

# A Faster Tree-Decomposition Based Algorithm for Counting Linear Extensions

**Kustaa Kangas**

Department of Computer Science, Aalto University, Espoo, Finland  
juho-kustaa.kangas@aalto.fi

**Mikko Koivisto**

Department of Computer Science, University of Helsinki, Helsinki, Finland  
mikko.koivisto@helsinki.fi

**Sami Salonen**

Department of Computer Science, University of Helsinki, Helsinki, Finland  
sami.m.salonen@helsinki.fi

---

## Abstract

We consider the problem of counting the linear extensions of an  $n$ -element poset whose cover graph has treewidth at most  $t$ . We show that the problem can be solved in time  $\tilde{O}(n^{t+3})$ , where  $\tilde{O}$  suppresses logarithmic factors. Our algorithm is based on fast multiplication of multivariate polynomials, and so differs radically from a previous  $\tilde{O}(n^{t+4})$ -time inclusion–exclusion algorithm. We also investigate the algorithm from a practical point of view. We observe that the running time is not well characterized by the parameters  $n$  and  $t$  alone, fixing of which leaves large variance in running times due to uncontrolled features of the selected optimal-width tree decomposition. For selecting an efficient tree decomposition we adopt the method of empirical hardness models, and show that it typically enables picking a tree decomposition that is significantly more efficient than a random optimal-width tree decomposition.

**2012 ACM Subject Classification** Theory of computation → Algorithm design techniques, Theory of computation → Parameterized complexity and exact algorithms

**Keywords and phrases** Algorithm selection, empirical hardness, linear extension, multiplication of polynomials, tree decomposition

**Digital Object Identifier** 10.4230/LIPIcs.IPEC.2018.5

**Funding** This work was partially supported by the Academy of Finland, Grant 276864, “Supple Exponential Algorithms.”

**Acknowledgements** We thank the anonymous reviewers for constructive suggestions, most of which are implemented in the present version of the paper.

## 1 Introduction

Consider a partially ordered set  $(V, \prec)$ , or *poset* for short, formed by an  $n$ -element set  $V$  and an irreflexive and transitive binary relation  $\prec$  on  $V$ , called a *partial order*. Another partial order  $<$  on  $V$  is a *linear extension* of  $\prec$  if it contains  $\prec$  and for any distinct elements  $x, y \in V$  either  $x < y$  or  $y < x$ . The problem of counting linear extensions ( $\#LE$ ) asks the number of linear extensions of a given poset; equivalent to  $\#LE$  is the problem of counting the topological sorts of a given directed (acyclic) graph. The problem has applications in numerous areas, for example, in sequence analysis [22], sorting [27], preference reasoning [21], convex rank tests [24], partial order plans [25], and learning graphical models [32, 26].



© Kustaa Kangas, Mikko Koivisto, and Sami Salonen;  
licensed under Creative Commons License CC-BY

13th International Symposium on Parameterized and Exact Computation (IPEC 2018).

Editors: Christophe Paul and Michal Pilipczuk; Article No. 5; pp. 5:1–5:13

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

The #LE problem is #P-complete [11] but admits a fully polynomial approximation scheme [13]. The best known asymptotic bounds for the expected running time are  $O(\epsilon^{-2}n^3 \log^2 \ell \log n)$  [6] and  $O(\epsilon^{-2}n^5 \log^2 n)$  [30], where  $\ell$  is the number of linear extensions and  $\epsilon$  the allowed relative error. These bounds, while polynomial, become prohibitively large in practice if, say, one requires an accuracy of  $\epsilon = 0.01$  and  $n$  is around one hundred.

Exact and parameterized algorithms offer an alternative approach to design practical algorithms for the problem. For the required number of arithmetic operations, the best known worst-case bound is  $O(2^n n)$  in general, which is achieved by a simple dynamic programming algorithm. For several special instance classes better bounds are known:  $O(n^w w)$  for width- $w$  posets,  $O(n^2)$  for series-parallel posets [23], and  $O(n^2)$  also for posets whose cover graph is a forest [5]; the *cover graph* of a poset  $(V, \prec)$  is the directed graph  $(V, E)$  where the edge set  $E$  is the transitive reduction of  $\prec$  (a.k.a. Hasse diagram). If parameterized by the treewidth of the cover graph,  $t$ , the problem can be solved with  $O(n^{t+3})$  arithmetic operations by an inclusion-exclusion algorithm [17]. On the other hand, the problem parameterized by  $t$  is W[1]-hard [14], and so it may be difficult, or even impossible, to find an algorithm that runs in time  $O(f(t)n^d)$  for some computable function  $f$  and constant  $d$ .

## 1.1 Theoretical contributions

In this paper, we give a new algorithm that exploits small treewidth. Write  $\tilde{O}(f)$  as a shorthand for  $f(\log f)^{O(1)}$ . We prove the following in Section 3.

► **Theorem 1 (Main).** *The linear extensions of a given  $n$ -element poset can be counted with  $\tilde{O}(n^{t+3})$  bit-operations, where  $t$  is the treewidth of the cover graph of the poset.*

We state the bound in terms of bit-complexity to emphasize the fact that the dominating computations deal with large integers. Indeed, the bounds stated in the previous paragraph refer to the number of arithmetic operations. Thus, in particular, our bound improves the previous bound of Kangas et al. [17] by a factor of  $n$ . For large  $n$  and small  $t$  the improvement is relatively significant; for instance, for  $t = 2$  the bound is reduced from  $\tilde{O}(n^6)$  to  $\tilde{O}(n^5)$ .

