

Approximate counting CSP seen from the other side*

Andrei A. Bulatov

School of Computing Science, Simon Fraser University, Canada
abulatov@sfu.ca

Stanislav Živný

Department of Computer Science, University of Oxford, UK
standa.zivny@cs.ox.ac.uk

Abstract

In this paper we study the complexity of counting Constraint Satisfaction Problems (CSPs) of the form $\#\text{CSP}(\mathcal{C}, -)$, in which the goal is, given a relational structure \mathbf{A} from a class \mathcal{C} of structures and an arbitrary structure \mathbf{B} , to find the number of homomorphisms from \mathbf{A} to \mathbf{B} . Flum and Grohe showed that $\#\text{CSP}(\mathcal{C}, -)$ is solvable in polynomial time if \mathcal{C} has bounded treewidth [FOCS'02]. Building on the work of Grohe [JACM'07] on decision CSPs, Dalmau and Jonsson then showed that, if \mathcal{C} is a recursively enumerable class of relational structures of bounded arity, then assuming $\text{FPT} \neq \#\text{W}[1]$, there are no other cases of $\#\text{CSP}(\mathcal{C}, -)$ solvable exactly in polynomial time (or even fixed-parameter time) [TCS'04].

We show that, assuming $\text{FPT} \neq \text{W}[1]$ (under randomised parametrised reductions) and for \mathcal{C} satisfying certain general conditions, $\#\text{CSP}(\mathcal{C}, -)$ is not solvable even *approximately* for \mathcal{C} of unbounded treewidth; that is, there is no fixed parameter tractable (and thus also not fully polynomial) randomised approximation scheme for $\#\text{CSP}(\mathcal{C}, -)$. In particular, our condition generalises the case when \mathcal{C} is closed under taking minors.

1 Introduction

The Constraint Satisfaction Problem (CSP) asks to decide the existence of a homomorphism between two given relational structures (or to find the number of such homomorphisms). It has been used to model a vast variety of combinatorial

*An extended abstract of this work will appear in the *Proceedings of the 44th International Symposium on Mathematical Foundations of Computer Science (MFCS'19)*. Andrei Bulatov was supported by an NSERC Discovery grant. Stanislav Živný was supported by a Royal Society University Research Fellowship. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 714532). The paper reflects only the authors' views and not the views of the ERC or the European Commission. The European Union is not liable for any use that may be made of the information contained therein.

problems and has attracted much attention. Since the general CSP is NP-complete (#P-complete in the counting case) and because one needs to model specific computational problems, various restricted versions of the CSP have been considered. More precisely, let \mathcal{C} and \mathcal{D} be two classes of relational structures. In this paper we will assume that structures from \mathcal{C}, \mathcal{D} only have predicate symbols of bounded arity. The *constraint satisfaction problem* (CSP) parametrised by \mathcal{C} and \mathcal{D} is the following computational problem, denoted by $\text{CSP}(\mathcal{C}, \mathcal{D})$: given $\mathbf{A} \in \mathcal{C}$ and $\mathbf{B} \in \mathcal{D}$, is there a homomorphism from \mathbf{A} to \mathbf{B} ? CSPs in which both input structures are restricted have not received much attention (with a notable exception of matrix partitions [19, 20] and assorted graph problems on restricted classes of graphs). However, the two most natural restrictions have been intensively studied over the last two decades. Let $-$ denote the class of all (bounded-arity) relational structures, or, equivalently, indicate that there are no restrictions on the corresponding input structure.

Problems of the form $\text{CSP}(-, \{\mathbf{B}\})$, where \mathbf{B} is a fixed finite relational structure, are known as *nonuniform* or *language-restricted* CSPs [33]. For instance, if $\mathbf{B} = K_3$ is the complete graph on 3 vertices then $\text{CSP}(-, \{\mathbf{B}\})$ is the standard 3-COLOURING problem [27]. The study of nonuniform CSPs has been initiated by Schaefer [38] who considered the case of $\text{CSP}(-, \{\mathbf{B}\})$ for 2-element structures \mathbf{B} . The complexity of $\text{CSP}(-, \{\mathbf{H}\})$, for a fixed graph \mathbf{H} , was studied under the name of \mathbf{H} -colouring by Hell and Nešetřil [32]. General nonuniform CSPs have been studied extensively since the seminal paper of Feder and Vardi [21] who in particular proposed the so-called Dichotomy Conjecture stating that every nonuniform CSP is either solvable in polynomial time or is NP-complete. The complexity of nonuniform CSPs has been resolved only recently in two independent papers by Bulatov [3] and Zhuk [39], which confirmed the dichotomy conjecture of Feder and Vardi and also its algebraic version [4].

CSPs restricted on the other side, that is, of the form $\text{CSP}(\mathcal{C}, -)$, where \mathcal{C} is a fixed (infinite) class of finite relational structures, are known as *structurally-restricted* CSPs. For instance, if $\mathcal{C} = \cup_{k \geq 1} \{K_k\}$ is the class of cliques of all sizes then $\text{CSP}(\mathcal{C}, -)$ is the standard CLIQUE problem [27]. In this case the complexity of CSPs is related to various “width” parameters of the associated class of graphs. For a relational structure \mathbf{A} let $G(\mathbf{A})$ denote the Gaifman graph of \mathbf{A} , that is, the graph whose vertices are the elements of \mathbf{A} , and vertices v, w are connected with an edge whenever v and w occur in the same tuple of some relation of \mathbf{A} . Then $G(\mathcal{C})$ denotes the class of Gaifman graphs of structures from \mathcal{C} , and we refer to the treewidth of $G(\mathbf{A})$ as the treewidth of \mathbf{A} . Dalmau, Kolaitis, and Vardi showed that $\text{CSP}(\mathcal{C}, -)$ is in PTIME if \mathcal{C} has bounded treewidth modulo homomorphic equivalence [10]. Grohe then showed that, assuming $\text{FPT} \neq \text{W}[1]$, there are no other cases of (bounded arity) $\text{CSP}(\mathcal{C}, -)$ solvable in polynomial time (or even fixed-parameter time, where the parameter is the size of the left-hand side structure) [29]. The case of structures with unbounded arity was extensively studied by Gottlob et al. who introduced the concept of bounded hypertree width in an attempt to characterise structurally restricted CSPs solvable in polynomial time [28]. The

search for a right condition is still going on, and the most general structural property that guarantees that $\text{CSP}(\mathcal{C}, -)$ is solvable in polynomial time is fractional hypertree width introduced by Grohe and Marx [30].

An important problem related to the CSP is counting: Given a CSP instance, that is, two relational structures \mathbf{A} and \mathbf{B} , find the number of homomorphisms from \mathbf{A} to \mathbf{B} . We again consider restricted versions of this problem. More precisely, for two classes \mathcal{C} and \mathcal{D} of relational structures, $\#\text{CSP}(\mathcal{C}, \mathcal{D})$ denotes the following computational problem: given $\mathbf{A} \in \mathcal{C}$ and $\mathbf{B} \in \mathcal{D}$, how many homomorphisms are there from \mathbf{A} to \mathbf{B} ? This problem is referred to as a *counting* CSP. Similar to decision CSPs, problems of the form $\#\text{CSP}(-, \mathcal{D})$ and $\#\text{CSP}(\mathcal{C}, -)$ are the two most studied ways to restrict the counting CSP, and the research on these problems follows a similar pattern as their decision counterparts.

For a fixed finite relational structure \mathbf{B} , the complexity of the nonuniform problem $\#\text{CSP}(-, \{\mathbf{B}\})$ was characterised for graphs by Dyer and Greenhill [17] and for 2-element structures by Creignou and Hermann [8]. The complexity of the general nonuniform counting CSPs was resolved by Bulatov [5] and Dyer and Richerby [18]. As in the case of the decision version the complexity of nonuniform counting CSPs is determined by their algebraic properties, and every such CSP is either solvable in polynomial time or is $\#\text{P}$ -complete. These dichotomy results were later extended to the case of weighted counting CSP, for which Cai and Chen obtained a complexity classification of counting CSPs with complex weights [6].

