



Fast Algorithms for Join Operations on Tree Decompositions

Johan M. M. van Rooij^(✉)

Department of Information and Computing Sciences, Utrecht University,
PO Box 80.089, 3508 TB Utrecht, The Netherlands
J.M.M.vanRooij@uu.nl

Abstract. Treewidth is a measure of how tree-like a graph is. It has many important algorithmic applications because many NP-hard problems on general graphs become tractable when restricted to graphs of bounded treewidth. Algorithms for problems on graphs of bounded treewidth mostly are dynamic programming algorithms using the structure of a tree decomposition of the graph. The bottleneck in the worst-case run time of these algorithms often is the computations for the so called join nodes in the associated nice tree decomposition.

In this paper, we review two different approaches that have appeared in the literature about computations for the join nodes: one using fast zeta and Möbius transforms and one using fast Fourier transforms. We combine these approaches to obtain new, faster algorithms for a broad class of vertex subset problems known as the $[\sigma, \rho]$ -domination problems. Our main result is that we show how to solve $[\sigma, \rho]$ -domination problems in $\mathcal{O}(s^{t+2}tn^2(t \log(s) + \log(n)))$ arithmetic operations. Here, t is the treewidth, s is the (fixed) number of states required to represent partial solutions of the specific $[\sigma, \rho]$ -domination problem, and n is the number of vertices in the graph. This reduces the polynomial factors involved compared to the previously best time bound (van Rooij, Bodlaender, Rossmanith, ESA 2009) of $\mathcal{O}(s^{t+2}(st)^{2(s-2)}n^3)$ arithmetic operations. In particular, this removes the dependence of the degree of the polynomial on the fixed number of states s .

Keywords: Tree decompositions · Dynamic programming · Fast Fourier transform · Möbius transform · Fast subset convolution · Sigma-rho domination

1 Introduction

Treewidth is an important concept in the theory of graph algorithms that measures how tree-like a graph is. While many problems that are \mathcal{NP} -hard on general graphs become efficiently solvable when restricted to trees, this often extends to these problems being polynomial or even linear-time solvable when restricted to graphs that have bounded treewidth. In general this is done in two steps:

1. Find a tree decomposition of the input graph of small treewidth.

2. Solve the problem by dynamic programming on this tree decomposition.

In this paper, we focus on the second of these two steps and show how to improve the running times of algorithms on tree decompositions using algebraic transforms. We apply these to the general case of the so called $[\sigma, \rho]$ -domination problems. This includes many well-known vertex subset problems such as INDEPENDENT SET, DOMINATING SET and TOTAL DOMINATING SET, but also problems such as INDUCED BOUNDED DEGREE SUBGRAPH and INDUCED p -REGULAR SUBGRAPH.

If we assume that a graph G is given with a tree decomposition T of G of width t , then the running time of an algorithm on tree decompositions is typically polynomial in the size of graph G , but exponential in the treewidth t . Early examples of such algorithms include algorithms on vertex partitioning problems (including the $[\rho, \sigma]$ -domination problems) [30], edge colouring problems such as CHROMATIC INDEX [5], or other problems such as STEINER TREE [21]. Often the worst-case running time of these algorithms involve large factors that depend on the treewidth t . This lead researchers to look for algorithms where these factors grow as slow as possible as a function of t . For several DOMINATING SET-like problems such as INDEPENDENT DOMINATING SET, TOTAL DOMINATING SET, PERFECT DOMINATING SET and PERFECT CODE, Alber et al. [1] give improved algorithms with special attention to the exponential dependence on the treewidth t : for example, they showed how to solve DOMINATING SET in $\mathcal{O}^*(4^t)$ time. This was improved by Van Rooij et al. [26] who first showed how to solve DOMINATING SET in $\mathcal{O}^*(3^t)$ time by giving an $\mathcal{O}(3^{t^2}n)$ -time algorithm. Van Rooij et al. also generalised this result solving the $[\rho, \sigma]$ -domination problems in $\mathcal{O}(s^{t+2}(st)^{2(s-2)}n^3)$ time. The result for DOMINATING SET seems to be optimal in some sense, as Lokshtanov et al. [22] showed that any $\mathcal{O}^*((3 - \epsilon)^t)$ -time algorithm would violate the *Strong Exponential-Time Hypothesis*; we expect the same for the other $[\rho, \sigma]$ -domination problems.

Since then, several results have appeared improving running times of dynamic programming algorithms on tree decompositions. For example, the algorithm by Van Rooij et al. [26] has been generalised to DISTANCE- r DOMINATING SET [12] and DISTANCE- r INDEPENDENT SET [18]. The most notable new results are the *Cut and Count* technique [16] giving randomised $\mathcal{O}^*(c^t)$ -time algorithms for many graph connectivity problems, mostly supported by matching lower bounds based on the Strong Exponential-Time Hypotheses, the *rank-based approach* [9, 14] and the *determinant-based approach* [9, 31] that derandomise these results at the cost of a greater base of the exponent c .

For many of these algorithms, the computations in the so called join nodes of a nice tree decomposition are the bottleneck of the worst-case run time. To speed up these computations, several approaches have been used, often based on algebraic transforms. One such method is using fast zeta and Möbius transforms in a way that is similar to the well-known fast subset convolution algorithm by Björklund et al. [2]. This method was first used in the context of tree decompositions by Van Rooij et al. [26] who also generalised the approach to work for the $[\sigma, \rho]$ -domination problems. At the same time, Cygan and Pilipczuk, showed that

the fast subset convolution result could also be based on Fourier transforms [17]; they also generalised it in a different way. A variant to this approach that we follow in this paper, directly applied to tree decompositions, can be found in the appendix of [15]. We will discuss both these approaches in more detail in this paper. Finally, faster joins are also obtained based on Clifford algebras [31], but these are beyond the scope of this paper.

1.1 Goal of This Paper

The goal of this paper is twofold. Firstly, we want to present a faster algorithm for the $[\sigma, \rho]$ -domination problems. This algorithm uses $\mathcal{O}(s^{t+2}tn^2(t \log(s) + \log(n)))$ arithmetic operations: this improves the polynomial factors compared to our earlier result [26] and removes the dependency of the degree of the polynomial on s , where s is the (fixed) number of states used. Secondly, we want to give a comprehensible overview of how Fourier and Möbius transforms can be used to obtain faster algorithms on tree decompositions.

We choose to take an algebraic perspective that allows for easier generalisation and easier combination of Fourier and Möbius transform-based approaches than that in [25, 26]. However, we consider the approach in [25, 26] to be more intuitive: it relies only on counting arguments (this is especially true for the first algorithm for DOMINATING SET in [25] that does not explicitly use any algebraic transform). In our overview, we will not give details on our earlier generalised convolution approach from [26]: after the initial examples, we directly go to the new and improved algorithm.

2 Preliminaries

2.1 Graphs and Tree Decompositions

Let $G = (V, E)$ be an n -vertex graph with m edges. A *terminal graph*¹ $G_X = (V, E, X)$ is a graph $G = (V, E)$ with an ordered sequence of distinct vertices that we call its terminals: $X = \{x_1, x_2, \dots, x_k\}$ with each $x_j \in V$. Two terminal graphs $G_X = (V_1, E_1, X_1)$ and $H_X = (V_2, E_2, X_2)$ with the same number of terminals k , but disjoint vertex and edge sets, can be *glued* together to form the terminal graph $G_X \oplus H_X$ by identifying each terminal x_i from X_1 with x_i from X_2 , for all $1 \leq i \leq k$. That is, if $X = X_1 = X_2$ through identification, then $G_X \oplus H_X = (V_1 \cup V_2, E_1 \cup E_2, X)$. A *completion* of a terminal graph G_X is a non-terminal graph G that can be obtained from G_X by gluing a terminal graph H_X on G_X and then ignoring which vertices are terminals in the result.

The treewidth of a (non-terminal) graph is a measure of how-tree like the graph is. From an algorithmic viewpoint this is a very useful concept because, where many \mathcal{NP} -hard problems on general graphs are linear time solvable on trees by dynamic programming, often similar style dynamic programming algorithms exist for graphs whose treewidth is bounded by a constant. We outline

¹ This is also known as a k -boundary graph.

the basics on treewidth and specifically on dynamic programming on tree decompositions below. More information can, amongst other places, be found in work by Bodlaender [4,6–8,10].

Definition 1 (tree decomposition and treewidth). A tree decomposition of an undirected graph $G = (V, E)$ is a tree T in which each node $i \in T$ has an associated set of vertices $X_i \subseteq V$ (called a bag), with $\bigcup_{i \in T} X_i = V$, such that the following properties hold:

- for every edge $\{u, v\} \in E$, there exist a bag X_i such that $\{u, v\} \subseteq X_i$;
- for every vertex v in G , the bags containing x form a connected subtree: i.e., if $v \in X_i$ and $v \in X_j$, then $v \in X_k$ for all nodes k on the path from i to j in T .

The width of a tree decomposition T is defined as $\max_{i \in T} \{|X_i|\} - 1$: the size of the largest bag minus one. The treewidth of a graph G is the minimum width over all tree decomposition of G .

For a tree decomposition T with assigned root node $r \in T$, we define the terminal graph $G_i = (V_i, E_i, X_i)$ for each node $i \in T$: let V_i be the union of X_i with all bags X_j where j is a descendant of i in T , and let $E_i \subseteq E$ be the set of edges with at least one endpoint in $V_i \setminus X_i$ (and as a result of Definition 1 with both endpoints in V_i). Now, G_i contains all edges between vertices in $V_i \setminus X_i$, and all edges between $V_i \setminus X_i$ and X_i , but no edges between two vertices in X_i .² Observe that, G is the completion of G_i formed through $G_i \oplus ((V \setminus V_i) \cup X_i, E \setminus E_i, X_i)$, and X_i can be seen as the *separator* separating $V_i \setminus X_i$ from $V \setminus V_i$ in G (where either side of the separator can be empty).

We now describe dynamic programming on a tree decomposition T . Given a graph problem that we are trying to solve \mathcal{P} , define a *partial solution* of \mathcal{P} on G_i to be the *restriction* to the subgraph G_i of a solution of \mathcal{P} on a completion of G_i (any completion of G_i , not only G itself). We say that the partial solution S' on G_i can be *extended* to a full solution S on a completion of G_i , where $S \setminus S'$ is the *extension* of S' . As an example, consider the MINIMUM DOMINATING SET problem: a solution for this problem is a vertex subset D in G such that for all $v \in V$ there is a $d \in D$ with $v \in N[d]$. A partial solution is a subset $D \subseteq V_i$ such that for all vertices in $v \in V_i \setminus X_i$ there is a $d \in D$ with $v \in N[d]$: for vertices in X_i there does not need to be a dominating neighbour in $d \in D$ as d can also be in an extension of D . A dynamic programming algorithm on a tree decomposition computes, for each node $i \in T$ in a bottom-up fashion, a *memoisation table* A_i containing all *relevant* (described in the next paragraph) partial solutions on G_i obtaining a solution to \mathcal{P} in the root of T .

To restrict the number of partial (relevant) solutions stored, an equivalence relation is defined on them: two partial solutions S'_1 and S'_2 on G_i are *equivalent*

² Often G_i is defined *including* all edges between vertices in X_i . We choose the alternative definition as it makes formulating the join algorithms in Sect. 4 easier: no bookkeeping of number of neighbours between vertices in X_i needs to be done, as they only become neighbours higher up in the tree.

with respect to \mathcal{P} if any extension of S_1 also is an extension of S'_2 and vice versa. When given two equivalent partial solutions S'_1 and S'_2 for an optimisation problem (minimisation or maximisation), we say that S'_1 *dominates* S'_2 if for any extension S_E of S'_1 and S'_2 , the solution value of $S'_1 \cup S_E$ is equal or better than the solution value of $S'_2 \cup S_E$. Clearly, a dynamic programming algorithm on a tree decomposition needs to store only one partial solution per equivalence class, and if we consider an optimisation problem it can store a partial solution that dominates all other partial solutions within its equivalence class.

Mostly, it is convenient to formulate a dynamic programming algorithm on a special kind of tree decomposition called a *nice tree decomposition* [20].³

Definition 2 (nice tree decomposition). A nice tree decomposition is a tree decomposition T with assigned root node $r \in T$ with $X_r = \emptyset$, in which each node is of one of the following types:

- Leaf node: a leaf i of T with $X_i = \emptyset$.
- Introduce node: an internal node i of T with one child node j and $X_i = X_j \cup \{v\}$ for some $v \in V \setminus V_j$.
- Forget node: an internal node i of T with one child node j and $X_i = X_j \setminus \{v\}$ for some $v \in X_j$.
- Join node: an internal node i of T with two child nodes l and r with $X_i = X_l = X_r$.

Given a tree decomposition consisting of $\mathcal{O}(n)$ nodes, a nice tree decomposition of $\mathcal{O}(n)$ nodes of the same width can be found in $\mathcal{O}(n)$ time [20]. Consequently, a dynamic programming algorithm on a nice tree decomposition can be used on general tree decompositions by applying this transformation. After computing A_i for all nodes $i \in T$, the solution to \mathcal{P} can be found as the unique value in A_r , where r is the root of T : here $G_i = G$ and there is only a single equivalence class as $X_i = \emptyset$.

This paper focuses on computing A_i for a *join node* i of a nice tree decomposition. This node is the most interesting as often it dominates the running time of the entire dynamic programming algorithm. For an example, consider [1] where an $\mathcal{O}^*(4^t)$ algorithm for MINIMUM DOMINATING SET for graphs with a tree decomposition of width t is given, while all computations except the computation for the join nodes can be performed in $\mathcal{O}^*(3^t)$ time.

2.2 Dynamic Programming for $[\sigma, \rho]$ -Domination Problems

The $[\sigma, \rho]$ -domination problems are a class of vertex-subset problems introduced by Telle [28–30] that generalise many well-known graph problems such as MAXIMUM INDEPENDENT SET, MINIMUM DOMINATING SET, and INDUCED BOUNDED DEGREE SUBGRAPH. See Table 1 for an overview.

³ Different version of the original definition [20] exists in literature (e.g. [16, 25]): the restrictions on the vertices in a bag of a leaf node and the root node often vary, and sometimes an additional type of node called an *edge introduce node* is used.

Definition 3 ($[\sigma, \rho]$ -dominating set). *Let $\sigma, \rho \subseteq \mathbb{N}$, a $[\sigma, \rho]$ -dominating set in a graph $G = (V, E)$ is a subset $D \subseteq V$ such that:*

- for every $v \in D: |N(v) \cap D| \in \sigma;$
- for every $v \in V \setminus D: |N(v) \cap D| \in \rho.$

We consider only $\sigma, \rho \subseteq \mathbb{N}$ that both are either finite or cofinite.

For given $\sigma, \rho \subseteq \mathbb{N}$ and the corresponding definition of a $[\sigma, \rho]$ -dominating set, one can define several different problem variants.

- *Existence* problem: given a graph G , does G have a $[\sigma, \rho]$ -dominating set?
- *Optimisation* problem (minimisation or maximisation): given a graph G , what is the smallest $[\sigma, \rho]$ -dominating set in G , or what is the largest $[\sigma, \rho]$ -dominating set in G ?
- *Counting* problem: given a graph G , how many $[\sigma, \rho]$ -dominating sets exist in G ?
- *Counting optimisation* problem (minimisation or maximisation): given a graph G , how many $[\sigma, \rho]$ -dominating sets in G exist of minimum/maximum size?

Many well-known NP-hard vertex subset problems in graphs correspond to the existence or optimisation variant of a $[\sigma, \rho]$ -domination problem, as can be seen from Table 1.

When solving a $[\sigma, \rho]$ -domination problem by dynamic programming on a tree decomposition, the equivalence classes for partial solutions stored in the memoisation table A_i (as defined in Sect. 2.1) can be uniquely identified by the following:

Table 1. Examples of $[\sigma, \rho]$ -domination problems (taken from [28–30]).

σ	ρ	Standard description
{0}	{0, 1, ...}	Independent Set/Stable Set
{0, 1, ...}	{1, 2, ...}	Dominating Set
{0}	{0, 1}	Strong Stable Set/2-Packing/Distance-2 Independent Set
{0}	{1}	Perfect Code/Efficient Dominating Set
{0}	{1, 2, ...}	Independent Dominating Set
{0, 1, ...}	{1}	Perfect Dominating Set
{1, 2, ...}	{1, 2, ...}	Total Dominating Set
{1}	{1}	Total Perfect Dominating Set
{0, 1, ...}	{0, 1}	Nearly Perfect Set
{0, 1}	{0, 1}	Total Nearly Perfect Set
{0, 1}	{1}	Weakly Perfect Dominating Set
{0, 1, ..., p}	{0, 1, ...}	Induced Bounded Degree Subgraph
{0, 1, ...}	{p, p + 1, ...}	p-Dominating Set
{p}	{0, 1, ...}	Induced p-Regular Subgraph

- the vertices in X_i that are in the partial solution D ;
- for every vertex in X_i (both in D and not in D), the number of neighbours in D .