Perhaps more importantly, the design of our algorithm is very different from that of the inclusion-exclusion algorithm. In the latter, the idea is to view a linear extension as a bijective mapping and then remove the global bijectivity constraint by inclusion-exclusion, similarly to previous applications to matrix permanent [29], Hamiltonian path [18], and set partitioning [19], but incurring only a polynomial overhead. Once the bijectivity constraint is removed, what remains is a collection of simpler subproblems with local constraints. The subproblems can be handled by standard routines that exploit small treewidth [7, 12]; see Section 2.4 for some additional details.

The present algorithm, in contrast, takes care of the bijectivity constraint within dynamic programming along a tree decomposition and is, with this respect, similar to a folklore  $t^{O(t)}n$ -time algorithm for the Hamiltonian path problem. However, #LE being W[1]-hard one may expect it to require a significantly larger dynamic programming table. We give a formulation, where each node of a tree decomposition is associated with  $\Theta(n^{t+1})$  counts. This formulation leads to a challenge: a step in the dynamic program that combines two (or more) arrays of such counts appears to require a quadratic number of arithmetic operations,  $\Theta(n^{2t+2})$ , if implemented in a straightforward manner. Fortunately, we discover that the key ingredient of the step takes a form of multidimensional convolution, which we can compute efficiently using known (deep) results for fast multiplication of multivariate polynomials.

Concerning space complexity we only make a couple of observations here: Both our algorithm and the inclusion-exclusion algorithm by Kangas et al. [17] require  $\tilde{O}(n^{t+2})$  bits of

space; this can be reduced by a factor about linear in  $n$  by carrying the computations modulo several small relative primes and constructing the final output using the Chinese remainder theorem. The simple dynamic programming algorithm requires  $\tilde{O}(2^n)$  bits of space.

## 1.2 Empirical contributions

In the second part of the paper we address the practical value of the algorithm. Given that the present algorithm is technically more convoluted than the inclusion–exclusion algorithm, it is natural to ask, whether the improvement in the asymptotic worst case time requirement can be realized in practice. Our interest is particularly in instances where  $t$  is small (at most four) and  $n$  ranging up to a few hundred.

A well known challenge in practical implementation of tree-decomposition based algorithms is that finding an optimal-width tree decomposition may be insufficient for minimizing the computational cost: the running time of the dynamic programming algorithm can be sensitive to the shape of the tree decomposition. Bodlaender and Fomin [9] addressed this issue from a theoretical viewpoint by studying the complexity of finding a tree decomposition that minimizes a sum of costs associated with each node of a tree decomposition. In their *f-cost* framework the cost of a node is allowed to depend only on the width of the node (i.e., the size of the associated bag; see Section 2). Recently, Abseher et al. [1, 2] presented a more general and more practical heuristic approach. Their `htd` library [1] allows a user to generate a variety of optimal-width tree decompositions and also (locally) optimize a given cost function. Moreover, they proposed and evaluated [2] a method to learn an appropriate cost function, or regression model, from collected empirical data on running times on varying instances. The method can be viewed as an instantiation of the method of empirical hardness models [20] for the algorithm selection problem [28].

Following these ideas we have implemented and tested our algorithm for `#LE` using a collection of synthetically generated instances (posets) together with a variety of tree decompositions generated by `htd` for each instance. We will report on and discuss our preliminary observations, which suggest that selecting the tree decomposition using a learned regression model can make a difference, at least for the smallest treewidth ( $t = 2$ ): compared to the median running time over generated tree decompositions, the selected one typically yields almost an order-of-magnitude speedup.

## 1.3 Organization

Some preliminary material is given in Section 2. Section 3 is devoted to proving Theorem 1. In Section 4 we describe some implementation details and report on empirical results. We conclude in Section 5 by highlighting the main observations and some open questions.

## 2 Preliminaries

In this section we introduce some basic terminology, notation, and facts.

### 2.1 Basic notation

We denote by  $\mathbb{N}$  the set of natural numbers  $\{0, 1, 2, \dots\}$ . For two sets  $S$  and  $U$  we write  $S^U$  for the set of functions from  $U$  to  $S$ . If  $m \in \mathbb{N}$  we write  $[m]$  for the set  $\{1, \dots, m\}$ . By  $S^m$  we denote the set of  $m$ -tuples  $a = (a_i)_{i=1}^m$  with  $a_i \in S$ .

The *restriction* of a function  $\alpha : U \rightarrow S$  to a subset  $A \subset U$  is defined in the standard manner and denoted by  $\alpha|_A$ ; conversely, we say that  $\alpha$  is an *extension* of  $\alpha|_A$ . We also denote by  $\alpha^{v \rightarrow i}$  the extension of  $\alpha$  that we obtain by mapping an added element  $v \notin U$  to  $i$ .

We use the Iverson's bracket notation: for a proposition  $P$ , the expression  $[P]$  evaluates to 1 if  $P$  is true, and to 0 otherwise.

## 2.2 Factorials and multiplication of polynomials

Let  $a = (a_1, \dots, a_k) \in \mathbb{N}^k$ . We denote  $a! := a_1! \cdots a_k!$  and  $|a| := a_1 + \dots + a_k$ . Since  $|a|!/a!$  is a multinomial coefficient, we have the following.

► **Fact 2.** *If  $a$  is a tuple of nonnegative integers, then  $a!$  divides  $|a|!$ .*

Another fact we need concerns the complexity of multiplying two multivariate polynomials that are sparse in the sense that their total degree is given an upper bound, while the degrees of the individual variables may be large (not larger than the total degree, of course). We assume that a polynomial is represented by a list of its coefficients. The following result can be obtained by a straightforward specialization of a more general and detailed bound due to van der Hoeven and Lecerf [31, Cor. 4].

► **Fact 3.** *Two  $k$ -variate polynomials with  $\tilde{O}(n)$ -bit integer coefficients and total degree at most  $n$  can be multiplied with  $\tilde{O}\left(\binom{n+k}{k}(n+k)\right)$  bit-operations.<sup>1</sup>*

For univariate polynomials the bound is quadratic in  $n$  because the coefficients can be large.