The complexity of counting CSPs with restrictions on the left hand side structures also turns out to be related to treewidth. Flum and Grohe showed that $\#\text{CSP}(\mathcal{C}, -)$ is solvable in polynomial time if \mathcal{C} has bounded treewidth [22]. Dalmau and Jonsson then showed that, assuming $\text{FPT} \neq \#\text{W}[1]$, there are no other cases of (bounded arity) $\#\text{CSP}(\mathcal{C}, -)$ solvable exactly in polynomial time (or, again, even fixed-parameter time) [9]. Note that the result of Dalmau and Jonsson states that the class \mathcal{C} itself has to be of bounded treewidth, while in Grohe’s characterisation of polynomial-time solvable decision CSPs of the form $\text{CSP}(\mathcal{C}, -)$ it is the class of cores of structures from \mathcal{C} that has to have bounded treewidth. To the best of our knowledge there has been no research on counting problems over structures of unbounded arity except for the work of Brault-Baron et al., who showed that the (unbounded arity) structurally-restricted $\#\text{CSP}(\mathcal{C}, -)$ are solvable in polynomial time for the class \mathcal{C} of β -acyclic hypergraphs [2].¹

The results we have mentioned so far concern exact counting; however, many applications of counting problems allow for approximation algorithms as well. For nonuniform CSPs the complexity landscape is much more complicated than the dichotomy results for decision CSPs or exact counting. The analogue of “easily solvable” problems in this case are those that admit a Fully Polynomial Randomised Approximation Scheme (FPRAS): a randomised algorithm that, given an instance and an error tolerance $\varepsilon \in (0, 1)$ returns in time polynomial in the size

¹Brault-Baron et al. [2] show their tractability results for so-called CSPs with default values, which in particular includes $\#\text{CSP}(\mathcal{C}, -)$ as defined here.

of the instance and ε^{-1} a result which is with high probability a multiplicative $(1 + \varepsilon)$ -approximation of the exact solution. The parametrised version of this algorithmic model is known as a Fixed Parameter Tractable Randomised Approximation Scheme (FPTRAS). Beyond counting nonuniform CSPs, however, it was conjectured by Dyer et al. [15] that there is an infinite hierarchy of approximation complexities attainable by such problems. Only a handful of results exist for the approximation complexity of counting nonuniform CSPs. The approximation complexity of $\#\text{CSP}(-, \{\mathbf{B}\})$ for 2-element structures \mathbf{B} was characterised by Dyer et al. [16], where a trichotomy theorem was proved: for every 2-element structure \mathbf{B} the problem $\#\text{CSP}(-, \{\mathbf{B}\})$ either admits an FPRAS, or is interreducible with $\#\text{SAT}$ or with the problem $\#\text{BIS}$ of counting independent sets in bipartite graphs. Apart from this only partial results are known. If \mathbf{B} is a connected graph and $\#\text{CSP}(-, \{\mathbf{B}\})$ does not admit an FPRAS, then Galanis, Goldberg and Jerrum [25] showed that $\#\text{CSP}(-, \{\mathbf{B}\})$ is at least as hard as $\#\text{BIS}$. Also, if every unary relation is a part of \mathbf{B} a complexity classification of $\#\text{CSP}(-, \{\mathbf{B}\})$ can be extracted from the results of Chen et al. [7],² see also [26].

Our Contribution It should be clear by now that the picture painted by the short survey above misses one piece: the approximation complexity of structurally restricted CSPs. This is the main contribution of this paper.

Let \mathcal{C} be a class of bounded-arity relational structures. If the treewidth of \mathcal{C} modulo homomorphic equivalence is unbounded then, by Grohe’s result [29], it is hard to test for the existence of a homomorphism from \mathbf{A} to \mathbf{B} , where $\mathbf{A} \in \mathcal{C}$, for any instance \mathbf{A}, \mathbf{B} of $\text{CSP}(\mathcal{C}, -)$. Using standard techniques (see, e.g., the proof of [34, Proposition 3.16]), this implies, assuming that $\text{FPT} \neq \#\text{W}[1]$ (under randomised parametrised reductions [14]), that there is not an FPTRAS for $\#\text{CSP}(\mathcal{C}, -)$, let alone an FPRAS. Consequently, the tractability boundary for approximate counting of $\#\text{CSP}(\mathcal{C}, -)$ lies between bounded treewidth and bounded treewidth modulo homomorphic equivalence.

As our main result, we show that for \mathcal{C} such that a certain class of graphs (to be defined later) is a subset of $G(\mathcal{C})$, $\#\text{CSP}(\mathcal{C}, -)$ cannot be solved even *approximately* for \mathcal{C} of unbounded treewidth, assuming $\text{FPT} \neq \#\text{W}[1]$ (under randomised parametrised reductions). Before we introduce the classes of graphs we use, we review how the hardness of $\text{CSP}(\mathcal{C}, -)$ or $\#\text{CSP}(\mathcal{C}, -)$ is usually proved.

We follow the hardness proof of Grohe for decision CSPs [29], which was lifted to exact counting CSPs by Dalmau and Jonsson [9]. In fact Grohe’s result had an important precursor [31]. The key idea is a reduction from the parametrised CLIQUE problem to $\text{CSP}(\mathcal{C}, -)$. Let $G = (V, E)$ and k be an instance of the p -CLIQUE problem, where k is the parameter. Broadly speaking, the reduction works as follows. For a class of unbounded treewidth, the Excluded Grid Theo-

²Chen et al. [7] studied the weighted version of $\#\text{CSP}(-, \{\mathbf{B}\})$, and although their result does not provide a complete characterisation of the weighted problem, it allows to determine the complexity of $\#\text{CSP}(-, \{\mathbf{B}\})$ as defined here.

rem of Robertson and Seymour [37] guarantees the existence of the $(k \times \binom{k}{2})$ -grid (as a minor of some structure $\mathbf{A} \in \mathcal{C}$), which is used to encode the existence of a k -clique in G as a certain structure \mathbf{B} . The encoding usually means that G has a k -clique if and only if there is a homomorphism from \mathbf{A} to \mathbf{B} whose image covers a copy of the grid built in \mathbf{B} . For decision CSPs, the correctness of the reduction — that there are no homomorphisms from \mathbf{A} to \mathbf{B} not satisfying this condition — is achieved by dealing with coloured grids [31] or by dealing with structures whose cores have unbounded treewidth (with another complication caused by minor maps) [29]. For the complexity of exact counting CSPs, the correctness of the reduction [9] is achieved by employing interpolation or the inclusion-exclusion principle, a common tool in exact counting.

None of these two methods can be applied to approximate solving $\#\text{CSP}(\mathcal{C}, -)$. We cannot assume that the class of cores of \mathcal{C} has unbounded treewidth, because then by [29] even the decision problem cannot be solved in polynomial time, which immediately rules out the existence of an FPRAS. Interpolation techniques such as the inclusion-exclusion principle are also well known to be incompatible with approximate counting. The standard tool in approximate counting to achieve the same goal of prohibiting homomorphisms except ones from a certain restricted type, is to use gadgets to amplify the number of homomorphisms of the required type. We give a reduction from p - $\#\text{CLIQUE}$ to $\#\text{CSP}(\mathcal{C}, -)$ by using “fan-grids”, formally introduced in Section 3.3. Unfortunately, due to the delicate nature of approximation preserving reductions, we cannot use minors and minor maps and have to assume that “fan-grids” themselves are present in $G(\mathcal{C})$. (In Section 5, we will briefly discuss how a weaker assumption can be used to obtain the same result.) By the Excluded Grid Theorem [37], if \mathcal{C} is *closed under taking minors*, then $G(\mathcal{C})$ contains all the fan-grids (details are given in Section 3.3 and in particular in Lemma 4). Thus, the classes \mathcal{C} for which we establish the hardness of $\#\text{CSP}(\mathcal{C}, -)$ includes the classes \mathcal{C} that are closed under taking minors.³

2 Preliminaries

\mathbb{N} denotes the set of positive integers. For every $n \in \mathbb{N}$, we let $[n] = \{1, \dots, n\}$.

2.1 Relational Structures and Homomorphisms

A *relational signature* is a finite set τ of relation symbols R , each with a specified arity $\text{ar}(R)$. A *relational structure* \mathbf{A} over a relational signature τ (or a τ -structure, for short) is a finite universe A together with one relation $R^{\mathbf{A}} \subseteq A^{\text{ar}(R)}$ for each

³We remark that the hardness for \mathcal{C} closed under taking minors follows from Grohe’s classification [29] of decision CSPs. Indeed, for \mathcal{C} of unbounded treewidth, the Excluded Grid Theorem [37] gives grids of arbitrary sizes. Since every planar graph is a minor of some grid [11], \mathcal{C} contains all planar graphs. As there exist graphs of arbitrary large treewidth that are also minimal with respect to homomorphic equivalence, Grohe’s result gives W[1]-hardness of $\text{CSP}(\mathcal{C}, -)$ and hence $\#\text{CSP}(\mathcal{C}, -)$ cannot have an FPRAS/FPTRAS.

symbol $R \in \tau$. The size $\|\mathbf{A}\|$ of a relational structure \mathbf{A} is defined as

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

Let R be a binary relational symbol. We will sometimes view graphs as $\{R\}$ -structures.

A *homomorphism* from a relational τ -structure \mathbf{A} (with universe A) to a relational τ -structure \mathbf{B} (with universe B) is a mapping $\varphi : A \rightarrow B$ such that for all $R \in \tau$ and all tuples $\mathbf{x} \in R^{\mathbf{A}}$ we have $\varphi(\mathbf{x}) \in R^{\mathbf{B}}$.

Two structures \mathbf{A} and \mathbf{B} are *homomorphically equivalent* if there is a homomorphism from \mathbf{A} to \mathbf{B} and a homomorphism from \mathbf{B} to \mathbf{A} .

Let \mathcal{C} be a class of relational structures. We say that \mathcal{C} has *bounded arity* if there is a constant $r \geq 1$ such that for every τ -structure $\mathbf{A} \in \mathcal{C}$ and $R \in \tau$, we have that $\text{ar}(R) \leq r$.

2.2 Treewidth and Minors

The notion of treewidth, introduced by Robertson and Seymour [36], is a well-known measure of the tree-likeness of a graph [11]. Let $G = (V(G), E(G))$ be a graph. A *tree decomposition* of G is a pair (T, β) where $T = (V(T), E(T))$ is a tree and β is a function that maps each node $t \in V(T)$ to a subset of $V(G)$ such that

1. $V(G) = \bigcup_{t \in V(T)} \beta(t)$,
2. for every $u \in V(G)$, the set $\{t \in V(T) \mid u \in \beta(t)\}$ induces a connected subgraph of T , and
3. for every edge $\{u, v\} \in E(G)$, there is a node $t \in V(T)$ with $\{u, v\} \subseteq \beta(t)$.