This corresponds exactly to the bookkeeping required to verify whether a partial solution locally satisfies the requirements imposed by the specific $[\sigma, \rho]$ -domination problem. As such, we can identify every equivalence class using an assignment of *labels* (sometimes also called *states*) that capture the above properties to the vertices in X_i : such an assignment is called a *state colouring*. Given $\sigma, \rho \subseteq \mathbb{N}$, define the set of labels $C = C_\sigma \cup C_\rho$ as follows (the meaning of a label is explained below):

$$\begin{aligned}
 C_\sigma &= \begin{cases} \{ |0|_\sigma, |1|_\sigma, |2|_\sigma, \dots, |\ell - 1|_\sigma, |\ell|_\sigma \} & \text{if } \sigma \text{ finite} & \text{where } \ell = \max\{\sigma\} \\ \{ |\geq 0|_\sigma \} & \text{if } \sigma = \mathbb{N} \\ \{ |0|_\sigma, |1|_\sigma, |2|_\sigma, \dots, |\ell - 1|_\sigma, |\geq \ell|_\sigma \} & \text{if } \sigma \neq \mathbb{N} \text{ cofinite} & \text{where } \ell = \max\{\mathbb{N} \setminus \sigma\} + 1 \end{cases} \\
 C_\rho &= \begin{cases} \{ |0|_\rho, |1|_\rho, |2|_\rho, \dots, |\ell - 1|_\rho, |\ell|_\rho \} & \text{if } \rho \text{ finite} & \text{where } \ell = \max\{\rho\} \\ \{ |\geq 0|_\rho \} & \text{if } \rho = \mathbb{N} \\ \{ |0|_\rho, |1|_\rho, |2|_\rho, \dots, |\ell - 1|_\rho, |\geq \ell|_\rho \} & \text{if } \rho \neq \mathbb{N} \text{ cofinite} & \text{where } \ell = \max\{\mathbb{N} \setminus \rho\} + 1 \end{cases}
 \end{aligned}$$

We will use the $|\cdot|_\rho$ and $|\cdot|_\sigma$ notation to denote labels from C_ρ , respectively C_σ . In general, when we write $|l|_\rho$ or $|l|_\sigma$, with a variable l , we mean the labels that are not equal to $|\geq \ell|_\rho$ or $|\geq \ell|_\sigma$. This allows us to refer to other labels by expressions such as $|l - 1|_\rho$. The symbol ℓ is reserved to indicate the last labels $|\ell|_\sigma, |\geq \ell|_\sigma, |\ell|_\rho, |\geq \ell|_\rho$ and is used similarly to form labels such as $|\ell - 1|_\rho$.

Let C^{X_i} be the set of assignments of labels from C to the vertices in X_i . A label from C_σ for a vertex $v \in X_i$ indicates that v is in the solution set D in the partial solution, a label from C_ρ indicates that v is not. Furthermore, the numbers in the labels indicate the number of neighbours that v has in D ; the \geq symbol in the label $|\geq 1|_\rho$ indicates that v has this number of neighbours (one in this case) in D or more. For an example, consider MINIMUM DOMINATING SET for which $\sigma = \mathbb{N}$ and $\rho = \mathbb{N} \setminus \{0\}$; for this problem $C = \{ |\geq 0|_\sigma, |0|_\rho, |\geq 1|_\rho \}$.

Now, the elements from C^{X_i} bijectively correspond to the above defined equivalence classes of partial solutions on G_i . Consequently, we can index the memoisation table A_i by C^{X_i} . To keep the dynamic programming recurrences in this paper simple, we will not store partial solutions in A_i , only the required partial solution values or counts. That is, from here on, let the table A_i be a function $A_i : C^{X_i} \rightarrow \{0, 1, \dots, M\} \cup \{\infty\}$ that assigns a number to each equivalence class of partial solutions. In an existence variant of a problem, we let $A_i(\mathbf{c})$, for $\mathbf{c} \in C^{X_i}$, be 0 or 1 indicating whether a partial solution of this equivalence class exists. In an optimisation variant, $A_i(\mathbf{c})$ indicates the size of a dominating partial solution in this equivalence class, or ∞ if no such partial solution exists. For convenience reasons⁴, we let $A_i(\mathbf{c})$, for $\mathbf{c} \in C^{X_i}$, contain the size of the partial solution D' restricted to $V' \setminus X'$, i.e., the size of a corresponding partial solution equals $A_i(\mathbf{c})$ plus the number of σ labels in \mathbf{c} . In a counting variant, $A_i(\mathbf{c})$ indicates the number of partial solutions in the equivalence class of \mathbf{c} .

⁴ In this way, we do not have to correct for double counting in join nodes in the rest of this paper.

Notice that for an existence variant, we can bound M by 1; for an optimisation variant, we can bound M by n ; and for a counting variant, we can bound M by 2^n .

Below, we give explicit recurrences for A_i for solving a minimisation variant of a $[\sigma, \rho]$ -dominating problem by dynamic programming on a nice tree decomposition T . Modifying the recurrences to the existence or counting variant of the problem is an easy exercise. Extensions to the recurrences in which partial solutions are stored (for existence and optimisation variants) are easy to make, but tedious to write down formally. This is also true for the extension to the optimisation counting variant where one needs to keep track of both the size and the number of such partial solutions.

Leaf Node. Let i be a leaf node of T . Since $X_i = \emptyset$, the only partial solution is \emptyset with size zero: this size is stored for the empty vector \square .

$$A_i(\square) = 0$$

Introduce Node. Let i be an introduce node of T with child node j . Let $X_i = X_j \cup \{v\}$ for some $v \in V \setminus V_j$. For $\mathbf{c} \in C^{X_j}$ and $c_v \in C$ the label for vertex v denote by $[\mathbf{c}, c_v]$ the vector \mathbf{c} with the element c_v appended to it such that $[\mathbf{c}, c_v] \in C^{X_i}$. Now:

$$A_i([\mathbf{c}, c_v]) = \begin{cases} A_j([\mathbf{c}]) & \text{if } c_v \in \{0|_\sigma, | \geq 0|_\sigma\} \text{ or } c_v \in \{0|_\rho, | \geq 0|_\rho\} \\ \infty & \text{otherwise} \end{cases}$$

Here, G_i equals G_j with one added isolated vertex v . Hence, v can be in the partial solution or not, and both choices do not influence the partial solution size on $V_i \setminus X_i$ (which equals $V_j \setminus X_j$). Note that only one of the labels from $\{0|_\sigma, | \geq 0|_\sigma\}$ and one from $\{0|_\rho, | \geq 0|_\rho\}$ is used, and which depends on the specific $[\sigma, \rho]$ -domination problem that we are solving.

Forget Node. Let i be a forget node of T with child node j . Let $X_i = X_j \setminus \{v\}$ for some $v \in X_j$.

By definition of G_i , G_i contains edges between v and vertices in X_i while G_j does not. To account for these edges, we start by updating the given table A_j such that it accounts for the additional edges: that is, for an edge $\{u, v\}$ with $u \in X_i$, we adjust the counts of the number of neighbours expressed in the state colourings for u and v . We do so before we construct table A_i .

Let $[\mathbf{c}, c_u, c_v] \in C^{X_j}$ be such that c_u and c_v are labels for u and v respectively. For every edge $\{u, v\}$ with $u \in X_i$, we update A_j twice, once for u and once for v . We update A_j for u as follows:

$$A_j([\mathbf{c}, c_u, c_v]) := \begin{cases} A_j([\mathbf{c}, c_u, c_v]) & \text{if } c_v \in C_\rho \\ \infty & \text{if } c_v \in C_\sigma, c_u \in \{0|_\rho, 0|_\sigma\} \\ A_j([\mathbf{c}, |l - 1|_\rho, c_v]) & \text{if } c_v \in C_\sigma, c_u = |l|_\rho, l > 0 \\ A_j([\mathbf{c}, |l - 1|_\sigma, c_v]) & \text{if } c_v \in C_\sigma, c_u = |l|_\sigma, l > 0 \\ \min\{A_j([\mathbf{c}, |l - 1|_\rho, c_v]), A_j([\mathbf{c}, | \geq \ell|_\rho, c_v])\} & \text{if } c_v \in C_\sigma, c_u = | \geq \ell|_\rho \\ \min\{A_j([\mathbf{c}, |l - 1|_\sigma, c_v]), A_j([\mathbf{c}, | \geq \ell|_\sigma, c_v])\} & \text{if } c_v \in C_\sigma, c_u = | \geq \ell|_\sigma \end{cases}$$

No update needs to be done if v is not in the partial solution D (first line). If c_u indicates that u has no neighbours in D while $v \in D$, then no such partial solution exists (second line). Otherwise, the counts in the label of u need to account for the extra neighbour. In the last four lines, we perform the required label update for all other labels giving special attention to the case where a $|\geq \ell|_\sigma$ or $|\geq \ell|_\rho$ label is used. Here, the minimum needs to be taken over two equivalence classes that through the added edge become equivalent: we take the minimum because we are solving the minimisation variant. Updating A_j for v goes identically with the roles of u and v switched, and as stated above, we perform this update for all edges incident to v in G_j .

Next, we compute A_i and start keeping track of equivalence classes based on X_i instead of based on X_j . To do so, we select a dominating solution from the partial solution equivalence classes for which v has a number of neighbours in D that corresponds to the specific $[\sigma, \rho]$ -domination problem:

$$A_i(\mathbf{c}) = \min_{c_v \text{ a valid label}} A_j([\mathbf{c}, c_v])$$

Here, a valid label c_v is any label that corresponds to having the correct number of neighbours in D as defined by the specific $[\sigma, \rho]$ -domination problem: c_v is a label $|\ell|_\sigma$ or $|\ell|_\rho$ for which $\ell \in \sigma$ or $\ell \in \rho$, respectively, or c_v is a label $|\geq \ell|_\rho$ or $|\geq \ell|_\sigma$ in case of cofinite σ or ρ .

Join Node. Let A_i be the memoisation table for a join node i of T with child nodes l and r . Here we give a simple algorithm for the join node; in Sect. 4, we survey more involved approaches.

A trivial algorithm to compute A_i would loop over all pairs of state colourings $\mathbf{c}_l, \mathbf{c}_r$ of X_i that agree on which vertices are in the solution set D , and then consider two corresponding partial solutions D_l on G_l and D_r on G_r and infer the state colouring \mathbf{c}_i of the partial solution $D_l \cup D_r$ on G_i . It then stores in A_i the minimum size of a solution for each equivalence class for G_i .

Note that the agreement on which vertices are in D is necessary for $D_l \cup D_r$ to be a valid partial solution: otherwise vertices that are no longer in X_i can obtain additional neighbours in D . At the same time the agreement is not a too tight restriction as any partial solution D on G_i can trivially be decomposed into partial solutions on G_l and G_r that agree on which vertices on X_i are in D .

Root Node. In the root node r of T (which is a forget node), $X_r = \emptyset, G_r = G$ and consequently $A_r(\square)$ is the minimum size of a $[\sigma, \rho]$ -dominating set on G . The result we set out to compute!

Lemma 1. *Let \mathcal{P} be the minimisation variant of a $[\sigma, \rho]$ -domination problem with label set C using $s = |C|$ labels. Let \mathcal{A} be an algorithm for the computations in a join node for problem \mathcal{P} that, given a join node i with $|X_i| = k$ and the memoisation tables A_l and A_r for its child nodes, computes the memoisation table A_i in $\mathcal{O}(f(n, k))$ arithmetic operations. Then, given a graph G with a tree*

decomposition T of width t , \mathcal{P} can be solved on G in $\mathcal{O}((s^{t+1}t + f(n, t + 1))n)$ arithmetic operations.

Proof. First transform T into a nice tree decomposition T' with $\mathcal{O}(n)$ nodes. If we show that the table A_j associated to any node j of T' can be computed in $\mathcal{O}(s^k k + f(n, k))$ arithmetic operations, then the result follows as $k \leq t + 1$. Consider the recurrences in the dynamic programming algorithm exposed above. The result trivially holds for leaf and root nodes, and also for the join nodes by definition of \mathcal{A} . It is easy to see that in the recurrences for the introduce and forget nodes, every value is computed using a constant amount of work. Since the tables are of size s^k , and for a forget node we need to do at most k update steps as we can add at most k edges, the result follows. \square

It is not difficult to modify the above algorithm to obtain:

Proposition 1. *Lemma 1 holds irrespective of \mathcal{P} being a existence, maximisation, minimisation, counting, counting minimisation or counting maximisation variant of a $[\sigma, \rho]$ -domination problem.*

3 Overview of Fast Transforms

To obtain fast algorithms for the computations in join nodes of a nice tree decomposition, we use several well-known algebraic transforms, specifically the Möbius transform and the Fourier transform. We opted for a reasonably extensive coverage of this standard material because of completeness reasons and because the details matter for some of the arguments in Sects. 4 and 5.

Recall that, in the introduction on dynamic programming for $[\sigma, \rho]$ -domination problems, we stored integers in the domain $\{0, 1, \dots, M\}$ for some large integer M . We present the algebraic transforms using computations in \mathbb{F}_p , the field of integers modulo a prime number p . Since we know that, for a join node i with child nodes l, r , all values in the memoisation tables A_i, A_l and A_r are in $\{0, 1, \dots, M\}$, we can do the computations in \mathbb{F}_p as long as $p > M$.

In the literature, the discrete Fourier transform is often defined on sequences in \mathbb{C} . We choose \mathbb{F}_p to avoid any analysis of rounding errors, especially when we combine it with the use of zeta and Möbius transforms. Using \mathbb{F}_p does require that p is chosen appropriately: \mathbb{F}_p must contain certain roots of unity required for the Fourier transforms. In the statements of definitions, propositions and lemmas in this section, we will sometimes say that p is *chosen appropriately* to state that \mathbb{F}_p contains the roots of unity required in the definition or in the following proof. A short discussion on how to choose a proper prime number p such that this condition is satisfied can be found in Sect. 3.3.

3.1 The Discrete Fourier Transforms Using Modular Arithmetic

Definition 4 (discrete Fourier transform). *Let $\mathbf{a} = (a_i)_{i=0}^{r-1}$ be a sequence of numbers in \mathbb{F}_p , and let ω_r be an r -th root of unity in \mathbb{F}_p . The discrete Fourier*

transform and inverse discrete Fourier transform are transformations between sequences of length r in \mathbb{F}_p defined as follows:

$$DFT(\mathbf{a})_i = \sum_{j=0}^{r-1} \omega_r^{ij} a_j \qquad DFT^{-1}(\mathbf{a})_i = \frac{1}{r} \sum_{j=0}^{r-1} \omega_r^{-ij} a_j$$

Recall that an r -th root of unity is an element $x \in \mathbb{F}_p$ such that $x^r = 1$ while $x^l \neq 1$ for all $l < r$.

These two transformations are inverses as their names suggest.

Proposition 2. $DFT^{-1}(DFT(\mathbf{a}))_i = a_i$

Proof. In the derivation below, we first fill in the definitions and rearrange the terms (1). Then, we split the sum based on $k = i$ and $k \neq i$ (from 1 to 2).

$$\begin{aligned} DFT^{-1}(DFT(\mathbf{a}))_i &= \frac{1}{r} \sum_{j=0}^{r-1} \omega_r^{-ij} \sum_{k=0}^{r-1} \omega_r^{jk} a_k = \frac{1}{r} \sum_{k=0}^{r-1} a_k \sum_{j=0}^{r-1} (\omega_r^{k-i})^j & (1) \\ &= \frac{1}{r} a_i \sum_{j=0}^{r-1} (\omega_r^{i-i})^j + \frac{1}{r} \sum_{\substack{k=0 \\ k \neq i}}^{r-1} a_k \sum_{j=0}^{r-1} (\omega_r^{k-i})^j = a_i \frac{1}{r} \sum_{j=0}^{r-1} 1 + \frac{1}{r} \sum_{\substack{k=0 \\ k \neq i}}^{r-1} a_k \cdot 0 = a_i & (2) \end{aligned}$$

Finally, we use that the first part of the sum is trivial as $\omega_r^{i-i} = \omega_r^0 = 1$, while the second part cancels as $\sum_{j=0}^{r-1} (\omega_r^{k-i})^j = \frac{1-\omega_r^{(k-i)r}}{1-\omega_r^{k-i}}$ is a geometric series with $\omega_r^{(k-i)r} = (\omega_r^r)^{k-i} = 1^{k-i} = 1$. □

There exist fast algorithms for the discrete Fourier transform and its inverse, called fast Fourier transforms (FFT's), e.g., see the Cooley-Tukey FFT algorithm [13] and Rader's FFT algorithm [24]. These algorithms are not particularly difficult to understand, but beyond the scope of this paper.

Proposition 3 (fast Fourier transform). *The discrete Fourier transform and its inverse for sequences of length r can be computed in $\mathcal{O}(r \log r)$ arithmetic operations.*

The definition of the discrete Fourier transform can be naturally extended from sequences to higher dimensional structures. Let \mathbb{Z}_r be the commutative ring of integers modulo r (here the modulus r can be non-prime), and let \mathbb{Z}_r^k be the \mathbb{Z}_r -module of k -tuples with elements from \mathbb{Z}_r .

