathbb{N}$, and let H be a graph which contains the $k \times \binom{k}{2}$ grid as a minor. (Later, we will be interested in determining the number of k -cliques in G .) We then define a new coloured graph G_H , and a colouring f of its vertices with $|V(H)|$ colours. In order to define our construction, it will be necessary to fix an arbitrary total order $<$ on $V(G)$. To assist in the explanation of this construction, an example of a pair of graphs G and H , for $k = 3$, is illustrated in Fig. 4.

Let A denote the $k \times \binom{k}{2}$ grid. We will label each of the k rows with a distinct element of the set $\{1, \dots, k\}$, and each of the $\binom{k}{2}$ columns with a distinct unordered pair of elements from the same set. We then denote by $a_{(i, \{j, l\})}$ the vertex in row i and column $\{j, l\}$.

Since the grid A is a minor of H , there exists a function $m : V(A) \rightarrow \mathcal{P}(V(H))$ such that the image of each vertex in A induces a connected subgraph in H , distinct vertices of A map to disjoint subsets of $V(H)$, and for each edge $aa' \in A$ there exist $u \in m(a)$ and $u' \in m(a')$ such that $uu' \in E(H)$. Note that the union of the image of m is not necessarily the whole of $V(H)$; if not, then we denote by V'_H the vertices of $V(H) \setminus \bigcup_{a \in V(A)} m(a)$.

In our new graph G_H , we will have multiple copies of the subgraph $H[m(a)]$ of H for each $a \in V(A)$, each one indexed by a pair consisting of a vertex and an edge from G . Specifically, for each $a_{(i, \{j, l\})} \in V(A)$, we will have a copy $H_{(i, \{j, l\})}^{(v, e)}$ of $H[m(a_{(i, \{j, l\})})]$ for each pair $(v, e) \in V(G) \times E(G)$ that satisfies the condition

$$\text{if } i \in \{j, l\} \text{ then } v \text{ is incident with } e. \tag{1}$$

In the example of Fig. 4, we will therefore have eight copies of $H[m(a_{(1, \{1, 2\})})]$ (each of which is a triangle): $H_{(1, \{1, 2\})}^{(v_1, e_1)}$, $H_{(1, \{1, 2\})}^{(v_1, e_3)}$, $H_{(1, \{1, 2\})}^{(v_2, e_1)}$, $H_{(1, \{1, 2\})}^{(v_2, e_2)}$, $H_{(1, \{1, 2\})}^{(v_3, e_2)}$, $H_{(1, \{1, 2\})}^{(v_3, e_3)}$, $H_{(1, \{1, 2\})}^{(v_3, e_4)}$, and $H_{(1, \{1, 2\})}^{(v_4, e_4)}$.

For each $u \in m(a_{(i, \{j, l\})})$, we denote by $u^{(v, e)}$ the corresponding vertex in $H_{(i, \{j, l\})}^{(v, e)}$ (note that each $u \in V(H)$ belongs to subgraphs $H_{(i, \{j, l\})}$ for only a single value of $(i, \{j, l\})$). For each $(i, \{j, l\}) \in [k] \times [k]^{(2)}$, we denote by $V_{(i, \{j, l\})}$ the union of the




Fig. 4. An example pair of graphs G (above) and H (below), with $k = 3$ (note that in general H will not be based on a square grid, as $k \neq \binom{k}{2}$ for larger k).

following set:

$$\left\{ V \left(H_{(i,\{j,l\})}^{(v,e)} \right) : (v, e) \in V(G) \times E(G), \text{ and } i \in \{j, l\} \implies v \text{ incident with } e \right\}.$$

We now set

$$V(G_H) = \bigcup_{(i,\{j,l\}) \in [k] \times [k]^{(2)}} V_{(i,\{j,l\})} \cup V'_H.$$

That is, the vertex-set of G_H is made up of all the vertices in subgraphs $H_{(i,\{j,l\})}^{(v,e)}$, together with any vertices in V'_H .

We now define the colouring f of $V(G_H)$. First, fix a colouring ω of $V(H)$ which gives every vertex a different colour from the set $\{1, \dots, |V(H)|\}$. For any $u \in V'_H$, we now set $f(u) = \omega(u)$, while for every vertex $u^{(v,e)} \in V(G_H)$ we set $f(u^{(v,e)}) = \omega(u)$.

To complete the definition of G_H , we now define its edge-set. Let $u, w \in V(G_H)$. Then $uw \in E(G_H)$ if and only if the following conditions are satisfied:

Condition 1. $\omega^{-1}(f(u))\omega^{-1}(f(w)) \in E(H)$ (i.e. we only join vertices if their colours occur at adjacent vertices in H), and

Condition 2. if $u \in H_{(i,\{j,l\})}^{(v,e)}$ and $w \in H_{(i',\{j',l'\})}^{(v',e')}$ then

1. if $i = i'$ then $v = v'$, and if $i < i'$ then $v < v'$, and
2. if $\{j, l\} = \{j', l'\}$ then $e = e'$.

Note that the second condition implies that there are no edges between $H_{(i,\{j,l\})}^{(v,e)}$ and $H_{(i,\{j,l\})}^{(v',e')}$ for $(v, e) \neq (v', e')$. In Figs. 5 and 6 we illustrate parts of this construction for the example in Fig. 4, showing the edges between $V_{(1,\{1,2\})}$ and $V_{(2,\{1,2\})}$ (Fig. 5) and between $V_{(2,\{1,2\})}$ and $V_{(2,\{2,3\})}$ (Fig. 6); for clarity in the diagram, we label each subgraph $H_{(i,\{j,l\})}^{(v,e)}$ with (v, e) only.

Finally, observe that, given G and H , the graph G_H and its colouring f can clearly be constructed in time $h(k) + n^{O(1)}$, where h (a computable function of k) represents the time required to find a grid minor in H .

4.1.2. Analysis of the construction

In this section, we prove the important properties of the construction defined above; these properties will form the basis of the proofs of our complexity results in Section 4.2. We begin by giving a necessary and sufficient condition for a colourful subgraph in G_H to be isomorphic to H .