## 2.3 Tree decomposition

► **Definition 4** (Tree decomposition). A *tree decomposition* of a graph  $G = (V, E)$  is a pair  $(T, B)$  where  $T = (I, F)$  is a tree and  $B$  maps each  $x \in I$  to a bag  $B_x \subseteq V$  such that

1. for each  $v \in V$ , the set  $\{x \in I : v \in B_x\}$  induces a nonempty connected subtree of  $T$ ,
2. for each  $uv \in E$ , there exists  $x \in I$  with  $u, v \in B_x$ .

The *width* of the tree decomposition is the size of its largest bag minus one,  $\max_{x \in I} |B_x| - 1$ , and the *treewidth* of a graph is the minimum width over all its tree decompositions.

For a graph of treewidth  $t$ , we can find a tree decomposition of width  $t$  in time  $t^{O(t^3)}n$  using Bodlaender's algorithm [8] or in time  $\tilde{O}(n^{t+2})$  using the approach of Arnborg et al. [4].

A tree decomposition is *rooted* if the edges of the tree are directed so that there is a unique node, the *root*, that has no parent. Clearly, one obtains a rooted tree decomposition by simply choosing one node as the root and directing the edges accordingly.

► **Definition 5** (Nice tree decomposition). A rooted tree decomposition  $(T, B)$  of a graph  $G = (V, E)$  is *nice* if each node  $x$  of  $T$  is of one of the following types:

- (**leaf**)  $x$  has no children and  $|B_x| = 1$ ;
- (**introduce**)  $x$  has a unique child  $y$  and  $B_x = B_y \cup \{v\}$  for some  $v \in V \setminus B_y$ ;
- (**forget**)  $x$  has a unique child  $y$  and  $B_x = B_y \setminus \{v\}$  for some  $v \in B_y$ ;
- (**join**)  $x$  has exactly two children  $y, z$  and  $B_x = B_y = B_z$ .

We can convert a given tree decomposition of width  $t$  and  $O(n)$  nodes into a nice tree decomposition of width  $t$  and  $O(tn)$  nodes in time  $t^{O(1)}n$  [10].

<sup>1</sup> In fact, this result assumes we have access to a  $\tilde{\Omega}(n)$ -bit prime  $p$  and to a primitive element (i.e., generator) of  $\mathbb{F}_p^*$ , a multiplicative field of order  $p$ . Finding such numbers seem hard, theoretically: no deterministic polynomial time algorithm is known for the former, and no polynomial time algorithm is known for the latter. The result thus concerns the non-uniform model of computational complexity.

## 2.4 Counting linear extensions via inclusion–exclusion

Kangas et al. [17] showed that the number of linear extensions of a poset  $(V, \prec)$  whose cover graph is  $G = (V, E)$  is given by the formula

$$\sum_{k=1}^n \binom{n}{k} (-1)^{n-k} \sum_{\tau} \prod_{uv \in E} [\tau(u) < \tau(v)],$$

where  $\tau$  runs over all functions from  $V$  to  $[k]$ .

Supposing  $G$  has treewidth  $t$ , it is well known that there is an elimination ordering  $v_1, \dots, v_n$  of the vertices, such that when removing the vertices from the graph in this order and always connecting the neighbors of the removed vertex, the size of the largest clique in the obtained graphs is  $t + 1$ . The  $n$ -dimensional inner summation over the variables  $\tau(v_i)$ , for  $i \in [n]$ , can be processed iteratively along such an ordering, the  $i$ th one-dimensional summation over  $\tau(v_i)$  requiring  $O(n_i k^{t+1})$  additions and multiplications of  $O(n \log n)$ -bit numbers, for some  $n_1, \dots, n_n$  that sum up to  $O(n)$ . In total, the evaluation of the inclusion–exclusion formula thus requires  $\tilde{O}(n^{t+4})$  bit-operations. We omit a more detailed treatment of the algorithm, as the method applied for computing the inner summation is standard. (The original analysis of Kangas et al. [17] use a looser bound of  $n_i = O(n)$  for each  $i$ , arriving at a bound that is larger by a factor of  $n$ .)

### 3 The algorithm – proof of Theorem 1

We implement a standard recipe of designing a tree-decomposition based algorithm. The outline of the algorithm is as follows.

- A1** Compute the cover graph of the input poset.
- A2** Find a minimum-width nice tree decomposition of the cover graph.
- A3** Run dynamic programming over the nice tree decomposition.

We will next consider each step in detail. We will see that the last step dominates our asymptotic running time bounds.

#### 3.1 Computing the cover graph

The cover graph  $G = (V, E)$  is obtained by computing the transitive reduction of the input poset  $(V, \prec)$ . The transitive reduction can be computed in time  $O(|V| \cdot |\prec|)$  [3], which is  $O(n^{t+2})$  for all  $t \geq 1$ .

#### 3.2 Finding a minimum-width nice tree decomposition

As mentioned in Section 2, if the cover graph has treewidth  $t$ , then a width- $t$  nice tree decomposition of the cover graph can be found in  $\tilde{O}(n^{t+2})$  time.

#### 3.3 Dynamic programming

Suppose now that a width- $t$  nice tree decomposition  $(T, B)$  of the cover graph  $G$  is available. Our idea will be to associate each node of  $T$  with an array of numbers such that (i) the numbers at the root node are sufficient for computing the number of linear extensions and (ii) the array of a node can be computed from the arrays of its child nodes.