Definition 5 (multidimensional discrete Fourier transform). *Let $Z = \mathbb{Z}_{r_1} \times \mathbb{Z}_{r_2} \times \dots \times \mathbb{Z}_{r_k}$, and let $R = \prod_{i=1}^k r_i$. Also, let $\mathbf{A} = (a_{\mathbf{x}})_{\mathbf{x} \in Z}$ be a tensor of rank k with elements in \mathbb{F}_p indexed by the k -tuple $\mathbf{x} = [x_1, x_2, \dots, x_k]$, where p is chosen appropriately. The multidimensional discrete Fourier transform and inverse multidimensional discrete Fourier transform are defined as follows:*

$$\begin{aligned}
 DFT_k(\mathbf{A})_{\mathbf{x}} &= \sum_{y_1=0}^{r_1-1} \omega_{r_1}^{x_1 y_1} \sum_{y_2=0}^{r_2-1} \omega_{r_2}^{x_2 y_2} \dots \sum_{y_k=0}^{r_k-1} \omega_{r_k}^{x_k y_k} a_{\mathbf{y}} \\
 DFT_k^{-1}(\mathbf{A})_{\mathbf{x}} &= \frac{1}{R} \sum_{y_1=0}^{r_1-1} \omega_{r_1}^{-x_1 y_1} \sum_{y_2=0}^{r_2-1} \omega_{r_2}^{-x_2 y_2} \dots \sum_{y_k=0}^{r_k-1} \omega_{r_k}^{-x_k y_k} a_{\mathbf{y}}
 \end{aligned}$$

When $r = r_1 = r_2 = \dots = r_k$, this simplifies to the following:

$$DFT_k(\mathbf{A})_{\mathbf{x}} = \sum_{\mathbf{y} \in \mathbb{Z}_r^k} \omega_r^{\mathbf{x} \cdot \mathbf{y}} a_{\mathbf{y}} \qquad DFT_k^{-1}(\mathbf{A})_{\mathbf{x}} = \frac{1}{r^k} \sum_{\mathbf{y} \in \mathbb{Z}_r^k} \omega_r^{-\mathbf{x} \cdot \mathbf{y}} a_{\mathbf{y}}$$

where the expressions in the exponents are the dot products on the tuples \mathbf{x} and \mathbf{y} in \mathbb{Z}_r^k .

Note that the dot products are in exponents of which the base is an r -th root of unity, hence they are computed modulo r : this agrees with the notation where \mathbf{x} and \mathbf{y} are taken from \mathbb{Z}_r^k .

Proposition 4 (fast multidimensional discrete Fourier transform). *Let $Z = \mathbb{Z}_{r_1} \times \mathbb{Z}_{r_2} \times \dots \times \mathbb{Z}_{r_k}$, and let $R = \prod_{i=1}^k r_i$. Also, let \mathbf{A} be a tensor of rank k with elements in \mathbb{F}_p , $\mathbf{A} = (a_{\mathbf{x}})_{\mathbf{x} \in Z}$, where p is chosen appropriately. The multidimensional discrete Fourier transform and inverse multidimensional discrete Fourier transform of \mathbf{A} can be computed in $\mathcal{O}(R \log(R))$ time.*

Proof. Denote by $\mathbf{x}[x_i \leftarrow y]$ the tuple \mathbf{x} with the i -th coordinate of \mathbf{x} replaced by y . We compute $DFT_k(\mathbf{A})$ with an algorithm that uses k -steps. Let $\mathbf{A}_0 = \mathbf{A}$. At the i -th step of the algorithm, let:

$$(\mathbf{A}_i)_{\mathbf{x}} = \sum_{j=0}^{r_i-1} \omega_{r_i}^{x_i j} a_{\mathbf{x}[x_i \leftarrow j]}$$

Notice that if $k = 1$, this formula equals the one dimensional discrete Fourier transform. It is not hard to see that \mathbf{A}_k is the k -dimensional Fourier transform of \mathbf{A} : if one repeatedly substitutes the formula for \mathbf{A}_{i-1} in the formula for \mathbf{A}_i starting at $i = k$, one obtains the (non-simplified) formula for the k -dimensional Fourier transform in Definition 5.

For the inverse multidimensional Fourier transform, almost the same procedure can be followed. Let $\mathbf{A}_0 = \mathbf{A}$ and use the following formula at the i -th step, finally obtaining the result \mathbf{A}_k . Here, again if $k = 1$, this formula equals the one dimensional inverse discrete Fourier transform.

$$(\mathbf{A}_i)_{\mathbf{x}} = \frac{1}{r_i} \sum_{j=0}^{r_i-1} \omega_{r_i}^{-x_i j} a_{\mathbf{x}[x_i \leftarrow j]}$$

For the running time, notice that step i preforms $\frac{R}{r_i}$ standard 1-dimensional (inverse) discrete Fourier transforms on a sequence of length r_i . By Proposition 3 this can be done in $\mathcal{O}(\frac{R}{r_i} r_i \log(r_i)) = \mathcal{O}(R \log(r_i))$ time. This leads to a total running time of $\mathcal{O}(R \sum_{i=1}^k \log(r_i)) = \mathcal{O}(R \log(R))$. \square

In the above proof, the sequences $\mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_k$ are created using 1-dimensional (inverse) discrete Fourier transforms. Because the 1-dimensional discrete Fourier transform and 1-dimensional inverse discrete Fourier transform are inverses, it directly follows that the sequence $\mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_k$ used in the k -dimensional discrete Fourier transform algorithm equals the sequence $\mathbf{A}_k, \mathbf{A}_{k-1}, \dots, \mathbf{A}_0$ used in the inverse k -dimensional discrete Fourier transform. I.e., as the name suggests, the k -dimensional inverse discrete Fourier transform is the inverse of the k -dimensional discrete Fourier transform.

We mainly use the multidimensional fast discrete Fourier transform in combination with the well-known convolution theorem.

Lemma 2 (multidimensional convolution theorem). *Let $Z = \mathbb{Z}_{r_1} \times \mathbb{Z}_{r_2} \times \dots \times \mathbb{Z}_{r_k}$, and let $\mathbf{A} = (a_x)_{x \in Z}$, $\mathbf{B} = (b_x)_{x \in Z}$ be tensors of rank k with elements in \mathbb{F}_p , where p is chosen appropriately. Let the tensor multiplication $\mathbf{A} \cdot \mathbf{B}$ be defined point wise, and let for a_x and b_y the sum $\mathbf{x} + \mathbf{y}$ be defined as the sum in Z (coordinate-wise with the i -th coordinate modulo r_i). Then:*

$$DFT_k^{-1}(DFT_k(\mathbf{A}) \cdot DFT_k(\mathbf{B}))_x = \sum_{z_1+z_2 \equiv x} a_{z_1} b_{z_2}$$

Proof. We prove the lemma for the simplified case where $r = r_1 = r_2 = \dots = r_k$ and hence $Z = \mathbb{Z}_r^k$, the more general case goes analogously but is notation-wise much more tedious as one needs to differentiate between multiple moduli and their corresponding roots of unity.

The proof follows the same pattern as in Proposition 2. That is, we first fill in the definitions (3) and rearrange the terms (4). Next (5), we observe that the sum over all $\mathbf{j} \in \mathbb{Z}_r^k$ can be written as a product of k smaller sums, each involving but one coordinate of \mathbf{j} .

$$DFT_k^{-1}(DFT_k(\mathbf{A}) \cdot DFT_k(\mathbf{B}))_x = \frac{1}{r^k} \sum_{y \in \mathbb{Z}_r^k} \omega_r^{-x \cdot y} \left(\sum_{z_1 \in \mathbb{Z}_r^k} \omega_r^{y \cdot z_1} a_{z_1} \right) \left(\sum_{z_2 \in \mathbb{Z}_r^k} \omega_r^{y \cdot z_2} b_{z_2} \right) \tag{3}$$

$$= \frac{1}{r^k} \sum_{z_1, z_2 \in \mathbb{Z}_r^k} a_{z_1} b_{z_2} \sum_{y \in \mathbb{Z}_r^k} \omega_r^{y \cdot (z_1 + z_2 - x)} \tag{4}$$

$$= \frac{1}{r^k} \sum_{z_1, z_2 \in \mathbb{Z}_r^k} a_{z_1} b_{z_2} \prod_{i=1}^k \left(\sum_{j=0}^{r-1} \omega_r^{j \cdot ((z_1)_i + (z_2)_i - x_i)} \right) \tag{5}$$

Here, x_i , $(z_1)_i$ and $(z_2)_i$ are the i -th components of \mathbf{x} , \mathbf{z}_1 and \mathbf{z}_2 respectively.

When $x_i \equiv (z_1)_i + (z_2)_i$ modulo r in the parenthesised sum of Eq. 5, this sum becomes $\sum_{j=0}^{r-1} \omega_r^0$ and thus equals r . Otherwise, when $x_i \not\equiv (z_1)_i + (z_2)_i$

the parenthesised sum is again a geometric series: $\sum_{j=0}^{r-1} (\omega_r^{(z_1)_i+(z_2)_i-x_i})^j$ that solves to $\frac{1-(\omega_r^{(z_1)_i+(z_2)_i-x_i})^r}{1-\omega_r^{(z_1)_i+(z_2)_i-x_i}} = 0$ as $(\omega_r^{(z_1)_i+(z_2)_i-x_i})^r = (\omega_r^r)^{(z_1)_i+(z_2)_i-x_i} = 1$ in the numerator.

Continuing from (5), we obtain:

$$\begin{aligned} \text{DFT}_k^{-1}(\text{DFT}_k(\mathbf{A}) \cdot \text{DFT}_k(\mathbf{B}))_{\mathbf{x}} &= \frac{1}{r^k} \sum_{z_1, z_2 \in \mathbb{Z}_r^k} a_{z_1} b_{z_2} \prod_{i=1}^k r[x_i = (z_1)_i + (z_2)_i] & (6) \\ &= \sum_{z_1, z_2 \in \mathbb{Z}_r^k} a_{z_1} b_{z_2} \prod_{i=1}^k [x_i = (z_1)_i + (z_2)_i] = \sum_{z_1+z_2 \equiv \mathbf{x}} a_{z_1} b_{z_2} & (7) \end{aligned}$$

completing the proof. □

Taking all the above together, we finally obtain the following result. To distinguish from \mathbb{Z}_r , let $\mathbb{N}_{<r} = \{0, 1, \dots, r-1\}$ be the integers up to r with standard operators without modulus (operations for which the result of standard operations on $\mathbb{N}_{<r}$ is outside $\mathbb{N}_{<r}$ are considered undefined, i.e., $2+2$ is undefined in $\mathbb{N}_{<3}$).

Lemma 3 (cyclic and non-cyclic convolution). *Let $N = \mathbb{N}_{<q_1} \times \mathbb{N}_{<q_2} \times \dots \times \mathbb{N}_{<q_l}$, and let $Q = \prod_{i=1}^l q_i$. Let $Z = \mathbb{Z}_{r_1} \times \mathbb{Z}_{r_2} \times \dots \times \mathbb{Z}_{r_k}$, and let $R = \prod_{i=1}^k r_i$. Let $f, g : Z \times N \rightarrow \mathbb{F}_p$, where p is chosen appropriately. And, let $h : Z \times N \rightarrow \mathbb{F}_p$ be the combined (partially cyclic and partially non-cyclic) convolution of f and g defined as:*

$$h(\mathbf{x}, \mathbf{i}) = \sum_{\mathbf{y}_1 + \mathbf{y}_2 \equiv \mathbf{x}} \sum_{\mathbf{j}_1 + \mathbf{j}_2 = \mathbf{i}} f(\mathbf{y}_1, \mathbf{j}_1) g(\mathbf{y}_2, \mathbf{j}_2)$$

where the sum $\mathbf{y}_1 + \mathbf{y}_2 \equiv \mathbf{x}$ is evaluated component-wise modulo r_i at coordinate i (sum in Z), and the sum $\mathbf{j}_1 + \mathbf{j}_2 = \mathbf{i}$ is evaluated component-wise without modulus (sum in N). Then, the combined convolution h can be computed in $\mathcal{O}(RQ2^l(\log(R) + \log(Q) + l))$ arithmetic operations.

Proof. We reduce the problem to a standard multidimensional convolution (with modulus) by padding the input with zeroes. To be precise, let $Z' = \mathbb{Z}_{2q_1} \times \mathbb{Z}_{2q_2} \times \dots \times \mathbb{Z}_{2q_l}$ (N with for each coordinate twice as many values and with modulo additions), and let $f', g' : Z \times Z' \rightarrow \mathbb{F}_p$ be equal to f and g on the intersection of their domains (where N is interpreted as subset of Z' by interpreting each $\mathbb{N}_{<q_i}$ as subset of \mathbb{Z}_{2q_i}) and zero otherwise. Use Proposition 4 and Lemma 2 to compute the standard multidimensional convolution of f' and g' . Because $Z \times Z'$ has $RQ2^l$ elements, this requires $\mathcal{O}(RQ2^l(\log(R) + \log(Q) + l))$ arithmetic operations. Because the padded zeroes prevent the circular convolution effect, we can extract h by taking the restriction of the result to $Z \times N$. □

Different than for previous propositions and lemmas, we have more freedom in choosing the prime p that is 'chosen appropriately' in the lemma above. For the given proof, appropriate means that in \mathbb{F}_p all r_i -th roots of unity exists

and all $2q_i$ -th roots of unity exist. However in our applications of the lemma, l is often fixed. This means that the running time does not change if we allow $Z' = \mathbb{Z}_{s_1} \times \mathbb{Z}_{s_2} \times \dots \times \mathbb{Z}_{s_l}$ with, for all i , $2q_i \leq s_i \leq cq_i$ for a small constant c . In other words, with respect to the different q_i , there must be some root of unity, but the order of this root of unity has a broad range in which it is acceptable for our results to be valid.

Corollary 1 (multidimensional non-cyclic convolution). *Let $N = \mathbb{N}_{<q_1} \times \mathbb{N}_{<q_2} \times \dots \times \mathbb{N}_{<q_l}$, and let $Q = \prod_{i=1}^l q_i$. Let $f, g : N \rightarrow \mathbb{F}_p$, where p is chosen appropriately. Let $h : N \rightarrow \mathbb{F}_p$ be the non-cyclic convolution of f and g defined as:*

$$h(\mathbf{i}) = \sum_{\mathbf{j}_1 + \mathbf{j}_2 = \mathbf{i}} f(\mathbf{j}_1)g(\mathbf{j}_2)$$

where the sum $\mathbf{j}_1 + \mathbf{j}_2 = \mathbf{i}$ is evaluated component-wise without modulus (sum in N). Then, h can be computed in $\mathcal{O}(Q 2^l(\log(Q) + l))$ arithmetic operations.

Proof. Direct consequence of Lemma 3 with $k = 0$. □

3.2 Möbius Inversion Using Fast Zeta and Fast Möbius Transforms

The zeta and Möbius transforms apply to functions on partially ordered sets.

Definition 6 (zeta and Möbius transform). *Let P be a partially ordered set. Given a function $f : P \rightarrow \mathbb{F}_p$, the zeta transform $\zeta(f)$ and the Möbius transform $\mu(f)$ are defined as follows:*

$$\zeta(f)(x) = \sum_{y \leq x} f(y) \quad \mu(f)(x) = \sum_{y \leq x} \mu(y, x)f(y)$$

where
$$\mu(x, y) = \begin{cases} 1 & \text{if } x = y \\ -\sum_{x < z \leq y} \mu(z, y) & \text{if } x < y \end{cases}$$

The recursively defined function $\mu(x, y)$ on pairs $x, y \in P$ with $x \leq y$ is the Möbius function of P .

The zeta transform $\zeta(f)$ and the Möbius transform are inverses, as we will now show.

Lemma 4 (Möbius inversion). *Let $f : P \rightarrow \mathbb{F}_p$ any function, then $\mu(\zeta(f))(x) = f(x)$.*

Proof. Let $x, y \in P$ and consider the sum $\sum_{x \leq z \leq y} \mu(z, y)$. If $x = y$, then this sum equals $\mu(x, x) = 1$. If $x < y$, then this sum equals $\mu(x, y) + \sum_{x < z \leq y} \mu(z, y) = 0$ by definition of $\mu(x, y)$. As such:

$$\mu(\zeta(f))(x) = \sum_{y \leq x} \mu(y, x) \sum_{z \leq y} f(z) = \sum_{z \leq x} f(z) \sum_{z \leq y \leq x} \mu(y, x) = \sum_{z \leq x} f(z)[z = x] = f(x)$$

The first equality is by expanding the definitions. The second follows by reordering terms. And, the third follows from the above, where $[z = x]$ is Iverson notation that is 1 if $z = x$ and 0 otherwise. □

In this paper, we will not define any Möbius transform explicitly. We will show how to compute zeta transforms $\zeta(f)$ of functions $f : P \rightarrow \mathbb{F}_p$ for some partial orders P . Then, given $\zeta(f)$, we show that we can reconstruct f . This reconstruction (implicitly) is an algorithm for the Möbius transform because a consequence of Lemma 4 is that the zeta transform has a unique inverse.