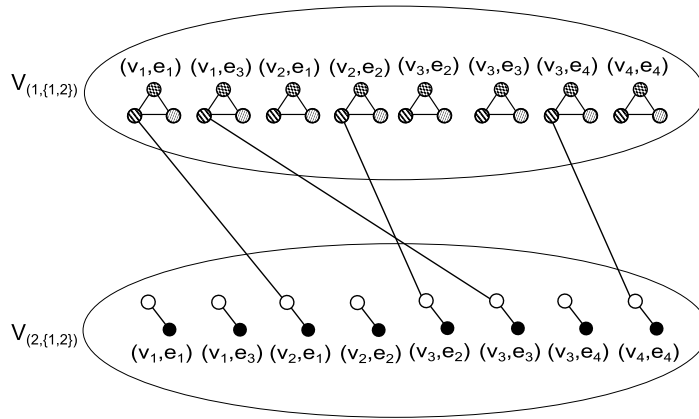


Fig. 5. Edges between $V_{(1, \{1,2\})}$ and $V_{(2, \{1,2\})}$.




Fig. 6. Edges between $V_{(2, \{1,2\})}$ and $V_{(2, \{2,3\})}$.

Lemma 4.2. Let Y be a colourful subset of $V(G_H)$. Then $G_H[Y]$ is isomorphic to a subgraph of H . Moreover, $G_H[Y]$ is isomorphic to H if and only if, for every $u \in Y \cap H_{(i, \{j, l\})}^{(v, e)}$ and $w \in Y \cap H_{(i', \{j', l'\})}^{(v', e')}$ such that $\omega^{-1}(f(u))\omega^{-1}(f(w)) \in E(H)$, the following conditions are satisfied:

1. if $i = i'$ then $v = v'$, and if $i < i'$ then $v < v'$, and
2. if $\{j, l\} = \{j', l'\}$ then $e = e'$.

Proof. Let $\pi : Y \rightarrow V(H)$ be the mapping given by $\pi(u) = \omega^{-1}(f(u))$. We claim first that Y defines an isomorphism from $G_H[Y]$ to a subgraph of H . To verify this claim, it suffices to check that, whenever $u, w \in Y$ with $uw \in E(G_H)$, we also have $\pi(u)\pi(w) \in E(H)$. However, by the first condition in the definition of $E(G_H)$, we can only have $uw \in E(G_H)$ if $\omega^{-1}(f(u))\omega^{-1}(f(w)) \in E(H)$, or in other words if $\pi(u)\pi(w) \in E(H)$, as required.

Now, suppose that Y is such that, for every $u \in Y \cap H_{(i, \{j, l\})}^{(v, e)}$ and $w \in Y \cap H_{(i', \{j', l'\})}^{(v', e')}$ with $\omega^{-1}(f(u))\omega^{-1}(f(w)) \in E(H)$, the following conditions are satisfied:

1. if $i = i'$ then $v = v'$, and if $i < i'$ then $v < v'$, and
2. if $\{j, l\} = \{j', l'\}$ then $e = e'$.

We claim that in this case π defines an isomorphism from $G_H[Y]$ to H . This claim holds provided that, for all $u, w \in Y$ such that $\pi(u)\pi(w) \in E(H)$ we also have $uw \in E(G_H)$. The assumption that $\pi(u)\pi(w) \in E(H)$ implies uw satisfies **Condition 1** in the definition of $E(G_H)$ above, and the condition that we are assuming is exactly **Condition 2** in this definition; thus it follows immediately that in this situation we will have $uw \in E(G_H)$ as required.

Conversely, suppose that the colourful subset $Y \subset V(G_H)$ induces a subgraph isomorphic to H ; we need to verify that, for every $u \in Y \cap H_{(i,\{j,l\})}^{(v,e)}$ and $w \in Y \cap H_{(i',\{j',l'\})}^{(v',e')}$ such that $\omega^{-1}(f(u))\omega^{-1}(f(w)) \in E(H)$, the conditions in the statement of the lemma are satisfied. By definition of π we know that, for any such u and w , $\pi(u)\pi(w) \in E(H)$. Moreover, since we have already established that π defines an isomorphism from $G_H[Y]$ to a subgraph of H , it follows in this case that π must be an isomorphism from $G_H[Y]$ to H ; thus the fact that $\pi(u)\pi(w) \in E(H)$ implies that $uw \in E(G_H)$. It then follows immediately from [Condition 2](#) that

1. if $i = i'$ then $v = v'$, and if $i < i'$ then $v < v'$, and
2. if $\{j, l\} = \{j', l'\}$ then $e = e'$,

as required. \square

Using this characterisation, we can now relate the number of colourful subgraphs in G_H that are isomorphic to H to the number of cliques in G .

Lemma 4.3. *Let G, H, G_H and f be as described above. Then*

$$\#ColSubInd(H, G_H, f) = \#Clique_k(G).$$

Proof. We will prove this result by showing that there is a one-to-one correspondence between colourful copies of H in G_H with colouring f and k -vertex subsets of $V(G)$ that induce cliques. We begin by giving an injective map from the set of k -vertex subsets inducing cliques in G to the colourful copies of H in G_H with colouring f ; we will then proceed to give an injective map in the opposite direction.

Given a subset $X \in V(G)^{(k)}$ that induces a clique in G , we denote by x_1, \dots, x_k the elements of X , with $x_1 < \dots < x_k$. We then claim that the mapping π , where

$$\pi(X) = V'_H \cup \bigcup_{\substack{i,j,l \in [k] \\ j \neq l}} V\left(H_{(i,\{j,l\})}^{(x_i, x_j x_l)}\right)$$

is an injective mapping from k -vertex subsets of $V(G)$ that induce cliques in G to colourful copies of H in G_H (with respect to the colouring f).

Observe first that this mapping is well-defined: if $\{x_1, \dots, x_k\}$ induces a clique in G , all edges $x_j x_l$ (with $j, l \in [k], j \neq l$) are edges of G , and moreover x_i is incident with $x_j x_l$ whenever $i \in \{j, l\}$, so the specified subgraphs $H_{(i,\{j,l\})}^{(x_i, x_j x_l)}$ do exist in G_H .

It is immediate from the definition of π that, for $X \neq X'$, we will have $\pi(X) \neq \pi(X')$. Thus, to show that π defines an injective map from the set of k -tuples inducing cliques in G to the set of colourful copies of H in G_H , it remains only to demonstrate that, whenever X induces a clique in G , $\pi(X)$ must induce a colourful copy of H in G_H (with respect to f).

We verify first that $\pi(X)$ is indeed colourful with respect to f . First note that all vertices of V'_H have distinct colours, and do not share colours with any vertex in $V(G_H) \setminus V'_H$. Next observe that, for each $(i, \{j, l\}) \in [k] \times [k]^{(2)}$, there exists a unique pair $(v, e) \in V(G) \times E(G)$ such that $\pi(X) \cap V_{(i,\{j,l\})} = V\left(H_{(i,\{j,l\})}^{(v,e)}\right)$. Thus, $\pi(X)$ contains exactly one vertex $u^{(v,e)}$ (for some $(v, e) \in V(G) \times E(G)$) for each $u \in V(H) \setminus V'_H$, and so by definition of f we see that $\pi(X)$ is indeed colourful with respect to f .