The following notation will be useful. Denote by  $V_x$  the set of vertices covered by the subtree of  $T$  rooted at  $x$ , that is,  $V_x$  is the union of the bags  $B_y$  of nodes  $y$  to which there is a directed path from  $x$ . Write  $n_x$  for the size  $|V_x|$  and  $E_x$  for set of edges in the induced graph  $G[V_x]$ .

Now, for each node  $x \in T$  and injection  $\alpha \in [n_x]^{B_x}$ , define  $\ell_x(\alpha)$  as the number of bijections  $\pi \in [n_x]^{V_x}$  such that  $\pi(v) = \alpha(v)$  for all  $v \in B_x$ , and  $\pi(u) < \pi(v)$  whenever  $uv \in E_x$ . In other words,  $\ell_x(\alpha)$  is the number of ways to extend  $\alpha$  to a linear extension of the induced poset  $(V_x, \prec \cap (V_x \times V_x))$ , where we view a linear extension as a bijection from  $V_x$  to  $[n_x]$  that satisfies the ordering constraints.

We begin by showing that the values  $\ell_x(\alpha)$  are sufficient for computing the number of linear extensions of the poset, that is, they satisfy the listed conditions (i) and (ii). After that we consider the time requirement of computing the values  $\ell_x(\alpha)$  for each node of the nice tree decomposition.

Consider first the root node.

► **Lemma 6 (Root).** *We have  $\ell(V) = \sum_{\alpha} \ell_r(\alpha)$ , where  $\alpha$  runs over all injections in  $[n_x]^{B_r}$ .*

**Proof.** Since  $V_r = V$ ,  $E_r = E$ , and  $n_r = n$ , we have that

$$\sum_{\alpha} \ell_r(\alpha) = \sum_{\pi} \prod_{uv \in E} [\pi(u) < \pi(v)] = \ell(V),$$

where  $\alpha$  and  $\pi$  run over all injections in  $[n]^{B_r}$  and  $[n]^V$ , respectively. ◀

Next we will show separately for each node type of the nice tree decomposition, how the values  $\ell_x(\alpha)$  are determined by the corresponding values for the child node or child nodes of  $x$ . For all but join nodes the results are immediate, and we omit the proofs.

► **Lemma 7 (Leaf).** *If  $x$  is a leaf node, then  $\ell_x(\alpha) = 1$  for the unique injection  $\alpha$  in  $[n_x]^{B_x}$ .*

For an introduce node, we simply restrict the injection  $\alpha$  to the bag in question and check that the ordering constraint holds.

► **Lemma 8 (Introduce).** *If  $x$  is an introduce node with child  $y$ , then*

$$\ell_x(\alpha) = \ell_y(\alpha|_{B_y}) \prod_{u,v \in B_x: uv \in E} [\alpha(u) < \alpha(v)].$$

For a forget node, we extend the injection  $\alpha$  to the larger bag by mapping the new vertex to some value.

► **Lemma 9 (Forget).** *If  $x$  is a forget node with child  $y$  and  $B_x = B_y \setminus \{v\}$ , then*

$$\ell_x(\alpha) = \sum_{a=1}^{n_y} \ell_y(\alpha^{v \rightarrow a}).$$

To handle a join node, we introduce some convenient notation. Let  $\alpha$  be an injection from a  $k$ -element set  $S$  to a range of integers  $[m]$ . Label the elements of  $S$  such that  $\alpha(v_1) < \dots < \alpha(v_k)$ . For  $i = 1, \dots, k-1$ , denote by  $\alpha_i$  the number of integers between  $\alpha(v_i)$  and  $\alpha(v_{i+1})$ , that is,  $\alpha_i := \alpha(v_{i+1}) - \alpha(v_i) - 1$ ; in addition, denote  $\alpha_0 := \alpha(v_1) - 1$  and  $\alpha_k := m - \alpha(v_k)$ . Observe that  $\alpha_0 + \dots + \alpha_k = m - k$ . Furthermore, if  $\beta$  is another injection from  $S'$  to  $[m']$ , write  $\beta \sim \alpha$  if  $\beta$  and  $\alpha$  specify the same linear order on  $S \cap S'$ , that is,  $\beta(u) < \beta(v)$  if and only if  $\alpha(u) < \alpha(v)$  for all  $u, v \in S \cap S'$ .

► **Lemma 10** (Join). *If  $x$  is a join node with children  $y$  and  $z$ , then*

$$\ell_x(\alpha) = \sum_{\beta} \sum_{\gamma} [\alpha \sim \beta \sim \gamma] \prod_{i=0}^{|B_x|} [\alpha_i = \beta_i + \gamma_i] \binom{\alpha_i}{\beta_i} \ell_y(\beta) \ell_z(\gamma),$$

where  $\beta$  and  $\gamma$  run over all injections in  $[n_y]^{B_y}$  and  $[n_z]^{B_z}$ , respectively.

**Proof.** By definition,

$$\ell_x(\alpha) = \sum_{\pi} \prod_{uv \in E_x} [\pi(u) < \pi(v)],$$

where  $\pi$  runs over all bijections from  $V_x$  to  $[n_x]$  that extend  $\alpha$ . Observe that by the tree decomposition properties, the sets  $V_y \setminus B_x$  and  $V_z \setminus B_x$  are disjoint and their union is  $V_x \setminus B_x$ . Thus we may represent any bijection  $\pi : V_x \rightarrow [n_x]$  that extends  $\alpha$  uniquely by a pair of injections  $\beta' : V_y \rightarrow [n_x]$  and  $\gamma' : V_z \rightarrow [n_x]$  whose restrictions to  $B_x$  are equal to  $\alpha$  and whose images  $\beta'(V_y)$  and  $\gamma'(V_z)$  cover  $[n_x]$ .

$$\ell_x(\alpha) = \sum_{\beta'} \sum_{\gamma'} [\beta'(V_y) \cup \gamma'(V_z) = [n_x]] \prod_{uv \in E_y} [\beta'(u) < \beta'(v)] \prod_{uv \in E_z} [\gamma'(u) < \gamma'(v)],$$

where  $\beta'$  and  $\gamma'$  run over all injections that extend  $\alpha$  in  $[n_x]^{V_y}$  and  $[n_x]^{V_z}$ , respectively.