Möbius inversion is often used in relation to lattices. A *meet-semilattice* is a partial order P on which, for any two elements in $x, y \in P$, the *meet* $x \wedge y$ (greatest lower bound) is properly defined. Similarly, a *join-semilattice* is a partial order P set on which, for any two elements in $x, y \in P$, the *join* $x \vee y$ (smallest upper bound) is properly defined. A lattice is a partial order P that is both a meet and a join semi-lattice. An example is the finite lattice $\mathbb{N}_{<r}^k$ with the coordinate-wise natural order and where the meet and join are the coordinate-wise minimum and maximum.

We will use Möbius inversion on partial orders that are Cartesian products P^k of a smaller partial order P . For $\mathbf{x}, \mathbf{y} \in P^k$, $\mathbf{x} = [x_1, x_2, \dots, x_k]$, $\mathbf{y} = [y_1, y_2, \dots, y_k]$, we write $\mathbf{x} \leq \mathbf{y}$ if and only if $x_i \leq y_i$ for all i . Additionally, our partial orders have the property that for every $x \in P$, the downward closed set $\{y \in P \mid y \leq x\}$ forms a join-semilattice. It is not hard to see that if for every $x \in P$, $\{y \in P \mid y \leq x\}$ forms a join-semilattice, then for every $\mathbf{x} \in P^k$, $\{\mathbf{y} \in P^k \mid \mathbf{y} \leq \mathbf{x}\}$ forms a join-semilattice as well, where the join operation is defined coordinate-wise.

On the subset lattice (isomorphic to $\mathbb{N}_{<2}^k$) there are well-known fast algorithms for the zeta and Möbius transforms, often referred to as Yates’ algorithm [32], see also [2, 19]. Below, we generalise these algorithms to partial orders P^k for which, for every $\mathbf{x} \in P^k$, the set $\{\mathbf{y} \in P^k \mid \mathbf{y} \leq \mathbf{x}\}$ forms a join-semilattice. For fast zeta and Möbius transforms on arbitrary finite lattices, see [3].

Proposition 5 (fast zeta and Möbius transforms). *The zeta transform and Möbius transform of a function $f : \mathbb{N}_{<r}^k \rightarrow \mathbb{F}_p$ can be computed in $\mathcal{O}(r^k k)$ arithmetic operations.*

Proof. We compute $\zeta(f)$ with an algorithm that uses k steps. Let $f_0 = f$, and let $\mathbf{x} = [x_1, \dots, x_k]$. Denote by $\mathbf{x}[x_i \leftarrow y]$ the tuple \mathbf{x} with the value on the i -th coordinate replaced by y . At the i -th step of the algorithm, we compute f_i recursively using the left formula below.

$$f_i(\mathbf{x}) = \begin{cases} f_{i-1}(\mathbf{x}) & \text{if } x_i = 0 \\ f_i(\mathbf{x}[x_i \leftarrow x_i - 1]) + f_{i-1}(\mathbf{x}) & \text{if } x_i > 0 \end{cases} \quad f_i(\mathbf{x}) = \sum_{j \leq x_i} f_{i-1}(\mathbf{x}[x_i \leftarrow j]) \tag{8}$$

The right formula above follows by induction on the left recurrence. By induction on the step number i , one easily sees that f_i satisfies the equation below, from which we can obtain $\zeta(f)$ since $f_k = \zeta(f)$. The result for $\zeta(f)$ follows because each step computes r^k values, each in constant time.

$$f_i(\mathbf{x}) = \sum_{y_1 \leq x_1} \sum_{y_2 \leq x_2} \cdots \sum_{y_i \leq x_i} f([y_1, y_2, \dots, y_i, x_{i+1}, x_{i+2}, \dots, x_k])$$

For $\mu(f)$, we use that $\mu(f)$ is the inverse of $\zeta(f)$: the sequence f_0, f_1, \dots, f_k used to compute $\zeta(f)$ from f can computationally be inverted to compute f from $\zeta(f)$. That is, let $f_k = \zeta(f)$, and let:

$$f_i(\mathbf{x}) = \begin{cases} f_{i+1}(\mathbf{x}) - f_{i+1}(\mathbf{x}[x_i \leftarrow x_i - 1]) & \text{if } x_i > 0 \\ f_{i+1}(\mathbf{x}) & \text{if } x_i = 0 \end{cases} \tag{9}$$

Assuming that the f_i were computed using Eq. 8, and substituting the right part of (8) into the case where $x_i > 0$ in Eq. 9, we see that (9) computes the inverse of Eq. 8:

$$\left(\sum_{j \leq x_i} f_{i-1}(\mathbf{x}[x_i \leftarrow j]) \right) - \left(\sum_{j \leq x_i - 1} f_{i-1}(\mathbf{x}[x_i \leftarrow j]) \right) = f_i(\mathbf{x})$$

Hence, we can reconstruct $f_0 = f$ again. Since the run time is the same, we obtain the result. □

It is easy to generalise the inductive proof above and obtain:

Lemma 5. *Given algorithms for the zeta and Möbius transform for functions $f : P \rightarrow \mathbb{F}_p$ that use $\mathcal{O}(|P|)$ arithmetic operations, there are algorithms for the zeta and Möbius transform for functions $f : P^k \rightarrow \mathbb{F}_p$ that require $\mathcal{O}(|P^k|k)$ arithmetic operations.*

The application of the zeta and Möbius transform that is important to us is the following lemma.

Lemma 6 (generalised covering product). *Let P be a finite partial order such that, for every $\mathbf{x} \in P^k$, the set $\{\mathbf{y} \in P^k \mid \mathbf{y} \leq \mathbf{x}\}$ forms a join-semilattice, and let $f, g : P^k \rightarrow \mathbb{F}_p$.*

Define the generalised covering product $h : P^k \rightarrow \mathbb{F}_p$ of f and g through:

$$h(\mathbf{x}) = \sum_{\mathbf{y}_1 \vee \mathbf{y}_2 = \mathbf{x}} f(\mathbf{y}_1) g(\mathbf{y}_2)$$

Then $\mu(\zeta(f) \cdot \zeta(g))(\mathbf{x}) = h(\mathbf{x})$, where the product $\zeta(f) \cdot \zeta(g)$ is defined by point-wise multiplication.

Proof. We will prove that $(\zeta(f) \cdot \zeta(g))(\mathbf{x}) = \zeta(h)(\mathbf{x})$, then the result follows from Lemma 4.

$$(\zeta(f) \cdot \zeta(g))(\mathbf{x}) = \left(\sum_{\mathbf{y} \leq \mathbf{x}} f(\mathbf{y}) \right) \left(\sum_{\mathbf{y} \leq \mathbf{x}} g(\mathbf{y}) \right) = \sum_{\mathbf{y}_1, \mathbf{y}_2 \leq \mathbf{x}} f(\mathbf{y}_1) g(\mathbf{y}_2) \tag{10}$$

Here, we first use the definition of the ζ -transform and then work out all the product terms. The result equals $\zeta(h)(\mathbf{x})$ as we now show by working out the definition of $\zeta(h)(\mathbf{x})$.

$$\zeta(h)(\mathbf{x}) = \sum_{z \leq \mathbf{x}} \sum_{\mathbf{y}_1 \vee \mathbf{y}_2 = z} f(\mathbf{y}_1) g(\mathbf{y}_2) = \sum_{\mathbf{y}_1, \mathbf{y}_2 \leq \mathbf{x}} f(\mathbf{y}_1) g(\mathbf{y}_2) \tag{11}$$

For the last equality, we reorder terms using that for any two $\mathbf{y}_1, \mathbf{y}_2 \leq \mathbf{x}$ there is a unique \mathbf{z} such that $\mathbf{y}_1 \vee \mathbf{y}_2 = \mathbf{z}$; this is well defined as the set $\{\mathbf{y} \in P^k \mid \mathbf{y} \leq \mathbf{x}\}$ forms a join-semilattice. \square

As a direct result, we obtain a generalisation of the covering product from [2].

Corollary 2. *The generalised covering product for $f, g : \mathbb{N}_{<r}^k \rightarrow \mathbb{F}_p$ defined in the statement of Lemma 6 can be computed in $\mathcal{O}(r^k k)$ arithmetic operations.*

Proof. Combine Lemma 6 with the fast evaluation algorithms of Proposition 5. \square

We conclude this part on zeta and Möbius transforms by a theorem on a combined covering product and convolution product that we will use in the sections to come.

Theorem 1. *Let P be a finite partial order where, for every $\mathbf{x} \in P^k$, the set $\{\mathbf{y} \in P^k \mid \mathbf{y} \leq \mathbf{x}\}$ forms a join-semilattice. Let $N = \mathbb{N}_{<q_1} \times \mathbb{N}_{<q_2} \times \dots \times \mathbb{N}_{<q_l}$, and let $Q = \prod_{i=1}^l q_i$. Let $f, g : P^k \times N \rightarrow \mathbb{F}_p$, where p is chosen appropriately. Define $h : P^k \times N \rightarrow \mathbb{F}_p$ as follows:*

$$h(\mathbf{x}, \mathbf{i}) = \sum_{\mathbf{y}_1 \vee \mathbf{y}_2 = \mathbf{x}} \sum_{\mathbf{j}_1 + \mathbf{j}_2 = \mathbf{i}} f(\mathbf{y}_1, \mathbf{j}_1) g(\mathbf{y}_2, \mathbf{j}_2)$$

If P allows zeta and Möbius transforms using $\mathcal{O}(|P|)$ arithmetic operations for functions $f' : P \rightarrow \mathbb{F}_p$, then h can be computed in $\mathcal{O}(|P|^k Q (k + \log(Q)))$ arithmetic operations.

In particular, if $P^k = \mathbb{N}_{<r}^k$, then h can be computed in $\mathcal{O}(r^k Q (k + \log(Q)))$ arithmetic operations.

Proof. For the functions f, g, h , with domains $P^k \times N$, we write $\zeta(f(-, \mathbf{i}))(\mathbf{x}) = \sum_{\mathbf{y} \leq \mathbf{x}} f(\mathbf{y}, \mathbf{i})$ to fix the second component when using the zeta transform. Following the same reasoning as in Eqs. 11 and 10 in the proof of Lemma 6, one easily obtains:

$$\begin{aligned} \zeta(h(-, \mathbf{i}))(\mathbf{x}) &= \sum_{\mathbf{y} \leq \mathbf{x}} \sum_{\mathbf{z}_1 \vee \mathbf{z}_2 = \mathbf{y}} \sum_{\mathbf{j}_1 + \mathbf{j}_2 = \mathbf{i}} f(\mathbf{z}_1, \mathbf{j}_1) g(\mathbf{z}_2, \mathbf{j}_2) \\ &= \sum_{\mathbf{j}_1 + \mathbf{j}_2 = \mathbf{i}} \left(\sum_{\mathbf{y} \leq \mathbf{x}} \sum_{\mathbf{z}_1 \vee \mathbf{z}_2 = \mathbf{y}} f(\mathbf{z}_1, \mathbf{j}_1) g(\mathbf{z}_2, \mathbf{j}_2) \right) \\ &= \sum_{\mathbf{j}_1 + \mathbf{j}_2 = \mathbf{i}} \left(\sum_{\mathbf{z}_1, \mathbf{z}_2 \leq \mathbf{x}} f(\mathbf{z}_1, \mathbf{j}_1) g(\mathbf{z}_2, \mathbf{j}_2) \right) \\ &= \sum_{\mathbf{j}_1 + \mathbf{j}_2 = \mathbf{i}} (\zeta(f(-, \mathbf{j}_1)) \cdot \zeta(g(-, \mathbf{j}_2))) (\mathbf{x}) \end{aligned}$$

Consequently, we can compute h by evaluating this expression and taking the Möbius transform.

That is, we can compute h by taking the following steps:

1. Compute a fast zeta transform of $f(-, \mathbf{j})$ and $g(-, \mathbf{j})$, for each fixed $\mathbf{j} \in N$ in $\mathcal{O}(|P|^k k)$ arithmetic operations using Lemma 5. This takes $\mathcal{O}(|P|^k Q k)$ arithmetic operations in total. For each $\mathbf{j}_1, \mathbf{j}_2 \in N$ and $\mathbf{x} \in P^k$, we now have $\zeta(f(-, \mathbf{j}_1))(\mathbf{x})$ and $\zeta(g(-, \mathbf{j}_2))(\mathbf{x})$.
2. For each fixed $\mathbf{x} \in P^k$, compute the sum over all $\mathbf{j}_1 + \mathbf{j}_2 = \mathbf{i}$ using the fast convolution algorithm of Corollary 1 in $\mathcal{O}(Q \log(Q))$ arithmetic operations. This takes $\mathcal{O}(|P|^k Q \log(Q))$ arithmetic operations in total. For each $\mathbf{x} \in P^k$ and $\mathbf{i} \in N$, we now have $\zeta(h(-, \mathbf{i}))(\mathbf{x})$.
3. Finally, for each fixed \mathbf{i} , take the Möbius transform of $\zeta(h(-, \mathbf{i}))(\mathbf{x})$ in $\mathcal{O}(|P|^k k)$ time using Lemma 5. Like the first step, this takes $\mathcal{O}(|P|^k Q k)$ arithmetic operations in total. As a result, we have the required values $h(\mathbf{x}, \mathbf{i})$.

The running time follows by summing the times required for each of the three steps. □

3.3 Modular Arithmetic

As discussed in the introduction to this section, we embed the integers $\{0, 1, \dots, M\}$ in the larger field \mathbb{F}_p , for a prime $p > M$. However, we need to choose p 'appropriately' such that the resulting field \mathbb{F}_p has the required root(s) of unity. Below we give a short description of how this can be done.

Let r_1, r_2, \dots, r_k be distinct integers. We look for a prime number p such that \mathbb{F}_p contains, for all i , an r_i -th root of unity. To find the prime p , we consider candidates $m_j = 1 + jR$, where $R = \prod_{i=1}^k r_i$, for j large enough such that $m_j > M$. By the prime number theorem for arithmetic progressions, the sequence $(m_j)_{j=1}^\infty$ contains $\mathcal{O}(\frac{1}{\phi(R)} \frac{x}{\ln(x)})$ prime numbers less than x , where ϕ is Euler's totient function. Since prime testing can be done in polynomial time, we can look for the first candidate $m_j > M$ that is prime and choose p as such.

By Euler's theorem, for any $x \in \mathbb{F}_p$, with p chosen as in the previous paragraph: $1 = x^{\phi(p)} = x^{p-1} = x^{jR}$. As such, for any $x \in \mathbb{F}_p$, x^l with $l = \frac{jR}{r_i}$ is an r_i -th root of unity if $(x^l)^i \neq 1$ for all $i < r_i$. Finding an appropriate x is not difficult for small r_i as an $\frac{1}{r_i}$ -th fraction of all elements $x \in \mathbb{F}_p$ results in x^l being an r_i -th root of unity. To see this, consider a generator g of the multiplicative subgroup of \mathbb{F}_p . The sequence g^1, g^2, \dots, g^{p-1} equals all elements in $\mathbb{F}_p \setminus \{0\}$. Putting this sequence to the power l gives $g^l, g^{2l}, \dots, g^{(p-1)l}$ which, by choice of l , equals $\omega_{r_i}^1, \omega_{r_i}^2, \dots, \omega_{r_i}^{(p-1)}$, where ω_{r_i} is an r_i -th root of unity in \mathbb{F}_p . Clearly, this forms l times the sequence $\omega_{r_i}^1, \omega_{r_i}^2, \dots, \omega_{r_i}^{r_i}$, as $\omega_{r_i}^{r_i} = 1$.

4 Fast Join Operations

Having introduced the basics of dynamic programming on tree decompositions for $[\sigma, \rho]$ -domination problems in Sect. 2, and the basics of the fast Fourier and fast Möbius transforms in Sect. 3, we are now ready to apply the fast transforms to obtain fast join operations. We will survey some known techniques based on both transforms.

Recall that, to compute the memoisation table for a join node i of a nice tree decomposition, we are given two memoisation tables A_l and A_r corresponding to the (left and right) child nodes of i . These tables store a number for each state colouring \mathbf{c} with labels from C (as defined in Sect. 2.2): this number indicates the existence (0 or 1) of a, the size of a, and/or the number of (minimum/maximum size) partial solution(s) on G_l and G_r for the partial-solution equivalence class corresponding to \mathbf{c} . Here, a partial-solution equivalence class is uniquely identified by \mathbf{c} in the following way: a label for a vertex v in \mathbf{c} defines whether v is in the solution set or not, and it defines how many neighbours v has in the solution set (see Sect. 2.2). Our goal is to compute A_i .