It remains to verify that $\pi(X)$ induces a copy of H . As we know that $\pi(X)$ is colourful, it suffices by [Lemma 4.2](#) to verify that, for every pair of vertices $u, w \in \pi(X) \setminus V'_H$ such that $u \in H_{(i,\{j,l\})}^{(v,e)}$, $w \in H_{(i',\{j',l'\})}^{(v',e')}$ and $\omega^{-1}(f(u))\omega^{-1}(f(w)) \in E(H)$, the following conditions are satisfied:

1. if $i = i'$ then $v = v'$, and if $i < i'$ then $v < v'$, and
2. if $\{j, l\} = \{j', l'\}$ then $e = e'$.

However, these two conditions follow immediately from the definition of π , so $\pi(X)$ does indeed induce a colourful copy of H , as required.

We now proceed to define a mapping in the opposite direction, that is an injective mapping from colourful copies of H in G_H (with colouring f_H) to k -vertex subsets of $V(G)$ that induce cliques in G . Suppose that the colourful subset $Y \subset V(G_H)$ induces a copy of H . In order to define our mapping, we prove a series of claims. We begin by arguing that the colourful subset Y can be used to define a mapping from $[k] \times [k]^{(2)}$ to pairs in $V(G) \times E(G)$.

Claim 1. *For each $(i, \{j, l\}) \in [k] \times [k]^{(2)}$, there is a unique pair $(v, e) \in V(G) \times E(G)$ such that $Y \cap V_{(i,\{j,l\})} = V\left(H_{(i,\{j,l\})}^{(v,e)}\right)$.*

Proof of Claim 1. It follows from [Lemma 4.2](#) and the connectedness of $H[m(u)]$ for each $u \in V(H)$ that, for each $(i, \{j, l\}) \in [k] \times [k]^{(2)}$, there exists a unique pair $(v, e) \in V(G) \times E(G)$ such that $Y \cap V_{(i,\{j,l\})} \subseteq V\left(H_{(i,\{j,l\})}^{(v,e)}\right)$. Thus, as Y is colourful, we must in fact have $Y \cap V_{(i,\{j,l\})} = V\left(H_{(i,\{j,l\})}^{(v,e)}\right)$ for this (v, e) , as required (otherwise we would miss some colour). \square ([Claim 1](#))

Hence we can define a mapping $\sigma_Y : [k] \times [k]^{(2)} \rightarrow V(G) \times E(G)$ by setting $\sigma_Y(i, \{j, l\})$ to be this unique pair (v, e) such that $Y \cap V_{(i, \{j, l\})} = V\left(H_{(i, \{j, l\})}^{(v, e)}\right)$; we denote by $\sigma_Y^1(i, \{j, l\})$ and $\sigma_Y^2(i, \{j, l\})$ respectively the first and second elements of this pair (so $\sigma_Y^1(i, \{j, l\}) \in V(G)$, and $\sigma_Y^2(i, \{j, l\}) \in E(G)$).

We now show that the value of σ_Y^1 depends only on its first argument, and σ_Y^2 only on its second.

Claim 2. For each $i \in [k]$ and every $\{j, l\}, \{j', l'\} \in [k]^{(2)}$, we have $\sigma_Y^1((i, \{j, l\})) = \sigma_Y^1((i, \{j', l'\}))$. Similarly, for each $\{j, l\} \in [k]^{(2)}$ and every $i, i' \in [k]$, we have $\sigma_Y^2((i, \{j, l\})) = \sigma_Y^2((i', \{j, l\}))$.

Proof of Claim 2. To prove this claim we suppose, for a contradiction, there exist $\{j, l\}$ and $\{j', l'\}$ such that $\sigma_Y^1((i, \{j, l\})) \neq \sigma_Y^1((i, \{j', l'\}))$; we may assume without loss of generality that $a_{(i, \{j, l\})}$ and $a_{(i, \{j', l'\})}$ are adjacent in the grid A . Now, by definition of the minor map m , there exist vertices $z, z' \in V(H)$ such that $z \in m(a_{(i, \{j, l\})})$, $z' \in m(a_{(i, \{j', l'\})})$ and $zz' \in E(H)$. Let y and y' be the vertices of Y whose colours match z and z' respectively, that is $y = f|_Y^{-1}\omega(z)$ and $y' = f|_Y^{-1}\omega(z')$ (as Y is colourful, the restriction of f to Y is a bijection), so $\omega^{-1}(f(y))\omega^{-1}(f(y')) \in E(H)$. Note that we must have $y \in H_{(i, \{j, l\})}^{\sigma_Y^1((i, \{j, l\}))}$ and $y' \in H_{(i, \{j', l'\})}^{\sigma_Y^1((i, \{j', l'\}))}$. By our initial assumption, $\sigma_Y^1((i, \{j, l\})) \neq \sigma_Y^1((i, \{j', l'\}))$, but by Lemma 4.2 this then means that Y cannot induce a copy of H in G_H , giving a contradiction. Thus we know that, for each $i \in [k]$, we must have $\sigma_Y^1((i, \{j, l\})) = \sigma_Y^1((i, \{j', l'\}))$ for all $\{j, l\}, \{j', l'\} \in [k]^{(2)}$. An analogous argument can be used to show that, for each $\{j, l\} \in [k]^{(2)}$, we must have $\sigma_Y^2((i, \{j, l\})) = \sigma_Y^2((i', \{j, l\}))$ for every $i, i' \in [k]$. \square (Claim 2)

Hence we can define the mapping $\tau_Y^1 : [k] \rightarrow V(G)$ by setting $\tau_Y^1(i)$ to be the unique vertex $v \in V(G)$ such that, for any value of $\{j, l\} \in [k]^{(2)}$, $\sigma_Y^1((i, \{j, l\})) = v$, and the mapping $\tau_Y^2 : [k]^{(2)} \rightarrow E(G)$ by setting $\tau_Y^2(\{j, l\})$ to be the unique edge $e \in E(G)$ such that, for any value of $i \in [k]$, $\sigma_Y^2((i, \{j, l\})) = e$.

Next, we demonstrate that the mapping τ_Y^1 is injective.

Claim 3. For $i \neq i'$, we have $\tau_Y^1(i) \neq \tau_Y^1(i')$.