Consider then a mapping that “compresses” any such injection  $\beta'$  into a bijection  $\beta'' : V_y \rightarrow [n_y]$  by letting  $\beta''(v) := |\{u \in V_y : \beta'(u) \leq \beta'(v)\}|$ ; let  $\gamma''$  denote the bijection obtained similarly from an injection  $\gamma'$ . Let  $\beta$  denote the restriction of  $\beta''$  to  $B_x$  and  $\gamma$  the restriction of  $\gamma''$  to  $B_x$ . We have that  $\beta'$  and  $\gamma'$  extend  $\alpha$  if and only if  $\beta \sim \alpha$  and  $\gamma \sim \alpha$ . Thus we get that

$$\ell_x(\alpha) = \sum_{\beta'' : \beta \sim \alpha} \sum_{\gamma'' : \gamma \sim \alpha} \prod_{i=0}^{|B_x|} [\alpha_i = \beta_i + \gamma_i] \binom{\alpha_i}{\beta_i} \prod_{uv \in E_y} [\beta''(u) < \beta''(v)] \prod_{uv \in E_z} [\gamma''(u) < \gamma''(v)],$$


where  $\beta''$  and  $\gamma''$  run over all bijections in  $[n_y]^{B_y}$  and  $[n_z]^{B_z}$ , respectively. The product of the binomial coefficients  $\binom{\alpha_i}{\beta_i}$  is the number of pairs  $(\beta', \gamma')$  that map to the same pair  $(\beta'', \gamma'')$ , that is, the number of interleavings of the  $\beta_i + \gamma_i$  elements to the range of  $\alpha_i$  elements.

To complete the proof, it suffices to write the summation over  $\beta''$  as a double-summation: the outer summation being over all injections  $\beta : B_y \rightarrow [n_y]$  and the inner summation being over all bijections  $\beta'' : V_y \rightarrow [n_y]$  that extend  $\beta$ ; similarly for the summation over  $\gamma''$ . ◀

► **Example 11** (Join in a tree). Consider the example illustrated in Figure 1. The cover graph is a tree with vertex set  $V = \{a, b, c, d, e, f, g\}$ . In a nice tree decomposition (not shown) the root node  $x$  is a join of nodes  $y$  and  $z$ , with  $B_y = B_z = B_x = \{c, f\}$ . The vertex sets associated with the nodes are  $V_x = V$ ,  $V_y = \{a, c, d, e, f\}$ , and  $V_z = \{b, c, f, g\}$ . For the shown injection  $\alpha$ , the value of  $\ell_x(\alpha)$  is obtained by a sum of  $\ell_y(\beta) \cdot \ell_z(\gamma)$  over valid pairs  $(\beta, \gamma)$ , multiplied by the number of possible interleavings, which is given by a product of binomial coefficient. Shown is one pair  $(\beta, \gamma)$ , for which  $\ell_y(\beta) = 1$  and  $\ell_z(\gamma) = 1$  and the number of interleavings is equal to  $\binom{2}{1} \binom{2}{2} \binom{1}{0} = 4$ .

It remains to bound the running time of the algorithm.

► **Lemma 12** (Time complexity). *Given a width- $t$  nice tree decomposition of the cover graph of an  $n$ -element poset, the linear extensions can be counted in  $\tilde{O}(n^{t+3})$  bit-operations.*



■ **Figure 1** An illustration of the recurrence for a join node; see Example 11 for a description.

**Proof.** Let  $x$  be a node in the nice tree decomposition. For brevity, denote  $k := |B_x|$ . If  $x$  is a leaf node, introduce node, or forget node, then the values  $\ell_x(\alpha)$  for all injections  $\alpha \in [n_x]^{B_x}$  can clearly be computed using  $O(kn^k) = O(tn^{t+1})$  basic operations, some of which are additions of two  $O(n \log n)$ -bit numbers, thus using  $\tilde{O}(n^{t+2})$  bit-operations (observe that the factor  $t$  is  $O(\log n^{t+2})$  and can thus be omitted).

Consider then the remaining case:  $x$  is a join node. Let  $y$  and  $z$  be the two children of  $x$ . Recall that  $B_x = B_y = B_z$ .

Represent an injection  $\alpha$  in  $[n_x]^{B_x}$  as a pair  $(\sigma, a)$ , where  $a = (a_i)_{i=1}^{k+1}$  with  $a_i = \alpha_{i-1}$  and  $\sigma$  is a bijection from  $B_x$  to  $[k]$  that captures the specified linear order, that is,  $\sigma(u) < \sigma(v)$  if  $\alpha(u) < \alpha(v)$ . Clearly, the mapping  $\alpha \mapsto (\sigma, a)$  is a bijection when we require that  $a_i \in \mathbb{N}$  and  $|a| = n_x - k$ . Using this representation and Lemma 10, write

$$\ell_x(\sigma, a) = \sum_b \sum_c \prod_{i=0}^k [a_i = b_i + c_i] \binom{a_i}{b_i} \ell_y(\sigma, b) \ell_z(\sigma, c),$$

where  $b$  and  $c$  run over  $\mathbb{N}^{k+1}$ . By writing  $\ell'_x(\sigma, a) := \ell_x(\sigma, a)/a!$ , we get the convolution form

$$\ell'_x(\sigma, a) = \sum_{a=b+c} \ell'_y(\sigma, b) \ell'_z(\sigma, c).$$

To treat this as a multiplication of multivariate polynomials, consider a fixed bijection  $\sigma$  and let  $P_x(r_1, \dots, r_k)$  be the  $k$ -variate polynomial where the coefficient of  $r_1^{a_1} \dots r_k^{a_k}$  equals  $n! \cdot \ell'_x(\sigma, a)$ ; we define  $P_y$  and  $P_z$  similarly. Note that  $k$  variables suffice, since  $a_{k+1}$  is determined by the fixed  $|a|$ . Here we multiplied by the factorial  $n!$  to get integer coefficients (by Fact 2, since  $n \geq |a|, |b|, |c|$ ). We have that  $n! \cdot P_x = P_y P_z$ , that is, we obtain  $P_x$  by multiplying  $P_y$  and  $P_z$  and dividing each coefficient of the resulting polynomial by  $n!$ .