As stated in Sect. 2.2, a trivial algorithm would loop over all combinations of state colourings $\mathbf{c}_l, \mathbf{c}_r$ of X_i that agree on which vertices are in the solution set D . Then, the algorithm considers two corresponding partial solutions D_l on G_l and D_r on G_r , and it infers the state colouring \mathbf{c}_i of the partial solution $D_l \cup D_r$ on G_i . Over all constructed partial solutions on G_i , it stores in A_i the minimum size of a solution for each equivalence class \mathbf{c}_i . It is not hard to show that one does not need to consider the partial solutions representing an equivalence class: given the state colourings \mathbf{c}_l and \mathbf{c}_r , the state colouring of $D_l \cup D_r$ can be inferred directly. This is done as follows. Since G_l, G_r and $G_i = G_l \oplus G_r$ do not contain edges between vertices in X_i , for any vertex $v \in X_i$, the number of neighbours in D_l and D_r add up to the resulting number in $D_r \cup D_l$. As such, for any vertex v , if v has label $|l|_\sigma$ in \mathbf{c}_l and $|l'|_\sigma$ in \mathbf{c}_r , then any combined partial solution has label $|l + l'|_\sigma$ for v in the state colouring that identifies the equivalence class (or label $|\geq \ell|_\sigma$ when $l + l' \geq \ell$ and the $|\geq \ell|_\sigma$ label is in C_σ). The same holds for labels $|l|_\rho$ from C_ρ .

We find it insightful to make tables, which we call ‘join tables’, that visualise the resulting label of a vertex in \mathbf{c}_i given its labels in \mathbf{c}_l and \mathbf{c}_r : see Fig. 1. In these tables, the patterns emerge that our fast join operations must fulfil. Here, one can see the running time that a trivial algorithm uses to perform the join: every non-empty cell represents a combination from A_l and A_r that can be made on each vertex coordinate (each vertex in X_i). As a result, this trivial algorithm performs the join in $\mathcal{O}^*(x^k)$ time, where $k = |X_i|$ and x is the number of non-empty cells in the join table.

	$ \geq 0 _\sigma$	$ 0 _\rho$	$ \geq 1 _\rho$		$ 0 _\sigma$	$ 0 _\rho$	$ 1 _\rho$		$ 0 _\sigma$	$ 1 _\sigma$	$ 2 _\sigma$	$ 3 _\sigma$	$ \geq 0 _\rho$		$ 0 _\sigma$	$ \geq 1 _\sigma$	$ 0 _\rho$	$ \geq 1 _\rho$	
$ \geq 0 _\sigma$	$ \geq 0 _\sigma$	$ \geq 0 _\sigma$	$ \geq 1 _\rho$	$ \geq 0 _\sigma$	$ 0 _\sigma$	$ 0 _\sigma$	$ 0 _\rho$	$ 1 _\rho$	$ 0 _\sigma$	$ 1 _\sigma$	$ 2 _\sigma$	$ 3 _\sigma$	$ \geq 0 _\rho$	$ \geq 0 _\rho$	$ 0 _\sigma$	$ \geq 1 _\sigma$	$ \geq 1 _\sigma$	$ 0 _\rho$	$ \geq 1 _\rho$
$ 0 _\rho$	$ \geq 0 _\sigma$	$ 0 _\rho$	$ \geq 1 _\rho$	$ 0 _\rho$	$ 0 _\sigma$	$ 0 _\sigma$	$ 0 _\rho$	$ 1 _\rho$	$ 1 _\sigma$	$ 2 _\sigma$	$ 3 _\sigma$	$ \geq 0 _\rho$	$ \geq 0 _\rho$	$ \geq 1 _\rho$	$ \geq 1 _\sigma$	$ \geq 1 _\sigma$	$ \geq 1 _\sigma$	$ 0 _\rho$	$ \geq 1 _\rho$
$ \geq 1 _\rho$	$ \geq 0 _\sigma$	$ \geq 1 _\rho$	$ \geq 1 _\rho$	$ \geq 1 _\rho$	$ 1 _\sigma$	$ 1 _\sigma$	$ 1 _\rho$	$ 1 _\rho$	$ 1 _\sigma$	$ 3 _\sigma$	$ 3 _\sigma$	$ \geq 0 _\rho$	$ \geq 0 _\rho$	$ \geq 1 _\rho$	$ \geq 1 _\sigma$	$ \geq 1 _\sigma$	$ \geq 1 _\sigma$	$ \geq 1 _\rho$	$ \geq 1 _\rho$

Fig. 1. Join tables corresponding, from left to right, to DOMINATING SET, STRONG STABLE SET, 3-REGULAR SUBGRAPH, and TOTAL DOMINATING SET.

4.1 Möbius Transforms for Dominating Set and Independent Dominating Set

For the DOMINATING SET problem, consider the first (leftmost) join table in Fig. 1. Notice that this is also the join table for the INDEPENDENT DOMINATING SET problem.

We first consider the (non-optimisation) counting variant of DOMINATING SET or INDEPENDENT DOMINATING SET: the table entries $A_i(\mathbf{c})$ represent the number of partial solutions to (independent) dominating set in the equivalence class represented by \mathbf{c} (partial solutions of any size).

Lemma 7 (based on [26]). *The join for the counting variant of DOMINATING SET can be computed in $\mathcal{O}(3^k k)$ arithmetic operations.*

Proof. We first loop over all 2^k subsets $X_\sigma \subseteq X_i$, and for each subset X_σ , we fix the labels of vertices in X_σ in \mathbf{c}_l , \mathbf{c}_r and \mathbf{c}_i to $|\geq 0|_\sigma$. We then consider the subproblem that remains using only the labels $|0|_\rho$ and $|\geq 1|_\rho$. Let $X' = X_i \setminus X_\sigma$ be the vertices without fixed label, let $k' = |X'|$, and let A'_i, A'_l, A'_r be the memoisation tables A_i, A_l , and A_r after fixing the vertices with label $|\geq 0|_\sigma$, i.e., A'_i, A'_l, A'_r are indexed by state colourings \mathbf{c}'_l and \mathbf{c}'_r on X'_i using only $|0|_\rho$ and $|\geq 1|_\rho$ -labels.

To compute the join, we now essentially want to take a coordinate-wise maximum of the state colourings \mathbf{c}'_l and \mathbf{c}'_r (identifying $|0|_\rho$ with 0 and $|\geq 1|_\rho$ with 1) to obtain the resulting state colouring \mathbf{c}'_i on X'_i . That is, to compute $A'_i(\mathbf{c}'_i)$, we want to efficiently evaluate the following formula:

$$A'_i(\mathbf{c}'_i) = \sum_{\mathbf{c}'_l \vee \mathbf{c}'_r = \mathbf{c}'_i} A'_l(\mathbf{c}'_l) A'_r(\mathbf{c}'_r) \tag{12}$$

where \vee is the above discussed coordinate-wise maximum (identifying C_ρ with $\mathbb{N}_{<2}^k$). Observe that this corresponds exactly to the covering product, generalised in Lemma 6 with $P = \mathbb{N}_{<2}^k$ and $f = A'_l, g = A'_r$. Consequently, this join can be computed in $\mathcal{O}(2^{k'} k')$ arithmetic operations by Corollary 2. Summing up the running time over all 2^k subsets of fixed labels, we obtain a running time of:

$$\begin{aligned} \mathcal{O} \left(\sum_{X' \subseteq X_i} 2^{|X'|} |X'| \right) &= \mathcal{O} \left(\sum_{k'=0}^k \binom{k}{k'} 2^{k' k'} \right) = \mathcal{O} \left(k \sum_{k'=0}^k \binom{k}{k'} 2^{k'} 1^{k-k'} \right) \\ &= \mathcal{O}(k(2+1)^k) = \mathcal{O}(3^k k) \end{aligned}$$

where we group the subsets $X' = X_i \setminus X_\sigma$ of the same size and then use the binomial theorem. \square

Corollary 3. *Given a graph G with a tree decomposition T of G of width t , the number of (independent) dominating sets can be computed in $\mathcal{O}(3^t t n)$ arithmetic operations on $\mathcal{O}(n)$ -bit numbers.*

Proof. Plug Lemma 7 into Lemma 1. We can use $\mathcal{O}(n)$ -bit numbers as the result is at most 2^n . \square

The above construction does not work directly for the minimisation version of (INDEPENDENT) DOMINATING SET, as then we are no longer counting combinations of partial solutions: we want to take the minimum over the sum of partial solution sizes. I.e., instead of Eq. 12, we need:

$$A'_i(\mathbf{c}'_i) = \min_{\mathbf{c}'_l \vee \mathbf{c}'_r = \mathbf{c}'_i} A'_l(\mathbf{c}'_l) + A'_r(\mathbf{c}'_r) \tag{13}$$

To obtain a similar result for the minimisation versions, we will embed this into a counting structure.

Lemma 8 (based on [26]). *The join for the minimisation variant of DOMINATING SET can be computed in $\mathcal{O}(3^k n(k + \log(n)))$ arithmetic operations.*

Proof. The proof is identical to the proof of Lemma 7, except that we need a different fast evaluation algorithm: one that corresponds to Eq. 13. To this end, we expand the memoisation tables A'_i , A'_l and A'_r by having the solution size as part of the index of the table. That is, we let:

$$A'_l(\mathbf{c}'_l, \kappa_l) = \begin{cases} 1 & \text{if } A'_l(\mathbf{c}'_l) = \kappa_l \\ 0 & \text{otherwise} \end{cases} \tag{14}$$

and similarly for A'_r . Let $A'_i(\mathbf{c}'_i, \kappa_i)$ be defined as follows, which can be computed using Theorem 1:

$$A'_i(\mathbf{c}'_i, \kappa_i) = \sum_{\mathbf{c}'_l \vee \mathbf{c}'_r = \mathbf{c}'_i} \sum_{\kappa_l + \kappa_r = \kappa_i} A'_l(\mathbf{c}'_l, \kappa_l) A'_r(\mathbf{c}'_r, \kappa_r) \tag{15}$$

It is easy to see that $A'_i(\mathbf{c}'_i, \kappa_i) > 0$ if and only if there exists $\mathbf{c}'_l, \kappa_l, \mathbf{c}'_r, \kappa_r$ such that $\mathbf{c}'_l \vee \mathbf{c}'_r = \mathbf{c}'_i$ and $A'_l(\mathbf{c}'_l) = \kappa_l$ $A'_r(\mathbf{c}'_r) = \kappa_r$. This allows us to obtain the result required by Eq. 13 by setting $A'_i(\mathbf{c}'_i)$ equal to the minimum value of κ_i for which $A'_i(\mathbf{c}'_i, \kappa_i) > 0$.

Observe that κ_i can range between 0 and n . Therefore, when we apply Theorem 1 with $N = \mathbb{N}_{<n+1}$, we can perform the join in $\mathcal{O}(3^k n(k + \log(n)))$ arithmetic operations. \square

Corollary 4. *Given a graph G with a tree decomposition T of G of width t , INDEPENDENT DOMINATING SET can be solved in $\mathcal{O}(3^t n^2(t + \log(n)))$ arithmetic operations on $\mathcal{O}(t + \log(n))$ -bit numbers.*

Proof. Plug Lemma 8 into Lemma 1. We need $\log(n)$ bit numbers for the sizes of partial solutions, while the sums in Eq. 15 can require up to $\mathcal{O}(k)$ -bit numbers. \square

We can gain linear dependence on n for DOMINATING SET by using a replacement property (see also [26]) that holds both for DOMINATING SET and TOTAL DOMINATING SET.

Definition 7 (replacement property for partial solutions). *An optimisation problem \mathcal{P} has the replacement property if the difference in size between the smallest and the largest partial solution for non-dominated equivalence classes is at most k .*

Mostly this holds for a problem \mathcal{P} when, if given two partial solutions, one can add or subtract all vertices in X_i from either one to obtain a solution that is at least as good as the other. This is the case for DOMINATING SET as adding all vertices in X_i to a partial solution D dominates all vertices any partial solution can dominate, thus being less restrictive than any other partial solution.

Corollary 5. *Given a graph G with a tree decomposition T of G of width t , DOMINATING SET can be solved in $\mathcal{O}(3^t t^2 n)$ arithmetic operations on $\mathcal{O}(t + \log(n))$ -bit numbers.*

Proof. We further modify the algorithm used in Lemma 8 (based on Lemma 7). Let ξ_l and ξ_r be the minimum values from $A'_l(c'_l)$ and $A'_r(c'_r)$. If we restrict the ranges of $\kappa_i, \kappa_l, \kappa_r$ in Eq. 15 to $[0, 1, \dots, k]$, then Theorem 1 allows us to evaluate the equation in $\mathcal{O}(2^{k'} k'^2)$ arithmetic operations. We can do so by subtracting ξ_l from all values in A'_l and ξ_r from all values in A'_r before adding the size-parameter to the index of the table. After the join, we can add $\xi_l + \xi_r$ to the results in A'_i . It is not hard to see that this does not influence the result of the algorithm, but now allows an $\mathcal{O}(3^k k^2)$ -time join operation. The result then follows from plugging this result into Lemma 1. \square

4.2 Count and Filter: Strong Stable Set, Perfect Code and Perfect Dominating Set

To obtain fast joins for the next set of problems, we now introduce a filtering trick based on counting. The algorithm we use here, is in essence the fast subset convolution algorithm by Björklund et al. [2]. Our different presentation is chosen so that we can use the same trick in the sections to follow.

First notice that the three problems mentioned in this section’s title have essentially the same join table (the second table in Fig. 1): even though PERFECT DOMINATING SET uses the $|\geq 0|_\sigma$ -label while the others use the $|0|_\sigma$ -label, the structure of the join tables is identical. Compared to the join table for DOMINATING SET, the difference in terms of Eq. 12 is that we now want to compute:

$$A'_i(c'_i) = \sum_{c'_l + c'_r = c'_i} A'_l(c'_l) A'_r(c'_r) \tag{16}$$

That is, where in Eq. 12 we sum over three combinations to obtain a $|\geq 1|_\rho$ -label in \mathbf{c}'_i , we may now only sum over two combinations ($|0|_\rho + |1|_\rho = |1|_\rho$, $|1|_\rho + |0|_\rho = |1|_\rho$).

Lemma 9. *The join for the maximisation variant of STRONG STABLE SET can be computed in $\mathcal{O}(3^k nk(k + \log(n)))$ arithmetic operations.*

Proof. We use the same construction as in Lemma 7 enumerating all subsets $X_\sigma \subseteq X_i$, $X' = X_i \setminus X_\sigma$, $k' = |X'|$; and for each subset X_σ , we fix the labels of the vertices in X_σ in \mathbf{c}_l , \mathbf{c}_r and \mathbf{c}_i to $|\geq 0|_\sigma$. For the remaining subproblem, let A'_i , A'_l , A'_r be the memoisation tables A_i , A_l , and A_r after fixing the vertices with label $|\geq 0|_\sigma$, i.e., they are indexed by state colourings \mathbf{c}'_l and \mathbf{c}'_r on X'_i using only $|0|_\rho$ and $|1|_\rho$ -labels. Next, we add the solution size to the index of these tables as in the proof of Lemma 8: $A'_i(\mathbf{c}'_l, \kappa_l) = 1$ if and only if $A'_i(\mathbf{c}'_l) = \kappa_l$. Observe that now the join can be computed by letting $A'_i(\mathbf{c}'_i)$ be the minimum value κ_i for which $A'_i(\mathbf{c}'_i, \kappa_i) > 0$, where:

$$A'_i(\mathbf{c}'_i, \kappa_i) = \sum_{\mathbf{c}'_l + \mathbf{c}'_r = \mathbf{c}'_i} \sum_{\kappa_l + \kappa_r = \kappa_i} A'_i(\mathbf{c}'_l, \kappa_l) A'_r(\mathbf{c}'_r, \kappa_r) \quad (17)$$

To compute the result of Eq. 17 efficiently, we add yet another parameter to the index of the tables. This parameter counts the number of $|1|_\rho$ -labels in the state colouring \mathbf{c} . In other words:

$$A'_i(\mathbf{c}'_l, \kappa_l, \iota_l) = \begin{cases} A'_i(\mathbf{c}'_l, \kappa_l) & \text{if } \#_{|1|_\rho}(\mathbf{c}_l) = \iota_l \\ 0 & \text{otherwise} \end{cases} \quad (18)$$

where $\#_{|1|_\rho}(\mathbf{c}_l) = \iota_l$ is our notation for stating that \mathbf{c}_l contains exactly ι_l $|1|_\rho$ -labels. We claim that $A'_i(\mathbf{c}'_i, \kappa_i)$ as defined in Eq. 17 equals $A'_i(\mathbf{c}'_i, \kappa_i, \#_{|1|_\rho}(\mathbf{c}'_i))$, where $A'_i(\mathbf{c}'_i, \kappa_i, \iota_i)$ is defined as:

$$A'_i(\mathbf{c}'_i, \kappa_i, \iota_i) = \sum_{\mathbf{c}'_l \vee \mathbf{c}'_r = \mathbf{c}'_i} \sum_{\kappa_l + \kappa_r = \kappa_i} \sum_{\iota_l + \iota_r = \iota_i} A'_l(\mathbf{c}'_l, \kappa_l, \iota_l) A'_r(\mathbf{c}'_r, \kappa_r, \iota_r) \quad (19)$$

where $\mathbf{c}'_l \vee \mathbf{c}'_r = \mathbf{c}'_i$ is again defined coordinate-wise and by identifying C_ρ with $\mathbb{N}_{<2}$.