Proof of Claim 3. We argue that, for $1 \leq i \leq k - 1$, we must have $\tau_Y^1(i) < \tau_Y^1(i + 1)$, implying that $\tau_Y^1(1) < \tau_Y^1(2) < \dots < \tau_Y^1(k)$ and in particular that no two of these vertices are the same. As in the proof of Claim 2, we observe that by definition of the minor map m , for each $i \in \{1, \dots, k - 1\}$, there exist $z, z' \in V(H)$ such that $z \in m(a_{(i, \{1, 2\})})$, $z' \in m(a_{(i+1, \{1, 2\})})$ and $zz' \in E(H)$. As before, we set $y = f|_Y^{-1}\omega(z)$ and $y' = f|_Y^{-1}\omega(z')$, and observe that $\omega^{-1}(f(y))\omega^{-1}(f(y')) \in E(H)$. Note also, by definition of τ_Y^1 and τ_Y^2 , that we have $y \in H_{(i, \{1, 2\})}^{(\tau_Y^1(i), \tau_Y^2(\{1, 2\}))}$ and $y' \in H_{(i+1, \{1, 2\})}^{(\tau_Y^1(i+1), \tau_Y^2(\{1, 2\}))}$. Since $i < i + 1$, it now follows from Lemma 4.2 that $\tau_Y^1(i) < \tau_Y^1(i + 1)$, as required. \square (Claim 3)

We now demonstrate a relationship between the mappings τ_Y^1 and τ_Y^2 .

Claim 4. Let $j, l \in [k]$ with $j \neq l$. Then $\tau_Y^2(\{j, l\}) = \tau_Y^1(j)\tau_Y^2(l)$.

Proof of Claim 4. Set $e = \tau_Y^2(\{j, l\})$. Then, by definition, $\sigma_Y(j, \{j, l\}) = (\tau_Y^1(j), e)$ and $\sigma_Y(l, \{j, l\}) = (\tau_Y^1(l), e)$. Thus, both subgraphs $H_{(j, \{j, l\})}^{(\tau_Y^1(j), e)}$ and $H_{(l, \{j, l\})}^{(\tau_Y^1(l), e)}$ must exist in G_H ; by definition of $V(G_H)$, since $j, l \in \{j, l\}$, this implies that both $\tau_Y^1(j)$ and $\tau_Y^1(l)$ must be incident with e . By Claim 3, we know that $\tau_Y^1(j) \neq \tau_Y^1(l)$, so this is only possible if in fact $e = \tau_Y^1(j)\tau_Y^1(l)$, as required. \square (Claim 4)

Using the facts proved so far, we can now give a mapping from colourful subsets $Y \subset V(G_H)$ with $G_H[Y] \cong H$ to k -tuples of $V(G)$ that induce cliques.

Claim 5. $V_Y = \{\tau_Y^1(i) : 1 \leq i \leq k\}$ induces a k -clique in G .

Proof of Claim 5. We already know from Claim 3 that, for $i \neq i'$, $\tau_Y^1(i) \neq \tau_Y^1(i')$ and hence that $|V_Y| = k$; thus it suffices to verify that, for each $1 \leq i < i' \leq k$, $\tau_Y^1(i)\tau_Y^1(i') \in E(G_H)$. But we know from Claim 4 that $\tau_Y^1(i)\tau_Y^1(i') = \tau_Y^2(\{i, i'\}) \in E(G_H)$, completing the proof of the claim. \square (Claim 5)

We have now given a well-defined mapping $Y \mapsto (\tau_Y^1(1), \dots, \tau_Y^1(k))$ from colourful subsets Y of $V(G_H)$ with $G_H[Y] \cong H$ to k -tuples of $V(G)$ that induce cliques, so in order to complete the proof of the lemma it remains only to argue that this mapping is injective.

Claim 6. The mapping $Y \mapsto (\tau_Y^1(1), \dots, \tau_Y^1(k))$ is injective.

Proof of Claim 6. Suppose that Y and Y' are distinct colourful subsets of G_H with respect to the colouring f , such that $G_H[Y], G_H[Y'] \cong H$. By the reasoning above, we know that there is a function $\tau_Y^1 : [k] \rightarrow V(G)$ such that, for each $(i, \{j, l\}) \in [k] \times [k]^{(2)}$,

$$Y \cap V_{(i, \{j, l\})} = V\left(H_{(i, \{j, l\})}^{(\tau_Y^1(i), \tau_Y^1(j)\tau_Y^1(l))}\right)$$

and

$$Y' \cap V_{(i,j,l)} = V \left(H_{(i,j,l)}^{(\tau_{Y'}^1(i), \tau_{Y'}^1(j), \tau_{Y'}^1(l))} \right).$$

Since Y and Y' are both colourful, we know that

$$\begin{aligned} Y &= V_H' \cup \bigcup_{(i,j,l) \in [k] \times [k]^{(2)}} (Y \cap V_{(i,j,l)}) \\ &= V_H' \cup \bigcup_{(i,j,l) \in [k] \times [k]^{(2)}} V \left(H_{(i,j,l)}^{(\tau_Y^1(i), \tau_Y^1(j), \tau_Y^1(l))} \right), \end{aligned}$$

and similarly

$$\begin{aligned} Y' &= V_H' \cup \bigcup_{(i,j,l) \in [k] \times [k]^{(2)}} (Y' \cap V_{(i,j,l)}) \\ &= V_H' \cup \bigcup_{(i,j,l) \in [k] \times [k]^{(2)}} V \left(H_{(i,j,l)}^{(\tau_{Y'}^1(i), \tau_{Y'}^1(j), \tau_{Y'}^1(l))} \right). \end{aligned}$$

Thus, if $\tau_Y^1 = \tau_{Y'}^1$, we would have $Y = Y'$, a contradiction to our initial assumption. Hence $\tau_Y^1 \neq \tau_{Y'}^1$, implying immediately that $(\tau_Y^1(1), \dots, \tau_Y^1(k)) \neq (\tau_{Y'}^1(1), \dots, \tau_{Y'}^1(k))$, as required. \square (Claim 6)

This claim completes the proof that there is a one-to-one correspondence between colourful subsets $Y \subset V(G_H)$ with $G_H[Y] \cong H$ and k -tuples of vertices that induce cliques in G , and hence completes the proof that

$$\#ColSubInd(H, G_H, f) = \#Clique_k(G),$$

as required. \square

Before proving our complexity results in the next section, we give a simple fact about the number of colourful copies of graphs from some collection \mathcal{H} in G_H , provided that \mathcal{H} satisfies certain conditions.