The *width* of the decomposition (T, β) is $\max\{|\beta(t)| \mid t \in V(T)\} - 1$. The *treewidth* $tw(G)$ of a graph G is the minimum width over all its tree decompositions.

Let \mathbf{A} be a relational structure over relational signature τ . The *Gaifman graph* (also known as *primal graph*) of \mathbf{A} , denoted by $G(\mathbf{A})$, is the graph whose vertex set is the universe of \mathbf{A} and whose edges are the pairs (u, v) for which there is a tuple \mathbf{x} and a relation symbol $R \in \tau$ such that u, v appear in \mathbf{x} and $\mathbf{x} \in R^{\mathbf{A}}$.

Let \mathcal{C} be a class of relational structures. We say that \mathcal{C} has *bounded treewidth* if there exists $w \geq 1$ such that $tw(\mathbf{A}) = tw(G(\mathbf{A})) \leq w$ for every $\mathbf{A} \in \mathcal{C}$. We say that \mathcal{C} has *bounded treewidth modulo homomorphic equivalence* if there exists $w \geq 1$ such that every $\mathbf{A} \in \mathcal{C}$ is homomorphically equivalent to \mathbf{A}' with $tw(\mathbf{A}') \leq w$.

A graph H is a *minor* of a graph G if H is isomorphic to a graph that can be obtained from a subgraph of G by contracting edges (for more details, see, e.g., [11]).

For $k, \ell \geq 1$, the $(k \times \ell)$ -grid is the graph with the vertex set $[k] \times [\ell]$ and an edge between (i, j) and (i', j') iff $|i - i'| + |j - j'| = 1$. Treewidth and minors are intimately connected via the celebrated Excluded Grid Theorem of Robertson and Seymour.

Theorem 1 ([37]). *For every k there exists a $w(k)$ such that the $(k \times k)$ -grid is a minor of every graph of treewidth at least $w(k)$.*

Let \mathcal{C} be a class of relational structures. We say that \mathcal{C} is *closed under taking minors* if for every $\mathbf{A} \in \mathcal{C}$ and for every minor H of $G(\mathbf{A})$, there is a structure $\mathbf{A}' \in \mathcal{C}$ such that $G(\mathbf{A}')$ is isomorphic to H .

3 Counting CSP

3.1 Exact Counting CSP

Let \mathcal{C} be a class of relational structures. We will be interested in the computational complexity of the following problem.

Name: $\#\text{CSP}(\mathcal{C}, -)$

Input: Two relational structures \mathbf{A} and \mathbf{B} over the same signature with $\mathbf{A} \in \mathcal{C}$.

Output: The number of homomorphisms from \mathbf{A} to \mathbf{B} .

We say that $\#\text{CSP}(\mathcal{C}, -)$ is in FP, the class of function problems solvable in *polynomial time*, if there is a deterministic algorithm that solves any instance \mathbf{A}, \mathbf{B} of $\#\text{CSP}(\mathcal{C}, -)$ in time $(\|\mathbf{A}\| + \|\mathbf{B}\|)^{O(1)}$.

We will also consider the parametrised version of $\#\text{CSP}(\mathcal{C}, -)$.

Name: $p\text{-}\#\text{CSP}(\mathcal{C}, -)$

Input: Two relational structures \mathbf{A} and \mathbf{B} over the same signature with $\mathbf{A} \in \mathcal{C}$.

Parameter: $\|\mathbf{A}\|$.

Output: The number of homomorphisms from \mathbf{A} to \mathbf{B} .

We say that $p\text{-}\#\text{CSP}(\mathcal{C}, -)$ is in FPT, the class of problems that are *fixed-parameter tractable*, if there is a deterministic algorithm that solves any instance \mathbf{A}, \mathbf{B} of $p\text{-}\#\text{CSP}(\mathcal{C}, -)$ in time $f(\|\mathbf{A}\|) \cdot \|\mathbf{B}\|^{O(1)}$, where $f : \mathbb{N} \rightarrow \mathbb{N}$ is an arbitrary computable function.

The class W[1], introduced in [12], can be seen as an analogue of NP in parameterised complexity theory. Proving W[1]-hardness of a problem (under a parametrised reduction which may be randomised), is a strong indication that the problem is not solvable in fixed-parameter time as it is believed that $\text{FPT} \neq \text{W}[1]$. For counting problems, $\#\text{W}[1]$ is the parametrised analogue of $\#\text{P}$. Similarly to the belief that $\text{FP} \neq \#\text{P}$, it is believed that $\text{FPT} \neq \#\text{W}[1]$. We refer the reader to [24] for the definitions of W[1] and $\#\text{W}[1]$, and for more details on parameterised complexity in general.

Dalmau and Jonsson established the following result.

Theorem 2 ([9]). *Assume $FPT \neq \#W[1]$ under parametrised reductions. Let \mathcal{C} be a recursively enumerable class of relational structures of bounded arity. Then, the following are equivalent:*

1. $\#CSP(\mathcal{C}, -)$ is in FP .
2. $p\text{-}\#CSP(\mathcal{C}, -)$ is in FPT .
3. \mathcal{C} has bounded treewidth.

The following problem is an example of a $\#W[1]$ -hard problem, as established by Flum and Grohe [23].

Name: $p\text{-}\#CLIQUE$

Input: A graph G and $k \in \mathbb{N}$.

Parameter: k .

Output: The number of cliques of size k in G .

Note that $p\text{-}\#CLIQUE$ can be modelled as $p\text{-}\#CSP(\mathcal{C}, -)$ if we set \mathcal{C} to be the set of cliques of all possible sizes. The decision version of $p\text{-}\#CLIQUE$ was shown to be $W[1]$ -hard by Downey and Fellows [13].

Name: $p\text{-}CLIQUE$

Input: A graph G and $k \in \mathbb{N}$.

Parameter: k .

Output: Decide if G contains a clique of size k .

3.2 Approximate Counting CSP

In view of our complete understanding of the exact complexity of $\#CSP(\mathcal{C}, -)$ for \mathcal{C} of bounded arity (cf. Theorem 2), we will be interested in *approximation* algorithms for $\#CSP(\mathcal{C}, -)$. In particular, are there any new classes \mathcal{C} of bounded arity for which the problem $\#CSP(\mathcal{C}, -)$ can be solved efficiently (if only approximately)? We will provide a partial answer to this question (cf. Theorem 3): for certain general bounded-arity classes \mathcal{C} (which include classes that are *closed under taking minors*), the answer is no!

The notion of efficiency for approximate counting is that of a fully polynomial randomised approximation scheme [35] and its parametrised analogue, a fixed parameter tractable randomised approximation scheme, originally introduced by Arvind and Raman [1]. We now define both concepts.

A *randomised approximation scheme* (RAS) for a function $f : \Sigma^* \rightarrow \mathbb{N}$ is a randomised algorithm that takes as input $(x, \varepsilon) \in \Sigma^* \times (0, 1)$ and produces as output an integer random variable X satisfying the condition $\Pr(|X - f(x)| \leq$

$\varepsilon f(x) \geq 3/4$. A RAS for a counting problem is called *fully polynomial* (FPRAS) if on input of size n it runs in time $p(n, \varepsilon^{-1})$ for some fixed polynomial p . A RAS for a parametrised counting problem is called *fixed parameter tractable* (FPTRAS) if on input of size n with parameter k it runs in time $f(k) \cdot p(n, \varepsilon^{-1})$, where p is a fixed polynomial and f is an arbitrary computable function.

To compare approximation complexity of (parametrised) counting problems two types of reductions are used. Suppose $f, g : \Sigma^* \rightarrow \mathbb{N}$. An *approximation preserving reduction* (AP-reduction) [15] from f to g is a probabilistic oracle Turing machine M that takes as input a pair $(x, \varepsilon) \in \Sigma^* \times (0, 1)$, and satisfies the following three conditions: (i) every oracle call made by M is of the form (w, δ) , where $w \in \Sigma^*$ is an instance of g , and $0 < \delta < 1$ is an error bound satisfying $\delta^{-1} \leq \text{poly}(|x|, \varepsilon^{-1})$; (ii) the TM M meets the specification for being a randomised approximation scheme for f whenever the oracle meets the specification for being a randomised approximation scheme for g ; and (iii) the running time of M is polynomial in $|x|$ and ε^{-1} .

Similar to [34] we also use the parametrised version of AP-reductions. Again, let $f, g : \Sigma^* \rightarrow \mathbb{N}$. A *parametrised approximation preserving reduction* (parametrised AP-reduction) from f to g is a probabilistic oracle Turing machine M that takes as input a triple $(x, k, \varepsilon) \in \Sigma^* \times (0, 1)$, and satisfies the following three conditions: (i) every oracle call made by M is of the form (w, k', δ) , where $w \in \Sigma^*$ is an instance of g , $k' \leq h(k)$ for some computable function h , and $0 < \delta < 1$ is an error bound satisfying $\delta^{-1} \leq \text{poly}(|x|, \varepsilon^{-1})$; (ii) the TM M meets the specification for being a randomised approximation scheme for f whenever the oracle meets the specification for being a randomised approximation scheme for g ; and (iii) M is fixed-parameter tractable with respect to k and polynomial in $|x|$ and ε^{-1} .