Notice that ι_l and ι_r track the number of $|1|_\rho$ -labels used. Therefore, the total number of $|1|_\rho$ -labels in a pair $(\mathbf{c}'_l, \mathbf{c}'_r)$ used as $\mathbf{c}'_l \vee \mathbf{c}'_r = \mathbf{c}'_i$ in a summand of the sum for $A'_i(\mathbf{c}'_i, \kappa_i, \iota_i)$ is exactly ι_i . Since we need at least one $|1|_\rho$ -label in \mathbf{c}'_l or \mathbf{c}'_r to realise each $|1|_\rho$ -label in \mathbf{c}_i , we know that $A'_i(\mathbf{c}'_i, \kappa_i, \#_{|1|_\rho}(\mathbf{c}'_i))$ uses $\#_{|1|_\rho}(\mathbf{c}'_i)$ $|1|_\rho$ -labels in total, and hence equals $A'_i(\mathbf{c}'_i, \kappa_i)$ from Eq. 17.

By Theorem 1, Eq. 19 can be evaluated in $\mathcal{O}(2^{k'} nk'(k' + \log(n)))$ arithmetic operations. Summing the running time over all 2^k subsets $X_\sigma \subseteq X_i$ of vertices for which we fixed the label, this leads to the claimed running time in exactly the same way as in the proof of Lemma 7. \square

Corollary 6. *Given a graph G with a tree decomposition T of G of width t , the optimisation variants of STRONG STABLE SET, PERFECT CODE and PERFECT DOMINATING SET can be solved in $\mathcal{O}(3^t n^2 t(t + \log(n)))$ arithmetic operations on $\mathcal{O}(t + \log(n))$ -bit numbers.*

Proof. Plug Lemma 9 into Lemma 1. We need $\log(n)$ -bit numbers for the sizes of partial solutions, while the sums in Eq. 19 can require up to $\mathcal{O}(k)$ -bit numbers. \square

4.3 Fourier Transforms for Induced Bounded Degree or p -Regular Subgraph

The results of the previous section can also be obtained using counting and filtering on top of Fourier transforms instead of on top of Möbius transforms. The resulting construction also results in fast joins for INDUCED BOUNDED DEGREE SUBGRAPH and p -REGULAR SUBGRAPH: these we present in this section. The join table for 3-REGULAR SUBGRAPH is given in Fig. 1: notice that it is identical to the join table for INDUCED BOUNDED DEGREE SUBGRAPH with degree bound three.

We should note that the results in this section can also be obtained using Cygan and Pilipczuk’s method [17] where they encode solutions as a polynomial and use FFT-based fast polynomial multiplication. The approach below is allows for easier combination with Möbius transforms (as we will see in Sect. 5), and saves a factor of t in the polynomial factors of the running time⁵.

We start in the same setting as before: we fix the vertices with label $|\geq 0|_\rho$ by looping over all subsets $X_\rho \subseteq X_i$ (now we fix states from C_ρ , previously from C_σ) and let $X' = X_i \setminus X_\rho$ and $k' = |X'|$. Let A'_i, A'_l and A'_r be the memoisation tables after fixing the vertices with the $|\geq 0|_\rho$ -label indexed by state colourings $\mathbf{c}'_i, \mathbf{c}'_l, \mathbf{c}'_r$ using labels from $C_\sigma = \{|0|_\sigma, |1|_\sigma, \dots, |\ell - 1|_\sigma, |\ell|_\sigma\}$.

Define the projection function π on labels in C_σ as to be $\pi(|l|_\sigma) = l$. Also, define addition on state colourings as follows: $\mathbf{c}'_l + \mathbf{c}'_r = \mathbf{c}'_i$ if and only if, for all j , $\pi((\mathbf{c}'_l)_j) + \pi((\mathbf{c}'_r)_j) = \pi((\mathbf{c}'_i)_j)$. Computing $A'_i(\mathbf{c}'_i)$ for the counting variant of the problem, or computing $A'_i(\mathbf{c}'_i, \kappa_i)$ for the optimisation variant of the problem using the solution size κ as part of the index, now comes down to evaluating:

$$A'_i(\mathbf{c}'_i) = \sum_{\mathbf{c}'_l + \mathbf{c}'_r = \mathbf{c}'_i} A'_l(\mathbf{c}'_l)A'_r(\mathbf{c}'_r) \quad A'_i(\mathbf{c}'_i, \kappa_i) = \sum_{\mathbf{c}'_l + \mathbf{c}'_r = \mathbf{c}'_i} \sum_{\kappa_l + \kappa_r = \kappa_i} A'_l(\mathbf{c}'_l, \kappa_l) A'_r(\mathbf{c}'_r, \kappa_r) \tag{20}$$

Observe that this looks very similar to the statement of Lemma 3. However, to obtain a non-cyclic convolution for $\mathbf{c}'_l + \mathbf{c}'_r = \mathbf{c}'_i$, a direct application of Lemma 3 would require $\mathcal{O}^*((\ell + 1)^k 2^k)$ arithmetic operations to evaluate either version of Eq. 20. We will use cyclic convolution with counting and filtering to obtain the result in $\mathcal{O}^*((\ell + 1)^k)$ arithmetic operations, resulting in an $\mathcal{O}^*((\ell + 2)^k) = \mathcal{O}^*(s^k)$ time join operation.

Lemma 10. *The join for the counting variant of INDUCED BOUNDED DEGREE SUBGRAPH can be computed in $\mathcal{O}((\ell + 2)^k k^2 \ell \log(\ell + 1))$ arithmetic operations.*

⁵ The approach in [17] uses a factor k^3 (compared to our k^2) by performing k^2 FFT-based polynomial multiplications that each cost $\mathcal{O}((\ell + 1)^k \log((\ell + 1)^k)) = \mathcal{O}((\ell + 1)^k k \log(\ell + 1))$ arithmetic operations.

Proof. For a state colouring $\mathbf{c}'_i \in C_\sigma^{k'}$, define $\Sigma(\mathbf{c}'_i) = \sum_{j=1}^{k'} \pi((\mathbf{c}'_i)_j)$. Now, similar to the proof of Lemma 9 where we added the number of 1-labels to the index of the table, we now add the sum of the labels $\Sigma(\mathbf{c}'_i)$ to the index of the table. That is, for both A'_l and A'_r , we set:

$$A'_i(\mathbf{c}'_i, \iota_i) = \begin{cases} A'_i(\mathbf{c}'_i) & \text{if } \Sigma(\mathbf{c}'_i) = \iota_i \\ 0 & \text{otherwise} \end{cases} \tag{21}$$

To compute the join efficiently, we use that $A'_i(\mathbf{c}'_i)$ as defined in Eq. 20 (left equation) equals $A'_i(\mathbf{c}'_i, \Sigma(\mathbf{c}'_i))$, where $A'_i(\mathbf{c}'_i, \iota_i)$ is the result of the following summation:

$$A'_i(\mathbf{c}'_i, \iota_i) = \sum_{\mathbf{c}'_l + \mathbf{c}'_r \equiv \mathbf{c}'_i} \sum_{\iota_l + \iota_r = \iota_i} A'_l(\mathbf{c}'_l, \iota_l) A'_r(\mathbf{c}'_r, \iota_r) \tag{22}$$

where $\mathbf{c}'_l + \mathbf{c}'_r \equiv \mathbf{c}'_i$ is now defined as, for all j , $\pi((\mathbf{c}'_l)_j) + \pi((\mathbf{c}'_r)_j) \equiv \pi((\mathbf{c}'_i)_j)$ modulo $\ell + 1$, and $\iota_l + \iota_r = \iota_i$ is the standard addition without modulus. To see that this is correct, notice that the parameters ι_l and ι_r track the sum of the labels in \mathbf{c}'_l and \mathbf{c}'_r , and $\iota_l + \iota_r = \Sigma(\mathbf{c}'_i)$ implies that each of the individual components of $\mathbf{c}'_l + \mathbf{c}'_r$ cannot cycle as that would result in $\iota_l + \iota_r > \Sigma(\mathbf{c}'_i)$.

By Lemma 3, Eq. 22 can be evaluated in $\mathcal{O}((\ell + 1)^{k'} k'^2 \ell \log(\ell + 1))$ arithmetic operations as the ι parameters range from 0 to $\ell k'$. The claimed running time follows by summing this running time over all $X_\rho \subseteq X_i$ for which we fixed the label (similar to in the proof of Lemma 7). \square

Corollary 7. *Given a graph G with a tree decomposition T of G of width t , the counting variant of INDUCED BOUNDED DEGREE SUBGRAPH can be solved in $\mathcal{O}((\ell + 2)^{t+1} t^2 n \ell \log(\ell + 1))$ arithmetic operations on $\mathcal{O}(n)$ -bit numbers.*

Proof. Plug Lemma 10 into Lemma 1. We can use $\mathcal{O}(n)$ -bit numbers as the result is at most 2^n . Since ℓ is a variable (not a constant), we need $t + 1$ in the exponent as $t \geq k - 1$ (see Definition 1). \square

Adapting the above lemma and corollary to the problem’s optimisation variant is a simple exercise.

4.4 Möbius Transforms with a Different Partial Order for Total Dominating Set

In the previous sections, we fixed the vertices with a $|\geq 0|_\sigma$ -label (or $|\geq 0|_\rho$ -label) and used a fast transform only on vertices with a label from C_ρ (or C_σ). Here, we give an example of a fast transform that deals with vertices with different labels from both C_σ and C_ρ simultaneously.

Consider the label set $C = \{|0|_\sigma, |\geq 1|_\sigma, |0|_\rho, |\geq 1|_\rho\}$ associated to the TOTAL DOMINATING SET problem. On this label set, we impose the following partial order: all labels are incomparable except that, we impose $|0|_\sigma \leq |\geq 1|_\sigma$ and $|0|_\rho \leq |\geq 1|_\rho$. Notice that, using this partial order, C^k does not form a lattice, e.g., in C^3 the join $[|0|_\sigma, |0|_\rho, |\geq 1|_\sigma] \vee [|\geq 1|_\rho, |\geq 1|_\rho, |0|_\sigma]$ is undefined due

to the first coordinate. However, it is not hard to see that, for every $\mathbf{x} \in C^k$, $\{\mathbf{y} \in P^k \mid \mathbf{y} \leq \mathbf{x}\}$ forms a join-semilattice: either x_i equals $|0|_\sigma$ or $|0|_\rho$ and then y_i must be equal to x_i , or x_i equals $|\geq 1|_\sigma$ or $|\geq 1|_\rho$ and then y_i has two choices which are comparable, and thus the join is defined (actually this is a lattice as the meet is also defined).

Proposition 6. *Let C and its partial order be defined as above. There are algorithms for the zeta and Möbius transform for functions $f : C^k \rightarrow \mathbb{F}_p$ that require $\mathcal{O}(4^k k)$ arithmetic operations.*

Proof. If $k = 1$, the zeta and Möbius transforms are:

$$\begin{aligned} \zeta(f)(|0|_\sigma) &= f(|0|_\sigma) & \mu(f)(|0|_\sigma) &= f(|0|_\sigma) \\ \zeta(f)(|\geq 1|_\sigma) &= f(|0|_\sigma) + f(|\geq 1|_\sigma) & \mu(|\geq 1|_\sigma) &= f(|\geq 1|_\sigma) - f(|0|_\sigma) \\ \zeta(f)(|0|_\rho) &= f(|0|_\rho) & \mu(f)(|0|_\rho) &= f(|0|_\rho) \\ \zeta(f)(|\geq 1|_\rho) &= f(|0|_\rho) + f(|\geq 1|_\rho) & \mu(|\geq 1|_\rho) &= f(|\geq 1|_\rho) - f(|0|_\rho) \end{aligned}$$

The result follows from Lemma 5 as these require a constant amount of arithmetic operations. □

Lemma 11 (based on [26]). *The join for the minimisation variant of TOTAL DOMINATING SET can be computed in $\mathcal{O}(4^k k^2)$ arithmetic operations.*

Proof. Given A_l and A_r , we can compute A_i by evaluating the following equation that is equivalent to Eq. 13 (notice that this corresponds exactly to the rightmost join table in Fig. 1):

$$A_i(\mathbf{c}_i) = \min_{\mathbf{c}_l \vee \mathbf{c}_r = \mathbf{c}_i} A_l(\mathbf{c}_l) + A_r(\mathbf{c}_r) \tag{23}$$

To obtain our result, we expand the memoisation tables A_i , A_l and A_r by having the solution size as part of the index of the table. That is, we let:

$$A_l(\mathbf{c}_l, \kappa_l) = \begin{cases} 1 & \text{if } A_l(\mathbf{c}_l) = \kappa_l \\ 0 & \text{otherwise} \end{cases} \tag{24}$$

and similarly for A_r . Then, we compute $A_i(\mathbf{c}_i, \kappa_i)$ as defined below, computed using Theorem 1:

$$A_i(\mathbf{c}_i, \kappa_i) = \sum_{\mathbf{c}_l \vee \mathbf{c}_r = \mathbf{c}_i} \sum_{\kappa_l + \kappa_r = \kappa_i} A_l(\mathbf{c}_l, \kappa_l) A_r(\mathbf{c}_r, \kappa_r) \tag{25}$$

Since $A_i(\mathbf{c}_i, \kappa_i) > 0$ if and only if there exists $\mathbf{c}_l, \kappa_l, \mathbf{c}_r, \kappa_r$ such that $\mathbf{c}_l \vee \mathbf{c}_r = \mathbf{c}_i$ and $A_l(\mathbf{c}_l) = \kappa_l, A_r(\mathbf{c}_r) = \kappa_r$, this allows us to obtain the result required by Eq. 23 by setting $A_i(\mathbf{c}_i)$ equal to the minimum value of κ_i for which $A_i(\mathbf{c}_i, \kappa_i) > 0$.

A direct application of Theorem 1 would allow us to evaluate Eq. 25 in $\mathcal{O}(4^k n^2(k + \log(n)))$ arithmetic operations. However, if we restrict the ranges of $\kappa_i, \kappa_l, \kappa_k$ to $[0, 1, \dots, k]$, then Theorem 1 allows us to evaluate Eq. 25 in

$\mathcal{O}(4^k k^2)$ arithmetic operations. We can do so because, just as in Corollary 5, TOTAL DOMINATING SET satisfies the replacement property (Definition 7): take the minimum value from A_l and A_r and subtract this from all values in A_l and A_r before adding the size-parameter to the index of the table. After the join, we can again add the sum of both minima to the results in A_i . The result now follows. \square

Corollary 8. *Given a graph G with a tree decomposition T of G of width t , TOTAL DOMINATING SET can be solved in $\mathcal{O}(4^t t^2 n)$ arithmetic operations on $\mathcal{O}(t + \log(n))$ -bit numbers.*

Proof. Plug Lemma 11 into Lemma 1 and observe that we need $\log(n)$ -bit numbers for the sizes of partial solutions, while the sums in Eq. 25 can require up to $\mathcal{O}(t)$ -bit numbers. \square

5 Bringing It Together: Faster Algorithms for $[\sigma, \rho]$ -Domination

In the previous section, we have surveyed a number of approaches to realise fast joins operations and have given some concrete examples. In this section, we will integrate several of these approaches into a new result obtaining the currently fastest algorithm for $[\sigma, \rho]$ -domination in its general form. The worst-case running time of the new algorithm is of the form $\mathcal{O}(s^t (nts)^{\mathcal{O}(1)})$, where the previously fastest algorithm by Van Rooij et al. [25, 26], has s as an exponent in the polynomial part of the running time, i.e., $\mathcal{O}(s^{t+2} (tn)^s (nts)^{\mathcal{O}(1)})$. Here, we write $s = |C|$, where C is the set of labels used in the dynamic programming algorithm for a specific $[\sigma, \rho]$ -domination problem (as in Sect. 2.2). To limit the (already heavy) notational burden, we give our result for a $[\sigma, \rho]$ -domination problem with *cofinite* σ and *finite* ρ . It is not hard to modify the proofs for the other cases.