Lemma 4.4. *Let \mathcal{H} be a collection of labelled graphs on k vertices such that $(H, \pi) \in \mathcal{H}$ and there is no proper subgraph H' of H such that, for some labelling π' , $(H', \pi') \in \mathcal{H}$. Then*

$$\#ColStrEmb(\mathcal{H}, G_H, f) = \#ColStrEmb(\mathcal{H}^H, G_H, f).$$

Proof. We know from Lemma 4.2 that any colourful subset in G_H induces a subgraph of H ; thus, by our assumption that \mathcal{H} does not contain any proper subgraphs of H , it follows immediately that the number of colourful labelled copies of graphs from \mathcal{H} in G_H (with respect to the colouring f) is exactly equal to the number of colourful labelled copies of graphs from \mathcal{H}^H in G_H (with respect to the colouring f). \square

4.2. Complexity results

In this section, we use the construction defined in Section 4.1.1, and in particular its properties established in Section 4.1.2, to prove our new complexity results. In several of the proofs below we will also need a result from [32]:

Lemma 4.5. *Let \mathcal{H} be a collection of labelled graphs, and $(H, \pi) \in \mathcal{H}$ a labelled k -vertex graph. Set*

$$\alpha_H = |\{\sigma : \sigma \text{ a permutation on } [k], \exists (H, \pi') \in \mathcal{H}^H \text{ such that } \pi' \circ \sigma^{-1} \circ \pi^{-1} \text{ defines an automorphism on } H\}|.$$

Then, for any graph G ,

$$\#StrEmb(\mathcal{H}^H, G) = \alpha_H \cdot \#SubInd(H, G).$$

Moreover, if G is equipped with a k -colouring f , then

$$\#ColStrEmb(\mathcal{H}^H, G, f) = \alpha_H \cdot \#ColSubInd(H, G, f).$$

Note that the value of α_H , as defined in the statement of this lemma, can be computed from H and \mathcal{H} in time bounded only by some computable function of k .

Exact counting

We begin with the hardness result for exact counting when $\min^*(\Phi)$ does not have bounded treewidth; this result applies in both the multicolour and uncoloured settings. Note that we do not need to assume that Φ is monotone.

Theorem 4.6. *Let Φ be a family (ϕ_1, ϕ_2, \dots) of functions $\phi_k : \mathcal{L}(k) \rightarrow \{0, 1\}$, such that the function mapping $k \mapsto \phi_k$ is computable. Suppose that $\min^*(\Phi)$ has unbounded treewidth. Then $\mathbf{p}\text{-}\#\text{ISWP}(\Phi)$ and $\mathbf{p}\text{-}\#\text{MISWP}(\Phi)$ are $\#W[1]$ -complete.*

Proof. We prove the result by means of an fpt-Turing reduction from the $\#W[1]$ -complete problem $\mathbf{p}\text{-}\#\text{CLIQUE}$. Note that, by Proposition 3.18, it suffices to give a reduction from $\mathbf{p}\text{-}\#\text{CLIQUE}$ to $\mathbf{p}\text{-}\#\text{MISWP}(\Phi)$.

Let the graph G be the input to an instance of $\mathbf{p}\text{-}\#\text{CLIQUE}$ with parameter k . By Proposition 4.1, we know that there is k' (bounded by some computable function of k) such that some $H \in \min^*(\phi_k)$ contains the $(k \times \binom{k}{2})$ grid as a minor and, moreover, that there is no proper subgraph H' of H which satisfies $\phi_{k'}$ (with any labelling). Let \mathcal{H} be the set of all k' -vertex labelled graphs satisfying $\phi_{k'}$.

Recall that the graph G_H and colouring f can be constructed in time $h(k) + n^{O(1)}$, where h is a computable function. Given the graph G_H , the colouring f and an oracle to $\mathbf{p}\text{-}\#\text{MISWP}(\Phi)$, it is now straightforward to compute the number of k -cliques in G , since

$$\begin{aligned} \#\text{Clique}_k(G, f) &= \#\text{ColSubInd}(H, G_H, f) \quad \text{by Lemma 4.3} \\ &= \frac{1}{\alpha_H} \#\text{ColStrEmb}(\mathcal{H}^H, G_H, f) \quad \text{by Lemma 4.5} \\ &= \frac{1}{\alpha_H} \#\text{ColStrEmb}(\mathcal{H}, G_H, f) \quad \text{by Lemma 4.4,} \end{aligned}$$

where

$$\alpha_H = |\{\sigma : \sigma \text{ a permutation on } [k], \exists(H, \pi') \in \mathcal{H}^H \text{ such that } \pi' \circ \sigma^{-1} \circ \pi^{-1} \text{ defines an automorphism on } H\}|.$$

Recall that α_H can be computed from H and \mathcal{H} in time depending only on a computable function of k' . Note also that the value of the parameter in the oracle call used to evaluate $\#\text{ColStrEmb}(\mathcal{H}, G_H, f)$ is at most $g(w(\binom{k}{2}))$, which is a computable function of k only. This therefore gives an fpt-Turing reduction from $\mathbf{p}\text{-}\#\text{CLIQUE}$ to $\mathbf{p}\text{-}\#\text{MISWP}(\Phi)$, as required. \square

Decision and approximate counting in the multicolour setting

We now turn our attention to the decision and approximate counting questions. We begin by giving a sufficient condition for $\mathbf{p}\text{-}\#\text{MISWP}(\Phi)$ to be $W[1]$ -complete, which does not make any assumptions about the monotonicity of Φ . For the definition of fpt-reductions for parameterised decision problems, we again refer the reader to [25].

Theorem 4.7. *Let Φ be a family (ϕ_1, ϕ_2, \dots) of functions $\phi_k : \mathcal{L}(k) \rightarrow \{0, 1\}$, such that the function mapping $k \mapsto \phi_k$ is computable, and suppose that $\min^*(\Phi)$ has unbounded treewidth. Then $\mathbf{p}\text{-}\#\text{MISWP}(\Phi)$ is $W[1]$ -complete.*

Proof. We prove $W[1]$ -completeness by means of an fpt-reduction from $\mathbf{p}\text{-}\#\text{CLIQUE}$. Let G be the input in an instance of $\mathbf{p}\text{-}\#\text{CLIQUE}$ with parameter k . By assumption, $\min^*(\Phi)$ has unbounded treewidth, so by Proposition 4.1 there exists $k' \in \mathbb{N}$ (where k' is at most some computable function of k) such that some $H \in \min^*(\phi_{k'})$ contains the $(k \times \binom{k}{2})$ grid as a minor.