We bound the bit-complexity of the polynomial multiplication using Fact 3. The total degrees of the polynomials are at most  $n - k$ . Each coefficient of the polynomials is a  $\tilde{O}(n)$ -bit integer. Thus the multiplication takes  $\tilde{O}\left(\binom{n}{k} n\right)$  bit-operations.

Multiplying the obtained bound by the number of bijections  $\sigma$ , we get that all  $\ell_x(\sigma, a)$  can be computed using  $\tilde{O}(n^{k+1}) = \tilde{O}(n^{t+2})$  bit-operations.

Because there are  $O(tn)$  nodes in the nice tree decomposition,  $\tilde{O}(n^{t+3})$  bit-operations suffice in total. ◀

## 4 Experiments

This section is devoted to empirical results. We first describe our implementation of the algorithm and the test instances used in the experiments. Then we show how the performance on the algorithm depends on the way we choose the tree decomposition.



## 4.1 Implementation

We have implemented the algorithm described in Section 3 in a C++ program `Countle`.<sup>2</sup> For multiplication of polynomials we used a C library `FLINT` [16]. For finding an optimal tree decomposition we used the C++ library `htd` [1]. We ran all experiments on machines with Intel Xeon E5540 CPUs; the same machines were used by Kangas et al. [17], which makes their running time measurements directly comparable to the present results.

Two implementation details are worth mentioning. First, for multiplication of multivariate polynomials, we transformed the polynomials into univariate polynomials using an appropriate Kronecker substitution. Specifically, separately for each node  $x$  of the tree decomposition and the considered vertex ordering (bijection)  $\sigma$ , we encode a  $k$ -variate polynomial in variables  $r_1, \dots, r_k$  as a univariate polynomial in variable  $s$  by substituting  $r_j := s^{(d_1+1)\cdots(d_{j-1}+1)}$ , where each  $d_i$  is an upper bound for the degree of  $r_i$  in the polynomial. Using knowledge associated with the node  $x$  and ordering  $\sigma$ , we aim at finding a value  $d_i$  that is smaller than the trivial upper bound  $n_x - k$ . To this end, we set  $d_i$  to the sum of the largest realized exponents of  $r_i$  in the already computed polynomials for the two child nodes of  $x$ .

Second, we wish to ignore any impossible ordering  $\sigma$  at a node  $x$  of the tree decomposition, and so save both time and space. The key observation is that, even if the value  $\ell_x(\sigma, a)$  is nonzero, we can ignore it if  $\sigma$  assigns some two vertices in the bag  $B_x$  an order that violates the partial order  $\prec$ , that is, for some  $u, v \in B_x$  we have  $u \prec v$  and  $\sigma(u) > \sigma(v)$ .

## 4.2 Instances


We generated random instances (posets) of different sizes for small values of treewidth  $t$ . We varied the number of elements  $n$  from 10 to 199 ( $t = 2$ ), 109 ( $t = 3$ ), and 59 ( $t = 4$ ). For each pair  $(t, n)$  we generated 5 posets; following Kangas et al. [17], we let each poset be a “grid tree,” constructed by randomly joining  $t$ -by- $t$  grids along the (boundary) edges, orienting the edges so that no directed cycles are introduced, and finally taking the transitive closure.

## 4.3 Results

Because the running time of `Countle` may be sensitive to the particular (nice) tree decomposition selected, we ran the program on 50 optimal-width tree decompositions, which we generated using `htd`; we checked that for every encountered instance, `htd` indeed generated tree decompositions of optimal width. We allowed each individual run to take up to 20 minutes of CPU time and 30 GB of memory.

First we considered the scaling of `Countle` in terms of the number of elements  $n$  and treewidth  $t$ . We observed that, while the growth of the running time follows the rate suggested by the worst case bound, there is significant variance in the running times for any fixed  $(n, t)$ , due to differences in the five posets and the 50 tree decompositions per poset (compare *median* to *best* in Figure 2). Compared to an implementation of the inclusion–exclusion algorithm by Kangas et al. [17], `VEIE`, we find that `Countle` is an order of magnitude faster. For example, `Countle` can solve a typical (median) poset with a typical (median) tree decomposition in about 20 seconds if  $n = 100$  and  $t = 2$ , or if  $n = 50$  and  $t = 3$ , while `VEIE` requires about 200 seconds on such instances [17, Fig. 8] (not reproduced here). We have observed that the same pattern also holds for  $t = 4$ ; we omit the detailed results here, as such posets are handled faster by another, exponential time algorithm of Kangas et al. [17, Fig. 8].

<sup>2</sup> `Countle` is free and publicly available at <https://bitbucket.org/samsalo/countledist/>.



■ **Figure 2** The running time of `Countle` (in seconds) on random “grid tree” posets of treewidth 2 and 3, with a varying number of elements  $n$ . For each poset, we generated 50 optimal-width tree decompositions and collected the median running time and the running times for the selected and the best tree decomposition. Shown are the medians of these three statistics over five independent posets for each value of  $n$ .