3.3 Main Result

The following concept plays a key role in this paper. Let $k, r, \ell_1, \ell_2 \in \mathbb{N}$. The *fan-grid* $L(k, r, \ell_1, \ell_2)$ is a graph with vertex set $L_1 \cup L_2$, where $L_1 = \{(i, p) \mid i \in [k], p \in [r]\}$, $L_2 = M_1 \cup \dots \cup M_{12}$, where M_1, \dots, M_{12} are disjoint and $|M_i| = \ell_1$ for $i \in [4]$, and $|M_i| = \ell_2$ for $i \in \{5, \dots, 12\}$. Vertices from L_1 will be called *grid vertices*. Vertices $u_1 = (1, 1)$, $u_2 = (1, r)$, $u_3 = (k, 1)$, $u_4 = (k, r)$, $u_5 = (1, 3)$, $u_6 = (1, r - 3)$, $u_7 = (k, 3)$, $u_8 = (k, r - 3)$, $u_9 = (3, 1)$, $u_{10} = (4, r)$, $u_{11} = (k - 2, 1)$, $u_{12} = (k - 3, r)$ will be called *fan vertices*, and u_1, u_2, u_3, u_4 will be called *corner vertices*. The edges of the fan grid are as follows: $(i, p)(i', p')$ for $|i - i'| + |p - p'| = 1$, and wu_i for each $w \in M_i$ and $i \in [12]$, see Figure 1.

We call a class \mathcal{C} of relational structures of bounded arity a *fan class* if either \mathcal{C} has bounded treewidth or for any parameters $k, r, \ell_1, \ell_2 \in \mathbb{N}$ we have that $G(\mathcal{C})$ contains the fan-grid $L(k, r, \ell_1, \ell_2)$.

The following is our main result.

Theorem 3 (Main). *Assume $\text{FPT} \neq \text{W}[1]$ under randomised parametrised reductions. Let \mathcal{C} be a recursively enumerable class of relational structures of bounded arity. If \mathcal{C} is a fan class then the following are equivalent:*




Figure 1: Fan-grid. Fan vertices are shown by larger dots.

1. $\#\text{CSP}(\mathcal{C}, -)$ is polynomial time solvable.
2. $\#\text{CSP}(\mathcal{C}, -)$ admits an FPRAS.
3. $p\text{-}\#\text{CSP}(\mathcal{C}, -)$ admits an FPTRAS.
4. \mathcal{C} has bounded treewidth.

Let \mathcal{C} be a recursively enumerable class of relational structures of bounded arity and closed under taking minors. We claim that \mathcal{C} is a fan class and thus Theorem 3 applies to such \mathcal{C} . For this we need Theorem 1. In particular, for any $k, r, \ell_1, \ell_2 \in \mathbb{N}$, if \mathcal{C} is not of bounded treewidth then, by Theorem 1, $G(\mathcal{C})$ contains an $(s \times s)$ -grid, where $s = \max(k + 2\ell_1, r + 2\ell_2)$, and thus also a $(k + 2\ell_1) \times (r + 2\ell_2)$ -grid. The following simple lemma then shows that fan-grids are minors of grids (of appropriate size).

Lemma 4. $L(k, r, \ell_1, \ell_2)$ is a minor of $(t \times t')$ -grid, where $t = k + 2\ell_1$, $t' = r + 2\ell_2$.

Proof. Take the subgraph of the $(t \times t')$ -grid as shown in Figure 2 and contract the paths shown with thicker edges. \square

4 Proof of Theorem 3

Conditions (1) and (4) in Theorem 3 are equivalent by [9]. Implications “(1) \Rightarrow (2) \Rightarrow (3)” are obvious; implication “(4) \Rightarrow (1)” is by the standard treewidth-based dynamic programming for exact counting. Our main contribution is to prove the “(3) \Rightarrow (4)” implication.

4.1 Construction

Let $G = (V, E)$ be a graph with $n = |V|$ and $m = |E|$. Let $k \in \mathbb{N}$. We construct a graph $H(G, k, W_1, W_2)$ for $W_1, W_2 > 2(n + m)$ as follows. Let $r = \binom{k}{2}$ and let ϱ be a correspondence between $[r]$ and the set of 2-element sets $\{\{i, j\} \mid i, j \in$




Figure 2: Fan-grid as a minor. In the subgraph of the bigger grid shown in solid lines contract the thick edges.

$[k], i \neq j\}$. For $i \in [k]$ and $p \in [r]$, we write $i \in p$ rather than $i \in \varrho(p)$. The vertex set of $H(G, k, W_1, W_2)$ is the union of two sets $H_1 \cup H_2$, defined by

$$H_1 = \{(v, e, i, p) \mid v \in V, e \in E, \text{ and } v \in e \iff i \in p\},$$

$$H_2 = K_1 \cup \dots \cup K_{12},$$

where K_1, \dots, K_{12} are disjoint and $|K_i| = W_1$ for $i \in [4]$, $|K_i| = W_2$ for $i \in \{5, \dots, 12\}$.

As in fan-grids, vertices of the form $(v, e, 1, 1), (v, e, 1, r), (v, e, k, 1), (v, e, k, r), (v, e, 1, 3), (v, e, 1, r-3), (v, e, k, 3), (v, e, k, r-3), (v, e, 3, 1), (v, e, 4, r), (v, e, k-2, 1), (v, e, k-3, r)$ will be called *fan vertices*, and vertices of the form $(v, e, 1, 1), (v, e, 1, r), (v, e, k, 1), (v, e, k, r)$ will be called *corner vertices*.

The edge set of $H(G, k, W_1, W_2)$ consists of the following pairs:

- $(v, e, i, p)(v', e, i', p)$ such that $|i - i'| = 1$;
- $(v, e, i, p)(v, e', i, p')$ such that $|p - p'| = 1$;
- $u(v, e, 1, 1)$ for $u \in S_1 \subseteq K_1$ and $(v, e, 1, 1) \in H_1$, where S_1 is an arbitrary subset of K_1 whose cardinality is such that the degree of $(v, e, 1, 1)$ is exactly W_1 ;
 similarly, $u(v, e, 1, r), u(v, e, k, 1), u(v, e, k, r), u(v, e, 1, 3), u(v, e, 1, r-3), u(v, e, k, 3), u(v, e, k, r-3), u(v, e, 3, 1), u(v, e, 4, r), u(v, e, k-2, 1), u(v, e, k-3, r)$ for $u \in S_j \subseteq K_j$ (for $j = 2, \dots, 12$ in this order) and $(v, e, 1, r), (v, e, k, 1), (v, e, k, r), (v, e, 1, 3), (v, e, 1, r-3), (v, e, k, 3), (v, e, k, r-3), (v, e, 3, 1), (v, e, 4, r), (v, e, k-2, 1), (v, e, k-3, r) \in H_1$, where S_2, \dots, S_{12} are arbitrary subsets whose cardinality is such that the degree of $(v, e, 1, r), (v, e, k, 1), (v, e, k, r)$ is exactly W_1 and the degree of the remaining vertices from the list is exactly W_2 .

We study homomorphisms from $L(k, r, \ell_1, \ell_2)$ to $H(G, k, W_1, W_2)$. A homomorphism $\varphi : L(k, r, \ell_1, \ell_2) \rightarrow H(G, k, W_1, W_2)$ is said to be corner-to-corner (or c-c for short) if

$$\varphi(1, 1), \varphi(1, r), \varphi(k, 1), \varphi(k, r) \in \{(v, e, 1, 1), (v, e, 1, r), (v, e, k, 1), (v, e, k, r) \mid v \in V, e \in E\}.$$

Homomorphism φ is called identity (skew identity) if $\varphi(i, p) \in \{(v, e, i, p) \mid v \in V, e \in E\}$ (respectively, $\varphi(i, p) \in \{(v, e, k - i + 1, p) \mid v \in V, e \in E\}$) for all $i \in [k]$ and $p \in [r]$. Sometimes we will abuse the terminology and call a (skew) identity homomorphism the restriction of such homomorphism to L_1 (the set of grid vertices).

We define the *weight* of a homomorphism φ from $L(k, r, \ell_1, \ell_2)$ restricted to L_1 (the set of grid vertices) to $H(G, k, W_1, W_2)$ as the number of extensions of φ to a homomorphism from $L(k, r, \ell_1, \ell_2)$.

4.2 Weights of Homomorphisms

We start with a simple lemma.

Lemma 5. *The weight of an identity or skew identity homomorphism is $W_1^{4\ell_1} W_2^{8\ell_2}$.*

Proof. The images of grid vertices (the set L_1) under identity and skew identity homomorphisms are fixed, while vertices from L_2 can be mapped to any neighbour of the corresponding fan vertex independently. Since the degree of a corner vertex (v, e, i, p) with $i \in \{1, k\}$ and $p \in \{1, r\}$ is W_1 , and the degree of any other fan vertex is W_2 , the result follows. \square

The next lemma, which will be proved using Lemma 5, is essentially [9, Lemma 3.1] adapted to our setting, which in turn builds on [29, Lemma 4.4].

Lemma 6. *Let N be the number of k -cliques in G . Then the total weight of identity and skew identity homomorphisms is $2NW_1^{4\ell_1} W_2^{8\ell_2} k!$.*

Proof. We will show that the total weight of identity homomorphisms is $NW_1^{4\ell_1} W_2^{8\ell_2} k!$. The same argument works for skew identity homomorphisms. First we give a description of all identity homomorphisms. Let v_1, \dots, v_k be the vertex set of a k -clique in G . For $p \in [r]$ with $\varrho(p) = \{a, b\}$, let $e_p = v_a v_b$ be the edge in G between v_a and v_b . We define $\varphi_{v_1, \dots, v_k} : L_1 \rightarrow H_1$ by

$$\varphi_{v_1, \dots, v_k}((i, p)) = (v_i, e_p, i, p)$$

for every $i \in [k]$ and $p \in [r]$.