For state colourings $\mathbf{c}_l, \mathbf{c}_r$ of X_i with labels from C , we define the operator $\mathbf{c}_l \oplus \mathbf{c}_r = \mathbf{c}_i$ as the coordinate-wise addition operator that keeps addition within both parts of the label set. This operator is only defined if \mathbf{c}_l and \mathbf{c}_r agree on which vertices are labelled with labels from C_σ and from C_ρ and if the result again is in C_σ or C_ρ . More formally, when we write $\mathbf{c}_l \oplus \mathbf{c}_r = \mathbf{c}_i$ and if $(\mathbf{c}_l)_j = |x_l|_\rho, (\mathbf{c}_r)_j = |x_r|_\rho$, then $(\mathbf{c}_i)_j = |x_l + x_r|_\rho$ if $|x_l + x_r|_\rho \in C_\rho$ and otherwise it is undefined. And, if $(\mathbf{c}_l)_j = |x_l|_\sigma, (\mathbf{c}_r)_j = |x_r|_\sigma$, then $(\mathbf{c}_i)_j = |x_l + x_r|_\sigma$ if $|x_l + x_r|_\sigma < \ell_\sigma$ and $(\mathbf{c}_i)_j = |\geq \ell|_\sigma$ otherwise. Observe that this corresponds exactly to the structure of how a join for a $[\sigma, \rho]$ -domination problem with cofinite σ and finite ρ should be performed. Besides the \oplus -operator, we will also use the standard $+$ -operator on state colourings: $\mathbf{c}_l + \mathbf{c}_r = \mathbf{c}_i$. Here, the underlying operation is the standard addition operator within each half of the label set, which is undefined if any $|\geq \ell|_\sigma$ or $|\geq \ell|_\rho$ -label is involved. That is, if $(\mathbf{c}_l)_j = |x_l|_\rho, (\mathbf{c}_r)_j = |x_r|_\rho$, then $(\mathbf{c}_i)_j = |x_l + x_r|_\rho$, which is defined only if $|x_l + x_r|_\rho \in C_\rho$. Addition with the $+$ -operator is similar for the σ -labels.

Let again A_l and A_r be indexed by both the state colouring and the solution size (similar to Eqs. 14 and 24). Performing the join for the minimisation variant of a $[\sigma, \rho]$ -domination problem can now be done by extracting, for each \mathbf{c}_i , the minimum value of κ_i for with the following expression is non-zero:

$$A_i(\mathbf{c}_i, \kappa_i) = \sum_{\mathbf{c}_l \oplus \mathbf{c}_r = \mathbf{c}_i} \sum_{\kappa_l + \kappa_r = \kappa_i} A_l(\mathbf{c}_l, \kappa_l) A_r(\mathbf{c}_r, \kappa_r) \tag{26}$$

To obtain a fast evaluation algorithm for Eq. 26, we use both zeta/Möbius transforms and Fourier transforms. To do so, we impose the following partial order \mathbb{p} on the label set C : all labels are incomparable except that, because σ is cofinite, we impose for all $|\ell|_\sigma \in C_\sigma$: $|\ell|_\sigma \leq |\geq \ell|_\sigma$.

Given a state colouring \mathbf{c}_i of the vertices in X_i with labels from $C = C_\sigma \cup C_\rho$, we write $\mathbf{c}_i = [\mathbf{c}_i^\sigma, \mathbf{c}_i^\rho] = [\mathbf{c}_i^{\geq \ell_\sigma}, \mathbf{c}_i^{< \ell_\sigma}, \mathbf{c}_i^\rho]$ to differentiate between the vertices with label from C_σ and C_ρ , and also to further differentiate between vertices with the label $|\geq \ell|_\sigma$ and vertices with a label from $\{ |0|_\sigma, |1|_\sigma, \dots, |\ell-1|_\sigma \}$. By splitting \mathbf{c}_i in this way we notationally split the different coordinates. This is just notation: we do not reorder anything in the dynamic programming table (e.g., \mathbf{c}_i^ρ can contain the first coordinate of \mathbf{c}_i and the last, while \mathbf{c}_i^σ contains the ones in between).

Using this notation, the zeta transform of a memoisation table A_i indexed by both a state colouring $\mathbf{c}_i = [\mathbf{c}_i^{\geq \ell_\sigma}, \mathbf{c}_i^{< \ell_\sigma}, \mathbf{c}_i^\rho]$ and additional indices \mathbf{x}_i (whose purpose will become clear later) becomes:

$$\zeta(A_i)(\mathbf{c}_i, \mathbf{x}_i) = \sum_{\mathbf{d} \leq \mathbf{c}_i} A_i(\mathbf{d}, \mathbf{x}_i) = \sum_{\mathbf{d}_1 \leq \mathbf{c}_i^{\geq \ell_\sigma}} A_i([\mathbf{d}_1, \mathbf{c}_i^{< \ell_\sigma}, \mathbf{c}_i^\rho], \mathbf{x}_i) \tag{27}$$

Proposition 7. *Given the memoisation table $A_i(\mathbf{c}_i, \mathbf{x}_i)$ indexed by state colourings $\mathbf{c}_i \in C^k$ over the label set C ($|C| = s$) and some additional indices \mathbf{x}_i with domain I , the zeta transform $\zeta(A_i)$ of A_i based on the partial order \mathbb{p} can be computed in $\mathcal{O}(s^k k |I|)$ arithmetic operations. Also, given $\zeta(A_i)$, A_i can be reconstructed in $\mathcal{O}(s^k k |I|)$ arithmetic operations.*

Proof. We will show that for $k = 1$ and hence $\mathbf{c}_i \in C^1$, we have zeta and Möbius transforms on $A_i(\mathbf{c}_i, \mathbf{x}_i)$ requiring $\mathcal{O}(s|I|)$ arithmetic operations. The result then follows from Lemma 5 and the fact that the transforms operate independent of the parameter \mathbf{x}_i .

By definition of the partial order \mathbb{p} , the following formulas compute $\zeta(A_i)$ from A_i and vice versa when $k = 1$:

$$\zeta(A_i)(\mathbf{c}_i, \mathbf{x}) = \begin{cases} A_i(\mathbf{c}_i, \mathbf{x}) & \text{if } \mathbf{c}_i \neq [|\geq \ell|_\sigma] \\ \sum_{z \in C_\sigma} A_i([z], \mathbf{x}) & \text{if } \mathbf{c}_i = [|\geq \ell|_\sigma] \end{cases}$$

$$A_i(\mathbf{c}_i, \mathbf{x}) = \begin{cases} \zeta(A_i)(\mathbf{c}_i, \mathbf{x}) & \text{if } \mathbf{c}_i \neq [|\geq \ell|_\sigma] \\ \zeta(A_i)(\mathbf{c}_i, \mathbf{x}) - \sum_{z \in C_\sigma \setminus \{|\geq \ell|_\sigma\}} \zeta(A_i)([z], \mathbf{x}) & \text{if } \mathbf{c}_i = [|\geq \ell|_\sigma] \end{cases}$$

Each requires $\mathcal{O}(s|I|)$ arithmetic operations as the sums are computed only when $\mathbf{c}_i = [|\geq \ell|_\sigma]$. □

Theorem 2 (bringing it all together). *For an optimisation variant of a $[\sigma, \rho]$ -domination problem with $s = |C|$, the join can be computed in $\mathcal{O}(s^{k+1}kn(k \log(s) + \log(n)))$ arithmetic operations.*

Proof. We want to evaluate Eq. 26 using a fast transform. Consider the what happens to this equation if we apply the zeta transform based on the partial order \mathbb{p} (Eq. 27) to it:

$$\begin{aligned} \zeta(A_i)(\mathbf{c}_i, \kappa_i) &= \sum_{\mathbf{d}_1 \leq \mathbf{c}_i^{\geq \ell_\sigma}} A_i([\mathbf{d}_1, \mathbf{c}_i^{\leq \ell_\sigma}, \mathbf{c}_i^\rho], \kappa_i) \\ &= \sum_{\mathbf{d}_1 \leq \mathbf{c}_i^{\geq \ell_\sigma}} \sum_{\mathbf{d}_l \oplus \mathbf{d}_r = [\mathbf{d}_1, \mathbf{c}_i^{\leq \ell_\sigma}, \mathbf{c}_i^\rho]} \sum_{\kappa_l + \kappa_r = \kappa_i} A_l(\mathbf{d}_l, \kappa_l) A_r(\mathbf{d}_r, \kappa_r) \end{aligned} \quad (28)$$

Continuing from here, we can decompose \mathbf{d}_l and \mathbf{d}_r coordinate-wise in the same way as we have decomposed \mathbf{c}_i as $[\mathbf{c}_i^{\geq \ell_\sigma}, \mathbf{c}_i^{\leq \ell_\sigma}, \mathbf{c}_i^\rho]$ (to be clear: we split the coordinates of \mathbf{d}_l and \mathbf{d}_r based on the labels in \mathbf{c}_i , not based on the actual labels in \mathbf{d}_l and \mathbf{d}_r). Let $\mathbf{d}_l = [\mathbf{d}_{1l}, \mathbf{d}_{2l}, \mathbf{d}_{3l}]$, $\mathbf{d}_r = [\mathbf{d}_{1r}, \mathbf{d}_{2r}, \mathbf{d}_{3r}]$. Now, observe that in $\mathbf{d}_l \oplus \mathbf{d}_r$, any pair \mathbf{d}_{1l} and \mathbf{d}_{1r} on the coordinates of $\mathbf{c}_i^{\geq \ell_\sigma}$ is summed over exactly once because for any pair there is exactly one \mathbf{d}_1 such that $\mathbf{d}_{1l} \oplus \mathbf{d}_{1r} = \mathbf{d}_1$. Also observe that because the other coordinates of \mathbf{d}_l and \mathbf{d}_r correspond to the vertices from the $\mathbf{c}_i^{\leq \ell_\sigma}$ and \mathbf{c}_i^ρ parts of \mathbf{c}_i their \oplus -sum is the standard (non-cyclic) $+$ -addition on labels. As such, we obtain:

$$\zeta(A_i)(\mathbf{c}_i, \kappa_i) = \sum_{\mathbf{d}_{1l}, \mathbf{d}_{1r} \leq \mathbf{c}_i^{\geq \ell_\sigma}} \sum_{\mathbf{d}_{2l} + \mathbf{d}_{2r} = \mathbf{c}_i^{\leq \ell_\sigma}} \sum_{\mathbf{d}_{3l} + \mathbf{d}_{3r} = \mathbf{c}_i^\rho} \sum_{\kappa_l + \kappa_r = \kappa_i} A_l(\mathbf{d}_l, \kappa_l) A_r(\mathbf{d}_r, \kappa_r) \quad (30)$$

We note that, in the first sum, we sum over all $\mathbf{d}_{1l}, \mathbf{d}_{1r} \leq \mathbf{c}_i^{\geq \ell_\sigma}$ which is consistent with earlier notation, but by definition of \mathbb{p} equals all $\mathbf{d}_{1l}, \mathbf{d}_{1r} \in C_\sigma$.

For our fast join operation, we need the sums $\mathbf{d}_{2l} + \mathbf{d}_{2r}$ and $\mathbf{d}_{3l} + \mathbf{d}_{3r}$ to be cyclic in order to use cyclic convolution. Therefore, we apply the same trick as in the proof of Lemma 10 and replace the tables A_l and A_r with expanded tables that include the sums of the labels in a state colouring as an additional parameter in the index. Here, we do so by defining this sum of the labels of a state colouring as the sum of the number in the labels, ignoring whether they are from C_σ or C_ρ and excluding the $|\leq \ell|_\sigma$ label. That is, let the projection function π on labels to be $\pi(|l|_\sigma) = \pi(|l|_\rho) = l$. Then, for a state colouring $\mathbf{c}_i \in C^k$, define $\Sigma(\mathbf{c}_i) = \Sigma([\mathbf{c}_i^{\geq \ell_\sigma}, \mathbf{c}_i^{\leq \ell_\sigma}, \mathbf{c}_i^\rho]) = \sum_{j=1}^{|\mathbf{c}_i^{\leq \ell_\sigma}|} \pi((\mathbf{c}_i^{\leq \ell_\sigma})_j) + \sum_{j=1}^{|\mathbf{c}_i^\rho|} \pi((\mathbf{c}_i^\rho)_j)$. For example, $\Sigma([|0|_\sigma, |0|_\rho, |\leq \ell|_\sigma]) = 0$ and $\Sigma([|2|_\sigma, |1|_\rho, |0|_\sigma]) = 3$. Now, we can define the expanded tables A_l, A_r as:

$$A_l(\mathbf{c}_l, \kappa_l, \iota_l) = \begin{cases} A_l(\mathbf{c}_l, \kappa_l) & \text{if } \Sigma(\mathbf{c}_l) = \iota_l \\ 0 & \text{otherwise} \end{cases} \quad (31)$$

Now, we can continue from (30) replacing the sums with sums coordinate-wise modulo ℓ_σ and $\ell_\rho + 1$, using the additional parameter to prevent the modular-cycling to happen for the result: if cycling occurs at a coordinate, the sums do

not add up any more. That is, we now compute $\zeta(A_i)(\mathbf{c}_i, \kappa_i)$ by evaluating the formula below using $\zeta(A_i)(\mathbf{c}_i, \kappa_i) = \zeta(A_i)(\mathbf{c}_i, \kappa_i, \Sigma(\mathbf{c}_i))$ where:

$$\zeta(A_i)(\mathbf{c}_i, \kappa_i, \iota_i) = \sum_{\mathbf{d}_{1l}, \mathbf{d}_{1r} \leq \mathbf{c}_i^{\geq \ell_\sigma}} \sum_{\mathbf{d}_{2l} + \mathbf{d}_{2r} \equiv \mathbf{c}_i^{< \ell_\sigma}} \sum_{\mathbf{d}_{3l} + \mathbf{d}_{3r} \equiv \mathbf{c}_i^\rho} \sum_{\kappa_l + \kappa_r = \kappa_i} \sum_{\iota_l + \iota_r = \iota_i} A_l(\mathbf{d}_l, \kappa_l, \iota_l) A_r(\mathbf{d}_r, \kappa_r, \iota_r) \tag{32}$$

Where $\mathbf{d}_{2l} + \mathbf{d}_{2r} \equiv \mathbf{c}_i^{< \ell_\sigma}$ is modulo ℓ_σ and $\mathbf{d}_{3l} + \mathbf{d}_{3r} \equiv \mathbf{c}_i^\rho$ is modulo $\ell_\rho + 1$ (the difference is due to the existence of the $|\geq \ell|_\sigma$ -label).

Next, we continue from Eq. 32 by change the order of summation, taking the outermost sum inwards, and by reordering the resulting inner terms:

$$= \sum_{\mathbf{d}_{2l} + \mathbf{d}_{2r} \equiv \mathbf{c}_i^{< \ell_\sigma}} \sum_{\mathbf{d}_{3l} + \mathbf{d}_{3r} \equiv \mathbf{c}_i^\rho} \sum_{\kappa_l + \kappa_r = \kappa_i} \sum_{\iota_l + \iota_r = \iota_i} \left(\sum_{\mathbf{d}_{1l} \leq \mathbf{c}_i^{\geq \ell_\sigma}} A_l(\mathbf{d}_l, \kappa_l, \iota_l) \right) \left(\sum_{\mathbf{d}_{1r} \leq \mathbf{c}_i^{\geq \ell_\sigma}} A_r(\mathbf{d}_r, \kappa_r, \iota_r) \right) \tag{33}$$

$$= \sum_{\mathbf{d}_{2l} + \mathbf{d}_{2r} \equiv \mathbf{c}_i^{< \ell_\sigma}} \sum_{\mathbf{d}_{3l} + \mathbf{d}_{3r} \equiv \mathbf{c}_i^\rho} \sum_{\kappa_l + \kappa_r = \kappa_i} \sum_{\iota_l + \iota_r = \iota_i} \zeta(A_l)([\mathbf{c}_i^{\geq \ell_\sigma}, \mathbf{d}_{2l}, \mathbf{d}_{3l}], \kappa_l, \iota_l) \zeta(A_r)([\mathbf{c}_i^{\geq \ell_\sigma}, \mathbf{d}_{2r}, \mathbf{d}_{3r}], \kappa_r, \iota_r) \tag{34}$$

Where in the last step, we apply the definition of the ζ -transform for \mathfrak{p} (Eq. 27). As a result, we obtain a standard convolution sum that can be evaluated using Lemma `reflem:combinedconv` given that the vertices with label $|\geq \ell|_\sigma$ in \mathbf{c}_i are fixed.

To be more precise, let us partition X_i into three parts $X_{\geq \ell_\sigma}$, $X_{< \ell_\sigma}$, X_ρ and say that a state colouring \mathbf{c}_i is compatible with this partition if: all vertices in $X_{\geq \ell_\sigma}$ have the $|\geq \ell|_\sigma$ -label; all vertices in $X_{< \ell_\sigma}$ have a label from $C_\sigma \setminus \{|\geq \ell|_\sigma\}$; and all vertices in X_ρ have a label from C_ρ . Then, given such a partition of X_i , Lemma 3 evaluates Eq. 34 for all \mathbf{c}_i compatible with $(X_{\geq \ell_\sigma}, X_{< \ell_\sigma}, X_\rho)$. Consequently, we can compute $\zeta(A_i)(\mathbf{c}_i, \kappa_i, \iota_i)$ for all values of \mathbf{c}_i , κ_i , and ι_i by enumerating all partitions of X_i into $(X_{\geq \ell_\sigma}, X_{< \ell_\sigma}, X_\rho)$ and evaluating Eq. 34 using Lemma 3 for each subset of compatible \mathbf{c}_i values, and then taking the results together.