Then we investigated whether one can efficiently select a near-optimal tree decomposition from a collection of generated candidates. The observed variance in running times suggests that, if successful, this could lead to a significant expedition of `Countle`, by up to one order of magnitude. To construct a “selector” we applied the method of Abseher et al. [2] in a straightforward manner:

- C1** We collected a data set of measured running times for multiple pairs of posets and tree decompositions. We used the procedure described in the previous section, except that we used a single poset (instead of five) for each combination of  $n$  and  $t$ . If a run was not completed within the 20-minute time limit, we simply discarded the instance (and thus introduced some bias).
- C2** We computed for each tree decomposition the values of several features, such as statistics of bag sizes (by node type), node depths (by node type), and distances between join nodes; for a full feature list, see Abseher et al. [2].
- C3** We fitted a multivariate linear regression model, separately for each  $t = 2, 3, 4$ , with the features as the predictor variables and the logarithm of the running time as the response variable. We used the machine learning software `WEKA 3.6.13` [15] with default options. To select a tree decomposition for a given new poset, we first generated 50 candidate tree decompositions for the poset, and then selected the one for which the model predicted the shortest running time.

We observed (Figure 2) that, for  $t = 2$ , the model is almost always able to select a top-3 tree decomposition, which yields a nearly as good (i.e., short) running time as the best among the 50 tree decompositions. For  $t = 3$  the performance degrades: the model is usually able to select a top-10 tree decomposition, which yields a running time that is systematically better than for a typical (median) tree decomposition, yet not quite achieving the performance of the best among the generated candidates. For  $t = 4$  the performance degrades further, yet being better than by selecting a random tree decomposition (results not shown). In more quantitative terms, the proportions of tree decompositions (among 50) better than the selected one were 2.5 %, 12 %, 28 % for  $t = 2, 3, 4$ , respectively; these numbers are medians of averages over 5 test posets (one per fixed  $n$  and  $t$ ).

## 5 Concluding remarks

We have presented a new tree-decomposition based algorithm for counting linear extensions. The algorithm relies on fast multiplication of multivariate polynomials, thus differing radically from the inclusion–exclusion approach of Kangas et al. [17]. For any constant treewidth  $t$  the obtained asymptotic speedup is about linear in the number of elements  $n$ .

A question not settled here is whether one could save another factor of  $n$ , that is, solve the problem in time  $\tilde{O}(n^{t+2})$ . The present authors find this question particularly intriguing for two reasons: One is that for finding an optimal-width tree decomposition, the best known time complexity bound is  $\tilde{O}(n^{t+2})$ , assuming we let  $t$  grow at least logarithmically in  $n$ . The other reason is that for posets whose cover graph is a tree ( $t = 1$ ), Atkinson’s [5] algorithm takes – at least seemingly – a different approach and runs in time  $\tilde{O}(n^3)$ . Furthermore, Atkinson’s algorithm is monotonic in the sense that all arithmetic operations are carried out with nonnegative numbers. This is in sharp contrast to both the present algorithm and the inclusion–exclusion algorithm, which crucially rely on a richer algebraic structure.

Our preliminary empirical study confirmed that the improvement in the asymptotic bound consistently transfers to the running times measured in practice. That said, the observed speedup, for  $n$  around one hundred, was by one order of magnitude rather than two. This “leak” of efficiency can, at least in part, be explained by the present algorithm’s higher sensitivity to the shape of the selected tree decomposition. Indeed, we observed that the best of 50 generated tree decompositions typically yields a 5- to 10-fold speedup in relation to a median tree decomposition. We also showed that there is an efficient way to select the best or close-to-best tree decomposition using a linear regression model that was fitted to a collected data set of instances along with the measured running times, following the method of Abseher et al. [2].

However, we observed that the performance of the regression method rapidly degraded as the treewidth  $t$  increases. This suggests that the general-purpose method may not suit well for the problem of counting linear extensions. A potential reason for suboptimal performance is that the default set of features [2] does not include perhaps the most informative quantity associated with a node  $x$  in a tree decomposition, namely the term  $n_x^k$  (or some variant of it), which combines the size  $k$  of the bag of  $x$  with the number of vertices in the subtree rooted at  $x$ . This issue could be addressed by extending the feature set accordingly, or, potentially, by using some nonlinear regression model. On the other hand, we did inspect how well a single feature can predict a well-performing tree decomposition. We observed that for  $t = 2$  and  $t = 3$  the average depth of join nodes alone yielded predictions that were almost as good as the predictions by the full regression model that used all the features. Nevertheless, these observations also encourage looking for a simpler solution: manually constructing a heuristic function that approximates the actual running time by adding up estimated contributions of each tree decomposition node. Note that here the  $f$ -cost framework [9] is insufficient, as in that framework the contribution of each node can only depend on the size of the bag.

---

## References

- 1 Michael Abseher, Nysret Musliu, and Stefan Woltran. htd – A Free, Open-Source Framework for (Customized) Tree Decompositions and Beyond. In *Proceedings of the 14th International Conference on Integration of AI and OR Techniques in Constraint Programming*, volume 10335 of *Lecture Notes in Computer Science*, pages 376–386. Springer, 2017.