We will need two claims; the first one follows directly from the definition.

Claim 1. $\varphi_{v_1, \dots, v_k}$ is an identity homomorphism from L_1 to H_1 .

Claim 2. If φ is an identity homomorphism from L_1 to H_1 then $\varphi = \varphi_{v_1, \dots, v_k}$ for some vertex set v_1, \dots, v_k of a k -clique in G .

Proof of Claim 2. Let φ be an identity homomorphism from L_1 to H_1 .

For every $i \in [k]$ and $p \in [r]$, we have $\varphi((i, p)) = (v, e, i, p)$ for some $v \in V$ and $e \in E$ with $v \in e \iff i \in p$. Let $\varphi((i', p)) = (v', e', i', p)$. We claim that (A) $e = e'$. We prove (A) for $i' = i + 1$, the rest follows by induction. Since φ is a homomorphism and $(i, p)(i', p)$ is an edge in $L(k, r, \ell_1, \ell_2)$, there is an edge $(v, e, i, p)(v', e', i', p)$ in $H(G, k, W_1, W_2)$. The definition of edges in $H(G, k, W_1, W_2)$ implies that $e = e'$. Similarly, let $\varphi((i, p)) = (v, e, i, p)$ and $\varphi((i, p')) = (v', e', i, p')$. We claim that (B) $v = v'$. For $p' = p + 1$, this again follows from the assumption that φ is a homomorphism and the definition of edges in $H(G, K, W_1, W_2)$; a simple induction establishes (B) for arbitrary values $p, p' \in [r]$.

Together, claims (A) and (B) imply that there are vertices $v_1, \dots, v_k \in V$ and edges $e_1, \dots, e_r \in E$ such that for all $i \in [k]$ and $p \in [r]$ we have $\varphi((i, p)) = (v_i, e_p, i, p)$. Since $\varphi((i, p)) \in H_1$, we have $v_i \in e_p \iff i \in p$. Hence v_1, \dots, v_k forms a k -clique in G . **(End of proof of Claim 2.)**

Claims 1 and 2 give us a complete description of identity homomorphisms from L_1 to H_1 : a mapping φ from L_1 to H_1 is an identity homomorphism if and only if $\varphi = \varphi_{v_1, \dots, v_k}$ for some vertex set v_1, \dots, v_k of a k -clique in G . Hence, the number of such mappings is the number of k -cliques in G multiplied by $k!$. By Lemma 5, each identity homomorphism can be extended in $W_1^{4\ell_1} W_2^{8\ell_2}$ distinct ways to a homomorphism from $L(k, r, \ell_1, \ell_2)$ to $H(G, k, W_1, W_2)$. \square

We will frequently use the following simple observation.

Observation 7. *Let φ be a homomorphism from a bipartite graph G to a graph H . If vertices u, v are of distance m in G then $\varphi(u), \varphi(v)$ are of distance at most m in H and the parity of the distances is the same.*

Next we establish an upper bound on the total weight of homomorphisms that are neither identity nor skew identity.

Lemma 8. *Let $G = (V, E)$ has $n = |V|$ vertices and $m = |E|$ edges, let $k = 4k'$ for some k' , and let $T = \log_{W_2} W_1$. If*

$$\ell_1 > \frac{8T\ell_2}{T-1},$$

then the total weight of homomorphisms that are neither identity nor skew identity is at most

$$W_1^{4\ell_1} W_2^{6\ell_2} (2n + m)^{2\ell_2} \cdot (4W_1 + 8W_2 + nmkr)^{kr}.$$

The key ideas in the proof of Lemma 8 are the following: Firstly, we show that c-c homomorphisms dominate non-c-c homomorphisms. Secondly, using crucially the special structure of fan grids and our choice of k being a multiple of four, we establish an upper bound on any c-c homomorphism that is neither identity nor skew identity. Finally, we give an upper bound on the number of all homomorphisms. These three ingredients together allows us to establish the required bound.

Proof. We prove this lemma in two steps. First, in Claims 1 and 2, we upper bound the weight of a homomorphism that is not identity or skew identity. Second, in Claim 3, we upper bound the number of such homomorphisms.

Claim 1. The weight of any c-c homomorphism is greater than the weight of any non c-c homomorphism.

Proof of Claim 1: The weight of a c-c homomorphism is lower bounded by $W_1^{4\ell_1}$, and the weight of any non c-c homomorphism is upper bounded by $W_1^{3\ell_1+8\ell_2}W_2^{\ell_1}$. Take the logarithm base W_2 of the two numbers. We need to show that

$$4T\ell_1 > T(3\ell_1 + 8\ell_2) + \ell_1,$$

or, equivalently,

$$T\ell_1 > 8T\ell_2 + \ell_1,$$

which is equivalent to the condition

$$\ell_1 > \frac{8T\ell_2}{T-1}$$

of the lemma. **(End of proof of Claim 1.)**

Claim 2. Let φ be a c-c homomorphism that is neither identity nor skew identity. Then its weight does not exceed $W_1^{4\ell_1}W_2^{6\ell_2}(2n+m)^{2\ell_2}$.

Proof of Claim 2: We consider several cases. First observe some symmetries in c-c homomorphisms. If ψ is the mapping of H_1 (the “grid” part of $H(G, k, W_1, W_2)$) mapping (v, e, i, p) to $(v, e, k-i+1, p)$, then the weight of $\psi \circ \varphi$ equals that of φ . Thus we may assume $\varphi(1, 1) \in \{(v, e, 1, 1), (v, e, 1, r) \mid v \in V, e \in E\}$, which gives CASE 1 and CASE 2 below, respectively. Note that by the assumption that k is a multiple of four, both $k-1$ and $r-1$, where $r = \binom{k}{2}$, are odd.

CASE 1. $\varphi(1, 1) = (v, e, 1, 1)$ for some $v \in V, e \in E$.

Since $(k, 1)$ is at distance $k-1$ from $(1, 1)$, by Observation 7, $\varphi(k, 1)$ is at odd distance not exceeding $k-1$ from $\varphi(1, 1)$. As φ is c-c, there is only one possibility $\varphi(k, 1) = (v', e', k, 1)$ for some $v' \in V, e' \in E$. Similarly, as $(1, r)$ is at odd distance from $(1, 1)$ and $\varphi(1, r)$ is a corner vertex, by Observation 7 it suffices to consider only two cases for $\varphi(1, r)$.

CASE 1.1. $\varphi(1, r) = (v'', e'', 1, r)$ for some $v'' \in V, e'' \in E$.

Since (k, r) is at distance $k-1$ from $(1, r)$, by Observation 7, $\varphi(k, r)$ is at odd distance not exceeding $k-1$ from $\varphi(1, r)$. As φ is c-c and we assume that $\varphi(1, r) = (v'', e'', 1, r)$, there is only one possibility $\varphi(k, r) = (v''', e''', k, r)$ for some $v''' \in V, e''' \in E$. This means that φ is identity, a contradiction.

CASE 1.2. $\varphi(1, r) = (v'', e'', k, 1)$ for some $v'' \in V, e'' \in E$.

As in CASE 1.1, $\varphi(k, r) = (v''', e''', 1, 1)$ for some $v''' \in V, e''' \in E$. In detail, since (k, r) is at distance $k-1$ from $(1, r)$, by Observation 7, $\varphi(k, r)$ is at

odd distance not exceeding $k - 1$ from $\varphi(1, r)$. As φ is c-c and we assume that $\varphi(1, r) = (v'', e'', k, 1)$, there is only one possibility $\varphi(k, r) = (v''', e''', 1, 1)$.

Since $(1, r), (2, r), \dots, (k, r)$ is the only shortest path from $(1, r)$ to (k, r) , homomorphism φ maps this path to $(v_k, e_k, k, 1), (v_{k-1}, e_{k-1}, k-1, 1), \dots, (v_1, e_1, 1, 1)$ for some $v_1, \dots, v_k \in V, e_1, \dots, e_k \in E$ (in fact, we can claim that $e''' = e_1 = \dots = e_k = e''$, but we do not need this). In particular, $\varphi(4, r) = (v_{k-3}, e_{k-3}, k-3, 1)$ and $\varphi(k-3, r) = (v_4, e_4, 4, 1)$; that is, these two vertices are mapped to non-fan vertices. Since both $(v_{k-3}, e_{k-3}, k-3, 1)$ and $(v_4, e_4, 4, 1)$ have at most $2n + m$ neighbours, the weight of φ is at most $W_1^{4\ell_1} W_2^{6\ell_2} (2n + m)^{2\ell_2}$.

CASE 2. $\varphi(1, 1) = (v, e, 1, r)$ for some $v \in V, e \in E$.