As a result, we can evaluate Eq. 26 using a fast transform that takes the following steps in the following amount of operations:

- Expand the tables A_l and A_r taking the sums of the labels using Eq. 31 to $A_l(\mathbf{c}_l, \kappa_l, \iota_l)$ and $A_r(\mathbf{c}_r, \kappa_r, \iota_r)$. This takes $\mathcal{O}(s^{k+1}kn)$ time, as \mathbf{c}_l takes $\mathcal{O}(s^k)$ values, κ_l takes $\mathcal{O}(n)$ values and ι_l takes $\mathcal{O}(sk)$ values.
- Compute $\zeta(A_l)$ and $\zeta(A_r)$ in $\mathcal{O}(s^{k+1}k^2n)$ arithmetic operations using Proposition 7.

- Enumerate all partitions of X_i into $(X_{\geq \ell_\sigma}, X_{< \ell_\sigma}, X_\rho)$. For each such partition, compute the part of the table $\zeta(A_i)(\mathbf{c}_i, \kappa_i, \iota_i)$ for all \mathbf{c}_i using Eq. 34 that are compatible with this partition using Lemma 3. Then, combine the results to obtain $\zeta(A_i)(\mathbf{c}_i, \kappa_i, \iota_i)$ for all \mathbf{c}_i, κ_i and ι_i . For each partition $(X_{\geq \ell_\sigma}, X_{< \ell_\sigma}, X_\rho)$, this takes $\mathcal{O}((|C_\sigma| - 1)^{|X_{< \ell_\sigma}|} |C_\rho|^{|X_\rho|} n s k (k \log(s) + \log(n)))$ arithmetic operations by Lemma 3. By summing over all partitions and using the multinomial theorem, we find $\mathcal{O}(s^{k+1} k n (k \log(s) + \log(n)))$ arithmetic operation for this whole step, as:

$$\begin{aligned} & \sum_{\substack{(X_{\geq \ell_\sigma}, X_{< \ell_\sigma}, X_\rho) \\ \text{partition of } X_i}} (|C_\sigma| - 1)^{|X_{< \ell_\sigma}|} |C_\rho|^{|X_\rho|} \\ &= \sum_{\substack{x_1+x_2+x_3 \\ =k}} \binom{k}{x_1, x_2, x_3} 1^{x_1} (|C_\sigma| - 1)^{x_2} |C_\rho|^{x_3} = (|C_\sigma| + |C_\rho|)^k = s^k \end{aligned}$$

- Extract non-cycling values using $\zeta(A_i)(\mathbf{c}_i, \kappa_i) = \zeta(A_i)(\mathbf{c}_i, \kappa_i, \Sigma(\mathbf{c}_i))$.
- Compute the Möbius transform of the result obtaining A_i as $\mu(\zeta(A_i)) = A_i$. This is done $\mathcal{O}(s^k k n)$ arithmetic operations using Proposition 7.

By summing over these steps we conclude that the algorithm requires $\mathcal{O}(s^{k+1} k n (k \log(s) + \log(n)))$ arithmetic operations. \square

Above, the state colourings \mathbf{c} are split in three components $\mathbf{c}_i = [\mathbf{c}_i^{\geq \ell_\sigma}, \mathbf{c}_i^{< \ell_\sigma}, \mathbf{c}_i^\rho]$. If both σ and ρ are finite, we need no Möbius transforms and it would suffice to use $\mathbf{c}_i = [\mathbf{c}_i^\sigma, \mathbf{c}_i^\rho]$. If both σ and ρ are co-finite, we need Möbius transforms for both parts of the label set and would need to use $\mathbf{c}_i = [\mathbf{c}_i^{\geq \ell_\sigma}, \mathbf{c}_i^{< \ell_\sigma}, \mathbf{c}_i^{\geq \ell_\rho}, \mathbf{c}_i^{< \ell_\rho}]$, and adjust the partial order \mathfrak{p} in a way similar to as in Sect. 4.4.

Corollary 9. *Given a graph G with a tree decomposition T of G of width t , the optimisation variant of a $[\sigma, \rho]$ -domination problem that involves $s = |C|$ labels can be solved in $\mathcal{O}(s^{t+2} t n^2 (t \log(s) + \log(n)))$ arithmetic operations on $\mathcal{O}(t \log(s) + \log(n))$ -bit numbers.*

Proof. Plug Theorem 2 into Lemma 1 and observe that all arithmetic operations can be done using $\mathcal{O}(t \log(s) + \log(n))$ -bit numbers: the sum of all the entries in A_l and A_r is at most s^k , hence $t \log(s)$ bits, while we need the additional $\log(n)$ bits to store partial solution sizes (see also Corollary 4). The $t + 2$ in the exponent comes from the fact that $t \geq k - 1$ (the minus one in Definition 1) \square

We conclude by summarising results for the other variants of the $[\sigma, \rho]$ -domination problems.

Theorem 3 (results for $[\sigma, \rho]$ -domination problem variants). *Given a graph G with a tree decomposition T of G of width t , the different problem variants of a $[\sigma, \rho]$ -domination problem involving $s = |C|$ labels can be solved with the following amount of effort:*

	0	1 ₁	2	1 ₂		0	1 ₁	2	1 ₂
0	0	1 ₁	2	1 ₂	0	0	1 ₁	2	1 ₂
1 ₁	1 ₁	2			1 ₁	1 ₁	2	1 ₂	0
2	2				2	2	1 ₂	0	1 ₁
1 ₂	1 ₂		2		1 ₂	1 ₂	0	1 ₁	2

Fig. 2. The right join table is for a join for the LONGEST/HAMILTONIAN PATH/CYCLE and (PARITAL) CYCLE COVER problems, as described in the appendix of [15]. It is obtained by using a variant of ‘counting and filtering’ (Sect. 4.2) on the left join table, for which a fast FFT-based join exists.

- Existence: $\mathcal{O}(s^{t+2}t^2n \log(s))$ operations on $\mathcal{O}(t \log(s))$ -bit numbers.
- Optimisation: $\mathcal{O}(s^{t+2}tn^2(t \log(s) + \log(n)))$ operations on $\mathcal{O}(t \log(s) + \log(n))$ -bit numbers.
- Counting: $\mathcal{O}(s^{t+2}t^2n \log(s))$ operations on $\mathcal{O}(n)$ -bit numbers.
- Counting optimisation: $\mathcal{O}(s^{t+2}tn^2(t \log(s) + \log(n)))$ operations on $\mathcal{O}(n)$ -bit numbers.

Proof. The result for the optimisation problem follows from Corollary 9, and underlying Theorem 2.

For the counting optimisation problem, we use the same construction, only without expanding tables by solution sizes and extracting the non-zero entries: at every step of the algorithm, we let $A_i(\mathbf{c}, \kappa)$ be the number of partial solutions of size κ that correspond to the equivalence class identified by \mathbf{c} . The join then also comes down to evaluating Eq. 26, resulting in the same amount of arithmetic operations. For the existence and counting problems, we observe that we can remove the parameter κ at step of the algorithm, as solution sizes do not matter. Redoing the analysis of the resulting algorithm gives in the claimed amount of arithmetic operations.

For both counting problems variants, we need $\mathcal{O}(n)$ -bit numbers as there can be $\mathcal{O}(2^n)$ solutions to count. For the existence problem, we need $\mathcal{O}(t \log(s))$ as the sum of all entries in a $\mathcal{O}(s^{t+1})$ table with zero-one entries can be at most $\mathcal{O}(s^{t+1})$. □

6 Conclusion

In this paper, we have shown how Möbius and Fourier transforms can be used to speed-up computations for dynamic programming algorithms on tree decompositions. This led us to the currently fastest algorithm for the general case of the $[\sigma, \rho]$ -domination problems on tree decompositions. Additionally, we generalised the covering product from [2] from being defined on the subset lattice to more general partial orders (Lemma 6 and Theorem 1).

The same algebraic transforms can, and have been, used for many different problems. For example, the Möbius-transform-based approach has been used for clique packing, partitioning and covering problems such as PARTITION INTO TRIANGLES or MINIMUM COVER BY CLIQUES; see [26]. Also,

the Fourier-transform-based approach (the variant from [17]) has been used for BANDWIDTH [17] and CONNECTED VERTEX COVER [27]. The Fourier-transform-based approach in this paper originates from [15], where it was used together with counting and filtering to obtain the join table in Fig. 2 for LONGEST/HAMILTONIAN PATH/CYCLE and (PARITAL) CYCLE COVER.

We want to emphasise the more general observation that almost any join for which the join table has a certain ‘max’ or ‘addition’ or ‘modulo’ structure, or a combination of those, can be done fast using the tools from this paper. For example, the Fourier transform and corresponding cyclic convolution theorem can be used to obtain algorithms for problems where there is some modulo relation in the definition of the problem’s solution set D , e.g., an odd number of neighbours in D .

The approaches in this paper have wider use beyond tree decompositions. For example, they can be applied to branch decomposition instead of tree decompositions, obtaining faster algorithms there as well, with faster exact $\mathcal{O}(c^{\sqrt{n}})$ -time algorithms on planar graphs as direct corollaries [11, 23].

Open Problems. What we see is that, in order to use fast algebraic transforms, we embed the problem into algebraic structures that we further parameterise by the solution size (κ in the algorithms in this paper). However, without the replacement property (Definition 7), this leads to algorithms with super-linear dependence on n , while algorithms that are exponentially-slower in t but linear in n exists. Can we somehow remove this super-linear dependence on n ?

Moreover, when we consider weighted versions of the problems, the weights will appear in the run times of the exponentially-optimal algorithms. For the exponentially-slower algorithms (e.g., those by Alber et al. [1]) weights play no role in the worst-case running times. Can we somehow remove the dependence on the weights and obtain $\mathcal{O}^*(s^t)$ -time algorithms for weighted problems?

References

1. Alber, J., Bodlaender, H.L., Fernau, H., Kloks, T., Niedermeier, R.: Fixed parameter algorithms for dominating set and related problems on planar graphs. *Algorithmica* **33**(4), 461–493 (2002)
2. Björklund, A., Husfeldt, T., Kaski, P., Koivisto, M.: Fourier meets Möbius: fast subset convolution. In: Johnson, D.S., Feige, U. (eds.) 39th Annual ACM Symposium on Theory of Computing, STOC 2007, pp. 67–74. ACM Press (2007)
3. Björklund, A., Husfeldt, T., Kaski, P., Koivisto, M., Nederlof, J., Parviainen, P.: Fast zeta transforms for lattices with few irreducibles. *ACM Trans. Algorithms* **12**(1), 4:1–4:19 (2015)
4. Bodlaender, H.L.: Dynamic programming on graphs with bounded treewidth. In: Lepistö, T., Salomaa, A. (eds.) ICALP 1988. LNCS, vol. 317, pp. 105–118. Springer, Heidelberg (1988). https://doi.org/10.1007/3-540-19488-6_110
5. Bodlaender, H.L.: Polynomial algorithms for graph isomorphism and chromatic index on partial k -trees. *J. Algorithms* **11**(4), 631–643 (1990)
6. Bodlaender, H.L.: A tourist guide through treewidth. *Acta Cybernetica* **11**(1–2), 1–22 (1993)

7. Bodlaender, H.L.: Treewidth: algorithmic techniques and results. In: Prívvara, I., Ružička, P. (eds.) MFCS 1997. LNCS, vol. 1295, pp. 19–36. Springer, Heidelberg (1997). <https://doi.org/10.1007/BFb0029946>
8. Bodlaender, H.L.: A partial k -arboretum of graphs with bounded treewidth. *Theoret. Comput. Sci.* **209**(1–2), 1–45 (1998)
9. Bodlaender, H.L., Cygan, M., Kratsch, S., Nederlof, J.: Deterministic single exponential time algorithms for connectivity problems parameterized by treewidth. *Inf. Comput.* **243**, 86–111 (2015)
10. Bodlaender, H.L., Koster, A.M.C.A.: Combinatorial optimization on graphs of bounded treewidth. *Comput. J.* **51**(3), 255–269 (2008)
11. Bodlaender, H.L., van Leeuwen, E.J., van Rooij, J.M.M., Vatshelle, M.: Faster algorithms on branch and clique decompositions. In: Hliněný, P., Kučera, A. (eds.) MFCS 2010. LNCS, vol. 6281, pp. 174–185. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-15155-2_17
12. Borradaile, G., Le, H.: Optimal dynamic program for r -domination problems over tree decompositions. In: Guo, J., Hermelin, D. (eds.) 11th International Symposium on Parameterized and Exact Computation, IPEC 2016. Leibniz International Proceedings in Informatics, vol. 63, pp. 8:1–8:23. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2017)
13. Cooley, J.W., Tukey, J.W.: An algorithm for the machine calculation of complex Fourier series. *Math. Comput.* **19**(90), 297–301 (1965)
14. Cygan, M., Kratsch, S., Nederlof, J.: Fast Hamiltonicity checking via bases of perfect matchings. In: Boneh, D., Roughgarden, T., Feigenbaum, J. (eds.) 42nd Annual ACM Symposium on Theory of Computing, STOC 2010, pp. 301–310. ACM Press (2013)
15. Cygan, M., Nederlof, J., Pilipczuk, M., Pilipczuk, M., van Rooij, J.M.M., Wojaszczyk, J.O.: Solving connectivity problems parameterized by treewidth in single exponential time. [arXiv.org](https://arxiv.org/abs/1103.0534). The Computing Research Repository abs/1103.0534 (2011)
16. Cygan, M., Nederlof, J., Pilipczuk, M., Pilipczuk, M., van Rooij, J.M.M., Wojaszczyk, J.O.: Solving connectivity problems parameterized by treewidth in single exponential time. In: Ostrovsky, R. (ed.) 52nd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2011, pp. 150–159. IEEE Computer Society (2011)
17. Cygan, M., Pilipczuk, M.: Exact and approximate bandwidth. *Theoret. Comput. Sci.* **411**(40–42), 3701–3713 (2010)
18. Katsikarelis, I., Lampis, M., Paschos, V.T.: Structurally parameterized d -scattered set. In: Brandstädt, A., Köhler, E., Meer, K. (eds.) WG 2018. LNCS, vol. 11159, pp. 292–305. Springer, Cham (2018)
19. Kennes, R.: Computational aspects of the Möebius transformation of graphs. *IEEE Trans. Syst. Man Cybern.* **22**(2), 201–223 (1992)
20. Kloks, T.: Treewidth: Computations and Approximations. LNCS, vol. 842. Springer, Heidelberg (1994). <https://doi.org/10.1007/BFb0045375>
21. Korach, E., Solel, N.: Linear time algorithm for minimum weight Steiner tree in graphs with bounded treewidth. Technical report 632, Technion, Computer Science Department, Israel Institute of Technology, Haifa, Israel (1990)
22. Lokshtanov, D., Marx, D., Saurabh, S.: Known algorithms on graphs of bounded treewidth are probably optimal. In: Randall, D. (ed.) 22st Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2010, pp. 777–789. Society for Industrial and Applied Mathematics (2011)

23. Pino, W.J.A., Bodlaender, H.L., van Rooij, J.M.M.: Cut and count and representative sets on branch decompositions. In: Guo, J., Hermelin, D. (eds.) 11th International Symposium on Parameterized and Exact Computation, IPEC 2016. Leibniz International Proceedings in Informatics, vol. 63, pp. 27:1–27:12. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2017)
24. Rader, C.M.: Discrete Fourier transforms when the number of data samples is prime. *Proc. IEEE* **56**(6), 1107–1108 (1968)
25. van Rooij, J.M.M.: Exact exponential-time algorithms for domination problems in graphs. Ph.D. thesis, Department of Information and Computing Sciences, Utrecht University, Utrecht, The Netherlands (2011)
26. van Rooij, J.M.M., Bodlaender, H.L., Rossmanith, P.: Dynamic programming on tree decompositions using generalised fast subset convolution. In: Fiat, A., Sanders, P. (eds.) ESA 2009. LNCS, vol. 5757, pp. 566–577. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-04128-0_51
27. van Rooij, S.B., van Rooij, J.M.M.: Algorithms and complexity results for the capacitated vertex cover problem. In: Catania, B., Kráľovič, R., Nawrocki, J., Pighizzini, G. (eds.) SOFSEM 2019. LNCS, vol. 11376, pp. 473–489. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-10801-4_37
28. Telle, J.A.: Complexity of domination-type problems in graphs. *Nord. J. Comput.* **1**(1), 157–171 (1994)
29. Telle, J.A.: Vertex partitioning problems: characterization, complexity and algorithms on partial k -trees. Ph.D. thesis, Department of Computer and Information Science, University of Oregon, Eugene, Oregon, USA (1994)
30. Telle, J.A., Proskurowski, A.: Algorithms for vertex partitioning problems on partial k -trees. *SIAM J. Discrete Math.* **10**(4), 529–550 (1997)
31. Włodarczyk, M.: Clifford algebras meet tree decompositions. *Algorithmica* **81**(2), 497–518 (2019)
32. Yates, F.: The design and analysis of factorial experiments. Technical Communication No. 35, Commonwealth Bureau of Soil Science (1937)