Now, we can construct the graph G_H and colouring f in time $h(k) + n^{O(1)}$, where h is a computable function. Moreover, we know from Lemma 4.3 that

$$\#\text{ColSubInd}(H, G_H, f) = \#\text{Clique}_{k'}(G),$$

so in particular the number of colourful subgraphs of G_H (with respect to the colouring f) that are isomorphic to H is non-zero if and only if the number of k -cliques in G is non-zero. Since the parameter, k' , in the instance of $\mathbf{p}\text{-}\#\text{MISWP}(\Phi)$ is bounded by a computable function of k , this gives a fpt-reduction from $\mathbf{p}\text{-}\#\text{CLIQUE}$ to $\mathbf{p}\text{-}\#\text{MISWP}(\Phi)$, as required, completing the proof of the theorem. \square

Together with Corollary 3.11, this immediately implies a dichotomy for the special case in which Φ is a uniformly monotone property (recall from the definition of a uniformly monotone property Φ that the class of labelled graphs $\min(\Phi)$ has bounded treewidth if and only if the class of unlabelled graphs $\min^*(\Phi)$ has bounded treewidth).

Corollary 4.8. *Let Φ be a family (ϕ_1, ϕ_2, \dots) of functions $\phi_k : \{0, 1\}^{\binom{k}{2}} \rightarrow \{0, 1\}$, such that the function mapping $k \mapsto \phi_k$ is computable, and suppose further that Φ is a uniformly monotone property. Then $\mathbf{p}\text{-}\#\text{MISWP}(\Phi)$ is in FPT if there exists an integer t such that all graphs in $\min(\Phi)$ have treewidth at most t ; otherwise $\mathbf{p}\text{-}\#\text{MISWP}(\Phi)$ is $W[1]$ -complete.*

Subject to the assumption that $W[1]$ is not equal to FPT under randomised parameterised reductions, Theorem 4.7 also implies an inapproximability result for the counting problem in the multicolour setting.

Corollary 4.9. Assume that $W[1]$ is not equal to FPT under randomised parameterised reductions, and let Φ be a family (ϕ_1, ϕ_2, \dots) of functions $\phi_k : \mathcal{L}(k) \rightarrow \{0, 1\}$ such that the function $k \mapsto \phi_k$ is computable. Then, if $\min^*(\Phi)$ has unbounded treewidth, there is no FPTRAS for $\mathbf{p}\text{-}\#\text{MISWP}(\Phi)$.

Proof. This follows immediately from Corollary 4.8, together with Corollary 3.17. \square

In the case that Φ is a uniformly monotone property, we also obtain a dichotomy result for the approximability of the counting problem.

Corollary 4.10. Assume that $W[1]$ is not equal to FPT under randomised parameterised reductions. Let Φ be a family (ϕ_1, ϕ_2, \dots) of functions $\phi_k : \mathcal{L}(k) \rightarrow \{0, 1\}$ such that the function $k \mapsto \phi_k$ is computable, and suppose further that Φ is a uniformly monotone property. Then there is an FPTRAS for $\mathbf{p}\text{-}\#\text{MISWP}(\Phi)$ if and only if $\min(\Phi)$ has bounded treewidth.

Proof. The existence of an FPTRAS is precisely Theorem 3.14, and the reverse direction is Corollary 4.9. \square

5. Future directions

There has been significant recent progress in the area of parameterised subgraph counting, but there remain many fundamental open questions, a number of which have already been touched on above. In this section we concentrate on uniformly monotone properties, which formed the main focus of this paper, and identify the key open questions.

We begin by considering the situation for exact counting. The only problems known to admit fpt-algorithms for exact counting are the “trivial” problems where the property can be decided, for any given k -vertex labelled subgraph, by examining edges incident with only a constant-sized subset of the vertices; in the case of uniformly monotone properties, this corresponds precisely to those properties Φ such that $\min(\Phi)$ is a class of graphs with bounded vertex-cover number. It is straightforward to adapt existing algorithms (such as the one given in [16] to solve $\mathbf{p}\text{-}\#\text{SUB}(\mathcal{H})$ when \mathcal{H} is a class of graphs of bounded vertex-cover number) to give polynomial-time algorithms for such trivial properties.

However, the best general hardness result for uniformly monotone problems is Theorem 4.6, which shows that if $\min(\Phi)$ is a class of graphs of unbounded treewidth then $\mathbf{p}\text{-}\#\text{ISWP}(\Phi)$ is $\#W[1]$ -complete. It is known that this hardness result is not the best possible: $\mathbf{p}\text{-}\#\text{CONNECTED INDUCED SUBGRAPH}$ is an example of a uniformly monotone property such that all minimal elements have bounded treewidth, but which is nevertheless known to be $\#W[1]$ -complete [31].

The recent result of Curticapean and Marx [16] is a significant breakthrough towards closing this gap, giving a complexity dichotomy for the problem $\mathbf{p}\text{-}\#\text{SUB}(\mathcal{C})$ (showing in fact that, unless $\text{FPT} \neq \#W[1]$, the trivial problems are the only ones to admit fpt-algorithms in this setting). However, the techniques used seem to be very specific to the situation in which there is a unique minimal element satisfying the property Φ , and so do not transfer easily to the more general case of uniformly monotone properties. Nevertheless, we conjecture that the same dichotomy holds for uniformly monotone properties:

Conjecture 1. Suppose that Φ is a uniformly monotone property, and that there is no constant c such that all graphs in $\min(\Phi)$ have vertex-cover number at most c . Then $\mathbf{p}\text{-}\#\text{ISWP}(\Phi)$ is $\#W[1]$ -complete.

Considering now the situation for decision problems, we were able to give a dichotomy result for the multicolour version (Corollary 4.8), showing that $\mathbf{p}\text{-}\text{MISWP}(\Phi)$ is $W[1]$ -complete if $\min(\Phi)$ is a class of graphs having unbounded treewidth, and belongs to FPT otherwise. However, there remains a gap when we consider the uncoloured version: we know (Corollary 3.11) that $\mathbf{p}\text{-}\text{ISWP}(\Phi)$ is in FPT if $\min(\Phi)$ has bounded treewidth, but the corresponding hardness result does not extend to the uncoloured version. Flum and Grohe [25] conjectured that $\mathbf{p}\text{-}\text{SUB}(\mathcal{C})$ is $W[1]$ -complete whenever \mathcal{C} is a class of graphs having unbounded treewidth, mentioning in particular the case in which $\mathcal{C} = \{K_{k,k} : k \in \mathbb{N}\}$ (the set of all complete bipartite graphs with k vertices in each vertex-class); this specific problem was recently shown to be $W[1]$ -complete by Lin [37], but the general conjecture remains open. We extend this conjecture to encompass all uniformly monotone properties:

Conjecture 2. Suppose that Φ is a uniformly monotone property, and that there is no constant c such that all graphs in $\min(\Phi)$ have treewidth at most c . Then $\mathbf{p}\text{-}\text{ISWP}(\Phi)$ is $W[1]$ -complete.