This case is symmetric to CASE 1 so we do not give full details. Using Observation 7 and the assumption that φ is c-c, we get $\varphi(k, 1) = (v', e', k, r)$ for some $v' \in V, e' \in E$. Also, we get that $\varphi(1, r) = (v'', e'', 1, 1)$ or $\varphi(1, r) = (v'', e'', k, 1)$ for some $v'' \in V, e'' \in E$. In the former case, as in CASE 1.1 we get that φ necessarily is skew identity, which is a contradiction. In the latter case, similarly to CASE 1.2, we get that $\varphi(3, 1) = (v_{k-2}, e_{k-2}, k-2, r), \varphi(k-2, 1) = (v_3, e_3, 3, r)$ for some $v_{k-2}, v_3 \in V, e_{k-2}, e_3 \in E$. Since $(v_{k-2}, e_{k-2}, k-2, r)$ and $(v_3, e_3, 3, r)$ are not fan vertices, as in CASE 1.2, the weight of φ does not exceed $W_1^{4\ell_1} W_2^{6\ell_2} (2n + m)^{2\ell_2}$. **(End of proof of Claim 2.)**

Claim 3. The number of homomorphisms of the $(k \times r)$ -grid to $H(G, k, W_1, W_2)$ is upper bounded by

$$(4W_1 + 8W_2 + nmkr)^{kr}.$$

Proof of Claim 3: Since $H(G, k, W_1, W_2)$ has no more than $4W_1 + 8W_2 + nmkr$ vertices and the $(k \times r)$ -grid has kr vertices, the claim follows. **(End of proof of Claim 3.)**

By Claims 1 and 2, the maximum weight of a homomorphism that is not identity or skew identity is $W_1^{4\ell_1} W_2^{6\ell_2} (2n + m)^{2\ell_2}$. By Claim 3, there are at most $(4W_1 + 8W_2 + kr)^{kr}$ such homomorphisms. The result follows. \square

We now have all results required to relate the number of k -cliques in a given graph G and the number of homomorphisms from $L(k, r, \ell_1, \ell_2)$ to $H(G, K, W_1, W_2)$, for appropriately chosen values of ℓ_1, ℓ_2, W_1, W_2 .

Lemma 9. Let $N \geq 0$ be the number of k -cliques in G , where $k = 4k'$ for some k' , $n = V(G)$, $m = E(G)$, and $2n + m > 6$. Let $M = M(\ell_1, \ell_2, W_1, W_2)$ be the number of homomorphisms from $L(k, r, \ell_1, \ell_2)$, $r = \binom{k}{2}$, to $H(G, k, W_1, W_2)$. If $W_2 = (2n + m)^2$, $W_1 = W_2^2$, $\ell_2 = 8kr$, and $\ell_1 = 17\ell_2$, then we have

$$N < \frac{M}{2W_1^{4\ell_1} W_2^{8\ell_2} k!} < N + \frac{1}{2}.$$

Proof. Let M_c be the total weight of identity and skew identity homomorphisms and let M_n be the total weight of the remaining homomorphisms. By Lemma 6, $M_c = 2W_1^{4\ell_1}W_2^{8\ell_2}k! \cdot N$. Therefore if $N \geq 1$ we only need to show that

$$M_n < \frac{M_c}{2N}. \quad (1)$$

Since $N \geq 1$,

$$\frac{M_c}{2N} = W_1^{4\ell_1}W_2^{8\ell_2}k! \geq W_1^{4\ell_1}W_2^{8\ell_2}, \quad (2)$$

and it suffices to show that $M_n < W_1^{4\ell_1}W_2^{8\ell_2}$. If $N = 0$ then it again suffices to show that

$$M_n < W_1^{4\ell_1}W_2^{8\ell_2}k! \geq W_1^{4\ell_1}W_2^{8\ell_2}. \quad !!!$$

On the other hand, $\ell_1 = 17\ell_2$ by the conditions of the lemma, that is, $\ell_1 > \frac{8T\ell_2}{T-1}$, where $T = \log_{W_2} W_1 = 2$. Therefore we are in the conditions of Lemma 8, and we have

$$M_n < W_1^{4\ell_1}W_2^{6\ell_2}(2n+m)^{2\ell_2} \cdot (4W_1 + 8W_2 + nmkr)^{kr}. \quad (3)$$

Note that for $n, m > 0, \varepsilon < 1$,

$$8W_2 = 8(2n+m)^2 < (2n+m)^4 = W_1. \quad (4)$$

Also, as $k \leq n, r \leq m$,

$$nmkr < (2n+m)^4 = W_1. \quad (5)$$

Using (4) and (5) in (3), we get

$$M_n < W_1^{4\ell_1}W_2^{6\ell_2}(2n+m)^{2\ell_2} \cdot (6W_1)^{kr}. \quad (6)$$

By (2) and (6), in order to establish (1) it suffices to prove

$$W_1^{4\ell_1}W_2^{6\ell_2}(2n+m)^{2\ell_2} \cdot (6W_1)^{kr} < W_1^{4\ell_1}W_2^{8\ell_2}, \quad (7)$$

or, equivalently, that

$$(2n+m)^{2\ell_2} \cdot (6W_1)^{kr} < W_2^{2\ell_2}. \quad (8)$$

Since $(2n+m)^{2\ell_2} = W_2^{\ell_2}$ and $W_1 = W_2^2$, inequality (8) is equivalent to

$$6^{kr} \cdot W_2^{2kr} < W_2^{\ell_2}. \quad (9)$$

Since $2n+m > 6$, we have

$$6^{kr} < (2n+m)^{kr}. \quad (10)$$

Finally, (10) together with $W_2 = (2n+m)^2$ and $\ell_2 = 8kr$ implies (9). \square

Finally, as Lemmas 8 and 9 are only proved for $k = 4k'$, we need to show that the problem for other values of the parameter can be reduced to k of such form. The following lemma takes care of that. Let $4p\text{-#CLIQUE}$ denote the following problem

Name: $4p\text{-#CLIQUE}$

Input: A graph G and $k \in \mathbb{N}$.

Parameter: k .

Output: The number of cliques of size $4k$ in G .

Lemma 10. *There is a parametrised AP-reduction from $p\text{-#CLIQUE}$ to $4p\text{-#CLIQUE}$.*

Proof. Let G, k be an instance of $p\text{-#CLIQUE}$ and let $\varepsilon \in (0, 1)$ be an error tolerance. If $k = 4k'$ for some k' then transform the instance to the instance G, k' of $4p\text{-#CLIQUE}$ with the same error tolerance ε . Otherwise repeat the following reduction as many times as required to obtain a parameter of the form $4k'$.

Suppose there is an FPRAS Alg that approximates the number of $(k + 1)$ -cliques in any graph. We construct graph $G^{+s} = (V', E')$ as follows. Let w_1, \dots, w_s be vertices not belonging to V . Then set $V' = V \cup \{w_1, \dots, w_s\}$ and $E' = E \cup \{vw_i \mid v \in V, i \in [s]\}$, that is, we connect all the new vertices with all vertices of G . The following claim is easy to verify.

Claim 1. Let N be the number of k -cliques in G and let N_1 be the number of $(k + 1)$ -cliques in G . Then the number of $(k + 1)$ -cliques in G^{+s} is $sN + N_1$.

Observe also that $N_1 < nN$, because every $(k + 1)$ -clique contains a k -clique, and for every k -clique C the number of $(k + 1)$ -cliques containing C is at most $n - k$.

The reduction works as follows: Apply Alg to the instance $G^{+s}, k + 1$, where $s = \frac{3n}{\varepsilon}$, with error tolerance $\varepsilon/3$. If it returns a number M output $Q = \lfloor \frac{M}{s} \rfloor$. We now show that $(1 - \varepsilon)N < Q < (1 + \varepsilon)N$. By Claim 1 we have

$$\left(1 - \frac{\varepsilon}{3}\right) (sN + N_1) < M < \left(1 + \frac{\varepsilon}{3}\right) (sN + N_1),$$

or equivalently,

$$\left(1 - \frac{\varepsilon}{3}\right) \left(N + \frac{N_1}{s}\right) < Q < \left(1 + \frac{\varepsilon}{3}\right) \left(N + \frac{N_1}{s}\right).$$

Since $(1 - \varepsilon)N < \left(1 - \frac{\varepsilon}{3}\right) \left(N + \frac{N_1}{s}\right)$, we obtain $(1 - \varepsilon)N < Q$. On the other hand, $N_1 < nN$, therefore

$$Q < \left(1 + \frac{\varepsilon}{3}\right) \left(N + \frac{N_1}{s}\right) < \left(1 + \frac{\varepsilon}{3}\right) \left(N + \frac{\varepsilon}{3}N\right) < (1 + \varepsilon)N.$$

The result follows. \square

In particular, Lemma 10 establishes #W[1]-hardness of the $4p\text{-#CLIQUE}$ problem.

4.3 Putting the Pieces Together

Proof of Theorem 3. As we mentioned earlier, conditions (1) and (4) are equivalent, the implications “(1) \Rightarrow (2) \Rightarrow (3)” are trivial and the implication “(4) \Rightarrow (1)” is known: if \mathcal{C} has bounded treewidth then, by [22, Proposition 7], $\#\text{CSP}(\mathcal{C}, -)$ belongs to FP.

The rest of the proof establishes “(3) \Rightarrow (4)”. Assume that $\#\text{CSP}(\mathcal{C}, -)$ admits an FPTRAS for a fan class \mathcal{C} . Our goal is to show that \mathcal{C} has bounded treewidth. For the sake of contradiction, assume that \mathcal{C} has unbounded treewidth. We will exhibit a parametrised reduction from $p\text{-}\#\text{CLIQUE}$ to $p\text{-}\#\text{CSP}(\mathcal{C}, -)$, which gives an FPTRAS for $p\text{-}\#\text{CLIQUE}$ assuming an FPTRAS for $p\text{-}\#\text{CSP}(\mathcal{C}, -)$. Under the assumption that $\text{FPT} \neq \text{W}[1]$ (under randomised parametrised reductions [14]), the $\text{W}[1]$ -hardness of $p\text{-}\#\text{CLIQUE}$ established in [13] implies, by [34, Corollary 3.17], the non-existence of an FPTRAS for the $p\text{-}\#\text{CLIQUE}$ problem, a contradiction.