- 2 Michael Abseher, Nysret Musliu, and Stefan Woltran. Improving the Efficiency of Dynamic Programming on Tree Decompositions via Machine Learning. *J. Artif. Intell. Res.*, 58:829–858, 2017.
- 3 Alfred V. Aho, M. R. Garey, and Jeffrey D. Ullman. The Transitive Reduction of a Directed Graph. *SIAM J. Comput.*, 1(2):131–137, 1972.
- 4 Stefan Arnborg, Derek G. Corneil, and Andrzej Proskurowski. Complexity of Finding Embeddings in a K-tree. *SIAM J. Algebra. Discr.*, 8(2):277–284, April 1987.
- 5 Mike D. Atkinson. On computing the number of linear extensions of a tree. *Order*, 7(1):23–25, 1990.
- 6 Jacqueline Banks, Scott Garrabrant, Mark L. Huber, and Anne Perizzolo. Using TPA to count linear extensions. *arXiv preprint arXiv:1010.4981*, 2017.
- 7 Umberto Bertelè and Francesco Brioschi. *Nonserial Dynamic Programming*. Academic Press, 1972.
- 8 Hans L. Bodlaender. A Linear-Time Algorithm for Finding Tree-Decompositions of Small Treewidth. *SIAM J. Comput.*, 25(6):1305–1317, 1996.
- 9 Hans L. Bodlaender and Fedor V. Fomin. Tree decompositions with small cost. *Discrete Appl. Math.*, 145(2):143–154, 2005.
- 10 Hans L. Bodlaender and Ton Kloks. Efficient and Constructive Algorithms for the Pathwidth and Treewidth of Graphs. *J. Algorithms*, 21(2):358–402, 1996.
- 11 Graham Brightwell and Peter Winkler. Counting linear extensions. *Order*, 8(3):225–242, 1991.
- 12 Rina Dechter. Bucket elimination: A unifying framework for reasoning. *Artif. Intell.*, 113(1–2):41–85, 1999.
- 13 Martin Dyer, Alan Frieze, and Ravi Kannan. A Random Polynomial-time Algorithm for Approximating the Volume of Convex Bodies. *J. ACM*, 38(1):1–17, 1991.
- 14 Eduard Eiben, Robert Ganian, Kustaa Kangas, and Sebastian Ordyniak. Counting Linear Extensions: Parameterizations by Treewidth. In *Proceedings of the 24th Annual European Symposium on Algorithms*, volume 57 of *LIPICs*, pages 39:1–39:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2016.
- 15 Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, and Ian H. Witten. The WEKA Data Mining Software: An Update. *SIGKDD Explor. Newsl.*, 11(1):10–18, November 2009.
- 16 William B. Hart. Fast Library for Number Theory: An Introduction. In *Proceedings of the Third International Congress on Mathematical Software*, pages 88–91, Berlin, Heidelberg, 2010. Springer-Verlag. URL: <http://flintlib.org>.
- 17 Kustaa Kangas, Teemu Hankala, Teppo Niinimäki, and Mikko Koivisto. Counting linear extensions of sparse posets. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*, pages 603–609. IJCAI/AAAI Press, 2016.
- 18 Richard M. Karp. Dynamic Programming Meets the Principle of Inclusion and Exclusion. *Oper. Res. Lett.*, 1(2):49–51, April 1982.
- 19 Mikko Koivisto. An  $O^*(2^n)$  Algorithm for Graph Coloring and Other Partitioning Problems via Inclusion–Exclusion. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science*, pages 583–590. IEEE Computer Society, 2006.
- 20 Kevin Leyton-Brown, Eugene Nudelman, and Yoav Shoham. Empirical hardness models: Methodology and a case study on combinatorial auctions. *J. ACM*, 56(4), 2009.
- 21 Thomas Lukasiewicz, Maria V. Martinez, and Gerardo I. Simari. Probabilistic Preference Logic Networks. In *Proceedings of the 21st European Conference on Artificial Intelligence*, volume 263 of *Frontiers in Artificial Intelligence and Applications*, pages 561–566. IOS Press, 2014.

- 22 Heikki Mannila and Christopher Meek. Global Partial Orders from Sequential Data. In *Proceedings of the Sixth International Conference on Knowledge Discovery and Data Mining*, pages 161–168. ACM, 2000.
- 23 Rolf H. Möhring. Computationally tractable classes of ordered sets. In I. Rival, editor, *Algorithms and Order*, pages 105–193. Kluwer Academic Publishers, 1989.
- 24 Jason Morton, Lior Pachter, Anne Shiu, Bernd Sturmfels, and Oliver Wienand. Convex Rank Tests and Semigraphoids. *SIAM J. Discrete Math.*, 23(3):1117–1134, 2009.
- 25 Christian J. Muise, J. Christopher Beck, and Sheila A. McIlraith. Optimal Partial-Order Plan Relaxation via MaxSAT. *J. Artif. Intell. Res.*, 57:113–149, 2016.
- 26 Teppo Niinimäki, Pekka Parviainen, and Mikko Koivisto. Structure Discovery in Bayesian Networks by Sampling Partial Orders. *J. Mach. Learn. Res.*, 17:57:1–57:47, 2016.
- 27 Marcin Peczarski. New Results in Minimum-Comparison Sorting. *Algorithmica*, 40(2):133–145, 2004.
- 28 John R. Rice. The algorithm selection problem. *Advances in Computers*, 15:65–118, 1976.
- 29 Herbert J. Ryser. *Combinatorial Mathematics*, volume 14 of *Carus Mathematical Monographs*. The Mathematical Association of America, 1963.
- 30 Topi Talvitie, Kustaa Kangas, Teppo Niinimäki, and Mikko Koivisto. Counting Linear Extensions in Practice: MCMC versus Exponential Monte Carlo. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*, 2018.
- 31 Joris van der Hoeven and Grégoire Lecerf. On the bit-complexity of sparse polynomial and series multiplication. *J. Symb. Comput.*, 50:227–254, 2013.
- 32 Chris S. Wallace, Kevin B. Korb, and Honghua Dai. Causal Discovery via MML. In *Proceedings of the 13th International Conference on Machine Learning*, pages 516–524. Morgan Kaufmann, 1996.