Observe that, by Corollary 3.17, and subject to the assumption that $W[1]$ is not equal to FPT under randomised parameterised reductions, this would also imply a dichotomy for the existence of an FPTRAS for uniformly monotone properties.

It was noted in Section 3.3 that recent developments in the theory of parameterised random complexity classes [39,11] are based on a more restricted notion of parameterised randomisation than is allowed in the definition of an FPTRAS. An interesting general direction for future research in parameterised counting complexity, therefore, would be to develop an alternative notion of approximability for parameterised counting problems that places corresponding restrictions on the permitted use of randomness; this would allow more precise conclusions to be drawn about the relationship between the approximability of counting problems and the complexity of the corresponding decision problems. Having developed such a framework for approximability, a natural first question to consider would be whether the parameterised subgraph counting problems admitting an FPTRAS also satisfy this stronger notion of approximability.

Acknowledgement

I am very grateful to Mark Jerrum for his helpful comments on several previous drafts of this paper.

References

- [1] Noga Alon, Shai Gutner, Balanced hashing, color coding and approximate counting, in: Jianer Chen, Fedor V. Fomin (Eds.), *Parameterized and Exact Computation*, in: *Lecture Notes in Computer Science*, vol. 5917, Springer, Berlin, Heidelberg, 2009, pp. 1–16.
- [2] Noga Alon, Shai Gutner, Balanced families of perfect hash functions and their applications, *ACM Trans. Algorithms* 6 (3) (2010) 54:1–54:12.
- [3] Noga Alon, Raphael Yuster, Uri Zwick, Color-coding, *J. ACM* 42 (4) (1995) 844–856.
- [4] N. Alon, R. Yuster, U. Zwick, Finding and counting given length cycles, *Algorithmica* 17 (3) (1997) 209–223 (English).
- [5] O. Amini, F. Fomin, S. Saurabh, Counting subgraphs via homomorphisms, *SIAM J. Discrete Math.* 26 (2) (2012) 695–717.
- [6] V. Arvind, Venkatesh Raman, Approximation algorithms for some parameterized counting problems, in: P. Bose, P. Morin (Eds.), *ISAAC 2002*, in: *LNCS*, vol. 2518, Springer-Verlag, Berlin, Heidelberg, 2002, pp. 453–464.
- [7] Luca Becchetti, Paolo Boldi, Carlos Castillo, Aristides Gionis, Efficient semi-streaming algorithms for local triangle counting in massive graphs, in: Ying Li, Bing Liu, Sunita Sarawagi (Eds.), *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Las Vegas, Nevada, USA, August 24–27, 2008, ACM, 2008, pp. 16–24.
- [8] Nadja Betzler, Rene van Bevern, Michael R. Fellows, Christian Komusiewicz, Rolf Niedermeier, Parameterized algorithms for finding connected motifs in biological networks, *IEEE/ACM Trans. Comput. Biol. Bioinformatics* 8 (5) (2011) 1296–1308.
- [9] Andreas Björklund, Thore Husfeldt, Petteri Kaski, Mikko Koivisto, Counting paths and packings in halves, in: Amos Fiat, Peter Sanders (Eds.), *Algorithms – ESA 2009*, in: *Lecture Notes in Computer Science*, vol. 5757, Springer, Berlin, Heidelberg, 2009, pp. 578–586 (English).
- [10] Markus Bläser, Radu Curticapean, Weighted counting of k -matchings is $\#W[1]$ -hard, in: Dimitrios M. Thilikos, Gerhard J. Woeginger (Eds.), *Parameterized and Exact Computation*, in: *Lecture Notes in Computer Science*, vol. 7535, Springer, Berlin, Heidelberg, 2012, pp. 171–181 (English).
- [11] Ankit Chauhan, B.V. Raghavendra Rao, *Parameterized analogues of probabilistic computation*, 2014. arxiv:1409.7790 [cs.CC].
- [12] Hubie Chen, Moritz Müller, The fine classification of conjunctive queries and parameterized logarithmic space complexity, in: *Proceedings of the 32nd Symposium on Principles of Database Systems (New York, NY, USA), PODS'13*, ACM, 2013, pp. 309–320.
- [13] Yijia Chen, Marc Thurley, Mark Weyer, Understanding the complexity of induced subgraph isomorphisms, in: Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfssdóttir, Igor Walukiewicz (Eds.), *Automata, Languages and Programming*, in: *Lecture Notes in Computer Science*, vol. 5125, Springer, Berlin, Heidelberg, 2008, pp. 587–596.
- [14] N. Chiba, T. Nishizeki, Arboricity and subgraph listing algorithms, *SIAM J. Comput.* 14 (1) (1985) 210–223.
- [15] Radu Curticapean, Counting matchings of size k is $\#W[1]$ -hard, in: Fedor V. Fomin, Rūsiņš Freivalds, Marta Kwiatkowska, David Peleg (Eds.), *Automata, Languages, and Programming*, in: *Lecture Notes in Computer Science*, vol. 7965, Springer, Berlin, Heidelberg, 2013, pp. 352–363.
- [16] Radu Curticapean, Dániel Marx, Complexity of counting subgraphs: Only the boundedness of the vertex-cover number counts, in: *55th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2014*, 2014.
- [17] Víctor Dalmau, Peter Jonsson, The complexity of counting homomorphisms seen from the other side, *Theoret. Comput. Sci.* 329 (1–3) (2004) 315–323.
- [18] Rod G. Downey, Michael R. Fellows, Fixed-parameter tractability and completeness II: On completeness for $W[1]$, *Theoret. Comput. Sci.* 141 (1–2) (1995) 109–131.
- [19] Rodney G. Downey, Michael R. Fellows, Kenneth W. Regan, Parameterized circuit complexity and the W hierarchy, *Theoret. Comput. Sci.* 191 (1–2) (1998) 97–115.
- [20] M. Dyer, C. Greenhill, The complexity of counting graph homomorphisms, *Random Structures Algorithms* 17 (3–4) (2000) 260–289. 9th International Conference on Random Structures and Algorithms, POZNAN, POLAND, AUG 02–06, 1999.
- [21] P. Erdős, G. Szekeres, A combinatorial problem in geometry, *Compos. Math.* 2 (1935) 464–470.
- [22] Tommy Färnqvist, Peter Jonsson, Bounded tree-width and csp-related problems, in: Takeshi Tokuyama (Ed.), *Algorithms and Computation*, in: *Lecture Notes in Computer Science*, vol. 4835, Springer, Berlin, Heidelberg, 2007, pp. 632–643 (English).
- [23] M. Fellows, D. Hermelin, F. Rosamond, S. Vialette, On the parameterized complexity of multiple-interval graph problems, *Theoret. Comput. Sci.* 410 (2009) 53–61.
- [24] J. Flum, M. Grohe, The parameterized complexity of counting problems, *SIAM J. Comput.* 33 (4) (2004) 892–922.
- [25] J. Flum, M. Grohe, *Parameterized Complexity Theory*, Springer, 2006.
- [26] Michael L. Fredman, János Komlós, Endre Szemerédi, Storing a sparse table with $O(1)$ worst case access time, *J. ACM* 31 (3) (1984) 538–544.
- [27] B. Gelbord, Graphical techniques in intrusion detection systems, *Information Networking*, 2001, in: *Proceedings. 15th International Conference on*, 2001, pp. 253–258.
- [28] Martin Grohe, The complexity of homomorphism and constraint satisfaction problems seen from the other side, *J. ACM* 54 (1) (2007) 1:1–1:24.
- [29] Martin Grohe, Thomas Schwentick, Luc Segoufin, When is the evaluation of conjunctive queries tractable?, in: *Proceedings of the 33rd ACM Symposium on Theory of Computing*, 2001, pp. 657–666.
- [30] Sylvain Guillemot, Florian Sikora, Finding and counting vertex-colored subtrees, *Algorithmica* 65 (4) (2013) 828–844.
- [31] Mark Jerrum, Kitty Meeks, The parameterised complexity of counting connected subgraphs and graph motifs, *J. Comput. System Sci.* 81 (4) (2015) 702–716.
- [32] Mark Jerrum, Kitty Meeks, Some hard families of parameterised counting problems, *ACM Trans. Comput. Theory* 7 (3) (2015) Article 11, 18 pages. <http://dx.doi.org/10.1145/2786017>.
- [33] Mark Jerrum, Kitty Meeks, The parameterised complexity of counting even and odd induced subgraphs, 2014. arXiv:1410.3375 [math.CO].
- [34] R.M. Karp, M. Luby, Monte-carlo algorithms for enumeration and reliability problems, in: *Proceedings of the 24th Annual IEEE Symposium on Foundations of Computer Science*, 1983, pp. 56–64.
- [35] Subhash Khot, Venkatesh Raman, Parameterized complexity of finding subgraphs with hereditary properties, *Theoret. Comput. Sci.* 289 (2) (2002) 997–1008.
- [36] Mihail N. Kolountzakis, Gary L. Miller, Richard Peng, Charalampos E. Tsourakakis, Efficient triangle counting in large graphs via degree-based vertex partitioning, *Internet Math.* 8 (1–2) (2012) 161–185.
- [37] Bingkai Lin, The parameterized complexity of k -Biclique, in: *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, 2015, pp. 605–615.
- [38] R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, U. Alon, Network motifs: Simple building blocks of complex networks, *Science* 298 (5594) (2002) 824–827.
- [39] Juan Andrés Montoya, Moritz Müller, Parameterized random complexity, *Theory of Comput. Syst.* 52 (2) (2013) 221–270.
- [40] J. Plehn, B. Voigt, Finding minimally weighted subgraphs, in: R. Möhring (Ed.), *Graph-Theoretic Concepts in Computer Science, WG'90*, in: *Lecture Notes in Computer Science*, vol. 484, Springer-Verlag, 1990, pp. 18–29.
- [41] Neil Robertson, P.D Seymour, Graph minors. V. Excluding a planar graph, *J. Combin. Theory Ser. B* 41 (1) (1986) 92–114.
- [42] C. Slot, P. van Emde Boas, On tape versus core an application of space efficient perfect hash functions to the invariance of space, in: *Proceedings of the Sixteenth Annual ACM Symposium on Theory of Computing (New York, NY, USA), STOC'84*, ACM, 1984, pp. 391–400.
- [43] Thomas Schank, Dorothea Wagner, Finding, counting and listing all triangles in large graphs, an experimental study, in: Sotiris E. Nikolettas (Ed.), *Experimental and Efficient Algorithms*, in: *Lecture Notes in Computer Science*, vol. 3503, Springer, Berlin, Heidelberg, 2005, pp. 606–609 (English).
- [44] Jeanette P. Schmidt, Alan Siegel, The spatial complexity of oblivious k -probe hash functions, *SIAM J. Comput.* 19 (5) (1990) 775–786.