Let $G = (V, E)$ and k be an instance of the $p\text{-}\#\text{CLIQUE}$ problem. By Lemma 10, we can assume that $k = 4k'$. First, we show that if G has any k -cliques at all, it can be assumed to have many k -cliques. Let $s \in \mathbb{N}$ and G_s is defined as follows. $V(G_s) = \{v_1, \dots, v_s \mid v \in V\}$ and $v_i w_j \in E(G_s)$, for $v, w \in V$ and $i, j \in [s]$, if and only if $vw \in E$. In other words, every vertex v of G is replaced with s distinct vertices v_1, \dots, v_s , and every edge vw is replaced with a complete bipartite graph $K_{s,s}$.

Claim 1. If N is the number of k -cliques in G , then G_s contains $s^k N$ k -cliques.

Proof of Claim 1. As is easily seen, for any indices $i_1, \dots, i_k \in [s]$ the vertices $v_{i_1}^1, \dots, v_{i_k}^k$ induce a clique in G_s if and only if v^1, \dots, v^k is a clique in G . Moreover, no clique in G_s contains vertices v_i, v_j for $v \in V$ and $i, j \in [s]$. The result follows. **(End of proof of Claim 1.)**

For a given instance $G = (V, E)$, k of $p\text{-}\#\text{CLIQUE}$ and error tolerance $\varepsilon \in (0, 1)$ using Claim 1, we first reduce it to the instance G_s, k of $p\text{-}\#\text{CLIQUE}$, where

$$s > \left(\frac{1 + \varepsilon/2}{\varepsilon} \right)^{\frac{1}{k}}.$$

Such a choice of s guarantees that if G_s contains any k -clique, the number of k -cliques it contains is at least $\frac{1 + \varepsilon/2}{\varepsilon}$. For simplicity we will have this assumption directly for G . We will also assume that if $n = |V|$ and $m = |E|$, then $2n + m > 6$.

Now we construct an instance \mathbf{A}, \mathbf{B} of $p\text{-}\#\text{CSP}(\mathcal{C}, -)$ such that an $\varepsilon/2$ -approximation of the number of homomorphisms from \mathbf{A} to \mathbf{B} yields an ε -approximation of the number of k -cliques in G . Structures \mathbf{A}, \mathbf{B} will be chosen to be (essentially) $\mathbf{A} = L(k, r, \ell_1, \ell_2)$ and $\mathbf{B} = H(G, k, W_1, W_2)$, where the parameters ℓ_1, ℓ_2, W_1, W_2 are set according to Lemma 9. Let $r = \binom{k}{2}$, $\ell_1 = 17\ell_2$, and $\ell_2 = 8kr$.

Since \mathcal{C} is a fan class and we assume that \mathcal{C} is not of bounded treewidth, there is a structure \mathbf{A} in \mathcal{C} such that $L(k, r, \ell_1, \ell_2)$ is the Gaifman graph $G(\mathbf{A})$ of \mathbf{A} .

We enumerate the class \mathcal{C} until we find such an \mathbf{A} . First we argue that \mathbf{A} can be assumed to be a τ -structure where τ consists of a single binary relation symbol; i.e., \mathbf{A} is a graph and hence $L(k, r, \ell_1, \ell_2)$ itself. Let \mathbf{A} be a τ -structure whose Gaifman graph $G(\mathbf{A})$ is $L(k, r, \ell_1, \ell_2)$. We show how to construct a τ -structure \mathbf{B} whose Gaifman graph $G(\mathbf{B})$ is $H(G, k, W_1, W_2)$ such that the set of homomorphisms from \mathbf{A} to \mathbf{B} is identical to the set of homomorphisms from $G(\mathbf{A})$ to $G(\mathbf{B}) = H(G, k, W_1, W_2)$, where $W_1 = (2n + m)^4$ and $W_2 = (2n + m)^2$. The universe of \mathbf{B} is the vertex set of $H(G, k, W_1, W_2)$. Let $R \in \tau$ and take any $\mathbf{x} \in R^{\mathbf{A}}$. Since $L(k, r, \ell_1, \ell_2)$ does not contain triangles, \mathbf{x} consists of at most two distinct elements, say $a, b \in A$. Let $I \subseteq [\text{ar}(R)]$ be the set of indices i with $\mathbf{x}[i] = a$. For every $u, v \in B$ with uv an edge in $H(G, k, W_1, W_2)$, we add (if it is not there already) to $R^{\mathbf{B}}$ the tuples \mathbf{y} and \mathbf{z} defined by $\mathbf{y}[i] = \mathbf{z}[j] = u$ and $\mathbf{y}[j] = \mathbf{z}[i] = v$ for every $i \in I$ and $j \notin I$. Now it is easy to see that a mapping $\varphi : A \rightarrow B$ is a homomorphism from \mathbf{A} to \mathbf{B} if and only if φ is a homomorphism from $G(\mathbf{A})$ to $G(\mathbf{B})$. !!!

Since the parameters $n, m, \ell_1, \ell_2, W_1, W_2$ satisfy the conditions of Lemma 9, by that lemma we have

$$N < \frac{M}{2W_1^{4\ell_1}W_2^{8\ell_2}k!} < N + \frac{1}{2}, \quad (11)$$

where N is the number of k -cliques in G , which we want to approximate within ε , and M is the number of homomorphisms from \mathbf{A} to \mathbf{B} , for which we have an FPTRAS by assumption. Let $Q = M/(2W_1^{4\ell_1}W_2^{8\ell_2}k!)$. The FPTRAS for p - $\#\text{CSP}(\mathcal{C}, -)$ applied with error tolerance $\varepsilon/2$ produces a number M' such that

$$(1 - \varepsilon/2)M < M' < (1 + \varepsilon/2)M. \quad (12)$$

We then return $\lfloor Q' \rfloor$, where

$$Q' = \frac{M'}{2W_1^{4\ell_1}W_2^{8\ell_2}k!}.$$

It remains to show that $(1 - \varepsilon)N < Q' < (1 + \varepsilon)N$. On one hand, we have

$$Q' > (1 - \varepsilon/2)Q \geq (1 - \varepsilon/2)N \geq (1 - \varepsilon)N,$$

where the first inequality follows from (12) and the definitions of Q and Q' , the second inequality follows from (11) and the definitions of Q and N , and the third inequality is trivial.

On the other hand, we have

$$Q' < (1 + \varepsilon/2)Q < (1 + \varepsilon/2) \left(N + \frac{1}{2} \right),$$

where again the first inequality follows from (12) and the second inequality follows from (11).

Assume first that $N = 0$. Then $Q' < \frac{1+\varepsilon/2}{2}$, and by the assumption $\varepsilon < 1$ we have $\lfloor Q' \rfloor = 0$ as required. Otherwise by the assumption on the number of k -cliques in G , $N > \frac{1+\varepsilon/2}{\varepsilon}$; therefore

$$Q' < (1+\varepsilon/2) \left(N + \frac{1}{2} \right) = (1+\varepsilon/2)N + \frac{1+\varepsilon/2}{2} < (1+\varepsilon/2)N + (\varepsilon/2)N = (1+\varepsilon)N.$$

Observe that the reduction runs in time $f(k) \cdot \text{poly}(n + m, \varepsilon^{-1})$ and is a parametrised AP-reduction. Thus, the reduction gives an FPTRAS for N . Theorem 3 is proved. \square

5 Conclusions

We do not know whether Theorem 3 holds for all classes of (bounded-arity) relational structures.

With more technicalities (but the same ideas as presented here), one can weaken the assumption on a fan class to obtain the same result (Theorem 3). In particular, it suffices to require that there are polynomials f_1, f_2, f_3, f_4 such that for any parameters $k, r, \ell_1, \ell_2 \in \mathbb{N}$, $G(\mathcal{C})$ contains the fan-grid $L(k', r', \ell'_1, \ell'_2)$, where $k' = f_1(k, r, \ell_1, \ell_2) \geq k$, $r' = f_2(k, r, \ell_1, \ell_2) \geq r$, $\ell'_1 = f_3(k, r, \ell_1, \ell_2) \geq \ell_1$, $\ell'_2 = f_4(k, r, \ell_1, \ell_2) \geq \ell_2$. This can be achieved by making use of Lemma 10 (as it would not be possible to test directly for cliques of all sizes) and by a modification of the construction from Section 4.1 (to accommodate for the fact that some fan-grids may not correspond to cliques due to incompatible numbers).

References

- [1] Vikraman Arvind and Venkatesh Raman. Approximation algorithms for some parameterized counting problems. In *Proceedings of the 13th International Symposium on Algorithms and Computation (ISAAC'02)*, volume 2518 of *Lecture Notes in Computer Science*, pages 453–464. Springer, 2002. doi:10.1007/3-540-36136-7_40.
- [2] Johann Brault-Baron, Florent Capelli, and Stefan Mengel. Understanding model counting for beta-acyclic CNF-formulas. In *Proceedings of the 32nd International Symposium on Theoretical Aspects of Computer Science (STACS'15)*, pages 143–156, 2015. doi:10.4230/LIPIcs.STACS.2015.143.
- [3] Andrei Bulatov. A dichotomy theorem for nonuniform CSP. In *Proceedings of the 58th Annual IEEE Symposium on Foundations of Computer Science (FOCS'17)*, pages 319–330. IEEE, 2017. doi:10.1109/FOCS.2017.37.

- [4] Andrei Bulatov, Peter Jeavons, and Andrei Krokhin. Classifying the Complexity of Constraints using Finite Algebras. *SIAM Journal on Computing*, 34(3):720–742, 2005. doi:[10.1137/S0097539700376676](https://doi.org/10.1137/S0097539700376676).
- [5] Andrei A. Bulatov. The complexity of the counting constraint satisfaction problem. *Journal of the ACM*, 60(5):34, 2013. doi:[10.1145/2528400](https://doi.org/10.1145/2528400).
- [6] Jin-Yi Cai and Xi Chen. Complexity of counting CSP with complex weights. *Journal of the ACM*, 64(3):19:1–19:39, 2017. doi:[10.1145/2822891](https://doi.org/10.1145/2822891).
- [7] Xi Chen, Martin E. Dyer, Leslie Ann Goldberg, Mark Jerrum, Pinyan Lu, Colin McQuillan, and David Richerby. The complexity of approximating conservative counting CSPs. *Journal of Computer and System Sciences*, 81(1):311–329, 2015. doi:[10.1109/FOCS.2017.37](https://doi.org/10.1109/FOCS.2017.37).
- [8] Nadia Creignou and Miki Hermann. Complexity of generalized satisfiability counting problems. *Information and Computation*, 125(1):1–12, 1996. doi:[10.1006/inco.1996.0016](https://doi.org/10.1006/inco.1996.0016).
- [9] Víctor Dalmau and Peter Jonsson. The complexity of counting homomorphisms seen from the other side. *Theoretical Computer Science*, 329(1-3):315–323, 2004. doi:[10.1016/j.tcs.2004.08.008](https://doi.org/10.1016/j.tcs.2004.08.008).
- [10] Víctor Dalmau, Phokion G. Kolaitis, and Moshe Y. Vardi. Constraint Satisfaction, Bounded Treewidth, and Finite-Variable Logics. In *Proceedings of the 8th International Conference on Principles and Practice of Constraint Programming (CP'02)*, volume 2470 of *Lecture Notes in Computer Science*, pages 310–326. Springer, 2002. doi:[10.1007/3-540-46135-3_21](https://doi.org/10.1007/3-540-46135-3_21).
- [11] Reinhard Diestel. *Graph Theory*. Springer, fourth edition, 2010.
- [12] Rodney G. Downey and Michael R. Fellows. Fixed-Parameter Tractability and Completeness I: Basic Results. *SIAM Journal on Computing*, 24(4):873–921, 1995. doi:[10.1137/S0097539792228228](https://doi.org/10.1137/S0097539792228228).
- [13] Rodney G. Downey and Michael R. Fellows. Fixed-parameter tractability and completeness II: On completeness for W[1]. *Theoretical Computer Science*, 141(1–2):109–131, 1995. doi:[10.1016/0304-3975\(94\)00097-3](https://doi.org/10.1016/0304-3975(94)00097-3).
- [14] Rodney G. Downey, Michael R. Fellows, and Kenneth W. Regan. Parameterized circuit complexity and the W hierarchy. *Theoretical Computer Science*, 191(1-2):97–115, 1998. doi:[10.1016/S0304-3975\(96\)00317-9](https://doi.org/10.1016/S0304-3975(96)00317-9).
- [15] Martin E. Dyer, Leslie Ann Goldberg, Catherine S. Greenhill, and Mark Jerrum. The relative complexity of approximate counting problems. *Algorithmica*, 38(3):471–500, 2004. doi:[10.1007/s00453-003-1073-y](https://doi.org/10.1007/s00453-003-1073-y).

- [16] Martin E. Dyer, Leslie Ann Goldberg, and Mark Jerrum. An approximation trichotomy for Boolean #CSP. *Journal of Computer and System Sciences*, 76(3-4):267–277, 2010. doi:10.1016/j.jcss.2009.08.003.
- [17] Martin E. Dyer and Catherine S. Greenhill. The complexity of counting graph homomorphisms. *Random Struct. Algorithms*, 17(3-4):260–289, 2000. doi:10.1002/1098-2418(200010/12)17:3/4<260::AID-RSA5>3.0.CO;2-W.
- [18] Martin E. Dyer and David Richerby. An effective dichotomy for the counting constraint satisfaction problem. *SIAM Journal on Computing*, 42(3):1245–1274, 2013. doi:10.1137/100811258.
- [19] Tomás Feder, Pavol Hell, Daniel Král’, and Jiří Sgall. Two algorithms for general list matrix partitions. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA’05)*, pages 870–876, 2005. URL: <http://dl.acm.org/citation.cfm?id=1070432.1070554>.
- [20] Tomás Feder, Pavol Hell, and Wing Xie. Matrix partitions with finitely many obstructions. *Electr. J. Comb.*, 14(1), 2007. URL: http://www.combinatorics.org/Volume_14/Abstracts/v14i1r58.html.
- [21] Tomás Feder and Moshe Y. Vardi. The Computational Structure of Monotone Monadic SNP and Constraint Satisfaction: A Study through Datalog and Group Theory. *SIAM Journal on Computing*, 28(1):57–104, 1998. doi:10.1137/S0097539794266766.
- [22] Jörg Flum and Martin Grohe. The Parameterized Complexity of Counting Problems. In *Proceedings of the 43rd Symposium on Foundations of Computer Science (FOCS’02)*, page 538. IEEE Computer Society, 2002. doi:10.1109/SFCS.2002.1181978.
- [23] Jörg Flum and Martin Grohe. The parameterized complexity of counting problems. *SIAM Journal on Computing*, 33(4):892–922, 2004. doi:10.1137/S0097539703427203.
- [24] Jörg Flum and Martin Grohe. *Parametrized Complexity Theory*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2006.
- [25] Andreas Galanis, Leslie Ann Goldberg, and Mark Jerrum. Approximately Counting H -Colorings is #BIS-Hard. *SIAM Journal on Computing*, 45(3):680–711, 2016. doi:10.1137/15M1020551.
- [26] Andreas Galanis, Leslie Ann Goldberg, and Mark Jerrum. A Complexity Trichotomy for Approximately Counting List H -Colorings. *ACM Transactions on Computation Theory*, 9(2):9:1–9:22, 2017. doi:10.1145/3037381.

- [27] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W.H. Freeman, 1979.
- [28] Georg Gottlob, Nicola Leone, and Francesco Scarcello. Hypertree decomposition and tractable queries. *Journal of Computer and System Sciences*, 64(3):579–627, 2002. doi:10.1006/jcss.2001.1809.
- [29] Martin Grohe. The complexity of homomorphism and constraint satisfaction problems seen from the other side. *Journal of the ACM*, 54(1):1–24, 2007. doi:10.1145/1206035.1206036.
- [30] Martin Grohe and Dániel Marx. Constraint solving via fractional edge covers. *ACM Transactions on Algorithms*, 11(1):4:1–4:20, 2014. doi:10.1145/2636918.
- [31] Martin Grohe, Thomas Schwentick, and Luc Segoufin. When is the evaluation of conjunctive queries tractable? In *Proceedings 33rd ACM Symposium on Theory of Computing (STOC’01)*, pages 657–666. ACM, 2001. doi:10.1145/380752.380867.
- [32] Pavol Hell and Jaroslav Nešetřil. On the Complexity of H -coloring. *Journal of Combinatorial Theory, Series B*, 48(1):92–110, 1990. doi:10.1016/0095-8956(90)90132-J.
- [33] Phokion G. Kolaitis and Moshe Y. Vardi. Conjunctive-Query Containment and Constraint Satisfaction. *Journal of Computer and System Sciences*, 61(2):302–332, 2000. doi:10.1006/jcss.2000.1713.
- [34] Kitty Meeks. The challenges of unbounded treewidth in parameterised subgraph counting problems. *Discrete Applied Mathematics*, 198:170–194, 2016. doi:10.1016/j.dam.2015.06.019.
- [35] Michael Mitzenmacher and Eli Upfal. *Probability and Computing*. Cambridge University Press, Cambridge, second edition, 2017. Randomization and Probabilistic techniques in Algorithms and Data Analysis.
- [36] Neil Robertson and Paul D. Seymour. Graph minors. III. Planar tree-width. *Journal of Combinatorial Theory, Series B*, 36(1):49–64, 1984. doi:10.1016/0095-8956(84)90013-3.
- [37] Neil Robertson and Paul D. Seymour. Graph minors. V. Excluding a planar graph. *Journal of Combinatorial Theory, Series B*, 41(1):92–114, 1986. doi:10.1016/0095-8956(86)90030-4.
- [38] Thomas J. Schaefer. The complexity of satisfiability problems. In *Proceedings of the 10th Annual ACM Symposium on Theory of Computing (STOC’78)*, pages 216–226, 1978. doi:10.1145/800133.804350.

- [39] Dmitriy Zhuk. The Proof of CSP Dichotomy Conjecture. In *Proceedings of the 58th Annual IEEE Symposium on Foundations of Computer Science (FOCS'17)*, pages 331–342. IEEE, 2017. [doi:10.1109/FOCS.2017.38](https://doi.org/10.1109/FOCS.2017.38).