- [45] V. Sekar, Y. Xie, D.A. Maltz, M.K. Reiter, H. Zhang, Toward a Framework for Internet Forensic Analysis, in: Third Workshop on Hot Topics in Networking (HotNets-III), 2004.
- [46] S. Staniford-chen, S. Cheung, R. Crawford, M. Dilger, J. Frank, J. Hoagl, K. Levitt, C. Wee, R. Yip, D. Zerkle, GrIDS—a graph based intrusion detection system for large networks, in: Proceedings of the 19th National Information Systems Security Conference, 1996, pp. 361–370.
- [47] Gopalakrishnan Sundaram, Steven S. Skiena, Recognizing small subgraphs, *Networks* 25 (4) (1995) 183–191.
- [48] C.E. Tsourakakis, Fast counting of triangles in large real networks without counting: Algorithms and laws, Data Mining, 2008. ICDM'08, in: Eighth IEEE International Conference on, Dec 2008, pp. 608–617.
- [49] S. Vadhan, The complexity of counting in sparse, regular, and planar graphs, *SIAM J. Comput.* 31 (2) (2001) 398–427.
- [50] L. Valiant, The complexity of enumeration and reliability problems, *SIAM J. Comput.* 8 (3) (1979) 410–421.
- [51] V. Williams, R. Williams, Finding, minimizing, and counting weighted subgraphs, *SIAM J. Comput.* 42 (3) (2013) 831–854.
- [52] Mingji Xia, Peng Zhang, Wenbo Zhao, Computational complexity of counting problems on 3-regular planar graphs, *Theoret. Comput. Sci.* 384 (1) (2007) 111–125. Theory and Applications of Models of Computation.