Counting linear extensions of restricted posets

Samuel Dittmer   Igor Pak

Department of Mathematics
UCLA
Los Angeles, CA, 90095, U.S.A.

`samuel.dittmer@gmail.com, pak@math.ucla.edu`

Abstract

We prove the 1991 conjecture by Brightwell and Winkler [BW91] that counting the number of linear extensions for *posets of height two* is #P-complete. We further extend this result to *incidence posets of graphs*.

Mathematics Subject Classifications: 06A07, 68Q17

1 Introduction

Counting *linear extensions* (#LE) of a finite poset is a fundamental problem in both Combinatorics and Computer Science. In 1991, Brightwell and Winkler showed that #LE is #P-complete [BW91]. They conjectured that the following problem is #P-complete:

> #H2LE (*Number of linear extensions of height-2 posets*)
> **Input:** A partially ordered set $P = (X, \prec)$ of height 2.
> **Output:** The number $e(P)$ of linear extensions.

Here *height two* means that P has two levels, i.e. no chains of length 3. This problem has been open for 27 years, most recently reiterated in [Hub14, LS17]. Its solution is the first result in this paper.

Theorem 1. #H2LE *is* #P-*complete.*

Our second result is an extension of Theorem 1. It was proposed recently by Lee and Skipper in [LS17], motivated by the optimization of nonlinear functions over the much-studied correlation polytope (see e.g. [DL97, LSS18]).

> #IPLE (*Number of linear extensions of incidence posets*)
> **Input:** A graph $G = (V, E)$.
> **Output:** The number $e(I_G)$ of linear extensions of the incidence poset I_G.

Here the incidence poset I_G is defined as a height 2 posets with vertices V on one level, edges E on another level, and the inequalities defined by adjacencies in G.

Theorem 2. #IPLE *is* #P-*complete.*

Theorem 2 implies Theorem 1, of course. Formally, the proofs of both results are independent, but use the same technical ideas of using number theory to to obtain targeted reductions modulo primes. Since the proof Theorem 1 is both technically and conceptually simpler, we chose to include both proofs.

Remark 3. The height-2 posets is an important and well studied class of posets. Brightwell and Winkler write: "We strongly suspect that Linear Extension Count for posets of height 2 is still #P-complete, but it seems that an entirely different construction is required to prove this" [BW91]. Incidence posets have also been studied quite intensely. We refer to recent papers [LS17, TW14] for an overview of the area and further references.

Notation

We assume the reader is familiar with basic definitions on posets, see e.g. [Tro95] and [Sta97, Ch. 3]. We describe a *linear extension* of a poset $\mathcal{P} = (X, \prec)$ on a set X with n elements informally as an *assignment* of the values $\{1, 2, \ldots, n\}$ to X, or formally as a function $\eta : X \to \{1, 2, \ldots, n\}$, s.t. $\eta(x) < \eta(y)$ for all $x \prec y$, $x, y \in X$. Finally, we refer to [MM11, Pap94] for notation, basic definitions and results in computational complexity.

2   Height two posets

Let $\mathcal{P} = (X, <)$ be a poset on a set X of n elements $\{x_1, \ldots, x_n\}$. Denote by $\Gamma = (X, E)$ its comparability graph, with oriented edges $(x_i, x_j) \in E$ if $x_i < x_j$ in $\mathcal{P}$. Denote by X' a identical copy of X with elements $\{x'_1, \ldots, x'_n\}$.

Define the poset $\mathcal{Q} = (X \cup X', \prec)$ on $2n$ elements, by having $x_i \prec x'_i$ for all $x_i \in X$, and $x_i \prec x'_j$ for all $x_i < x_j$, with $x_i, x_j \in X$. In particular, the Hasse diagram of $\mathcal{Q}$ consists of $n + |E|$ edges. Note that $\mathcal{Q}$ is a poset of height 2, see Figure 2.
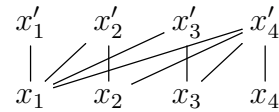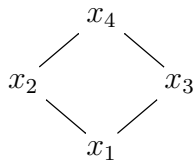

Figure 1: The Hasse diagram of a poset $\mathcal{P}$.

Figure 2: Poset $\mathcal{Q}$ associated to poset $\mathcal{P}$.

For every prime p between n and n^2, we construct the modified poset $\mathcal{Q}_p$ by adding, for all i and j satisfying $1 \leqslant i \leqslant n$ and $1 \leqslant j \leqslant p-2$, the element x_{ij} and the relation $x_{ij} \prec x'_i$. Note that $\mathcal{Q}_p$ is still of height 2 and has pn elements (see Figure 3).

We will use the numbers of linear extensions of posets $\mathcal{Q}$ and $\mathcal{Q}_p$, to compute the number of linear extensions of $\mathcal{P}$. Consider first the number $\boldsymbol{e}(\mathcal{Q})$ of linear extensions

of $\mathcal{Q}$. Let $A \in \binom{[2n]}{n}$, i.e. A is a n-subset of $[2n] = \{1, 2, \ldots, 2n\}$. Denote by $\boldsymbol{e}_A(\mathcal{Q})$ be the number of linear extensions ℓ of $\mathcal{Q}$ such that $\ell(X') = A$. In other words, $\boldsymbol{e}_A(\mathcal{Q})$ counts linear extensions that assign values in A to elements of $X' = \{x'_1, \ldots, x'_n\}$, defined as above. A linear extension ℓ of $\mathcal{Q}$ is counted in exactly one $\boldsymbol{e}_A(\mathcal{Q})$, so we have:

$$\boldsymbol{e}(\mathcal{Q}) = \sum_{A \in \binom{[2n]}{n}} \boldsymbol{e}_A(\mathcal{Q}).$$
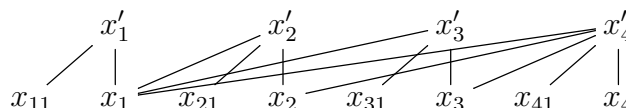
Figure 3: $\mathcal{Q}_p$ for $p = 3$.

Lemma 4. $\boldsymbol{e}(\mathcal{P}) = \boldsymbol{e}_{\{2,4,6,\ldots,2n\}}(\mathcal{Q})$.

Proof. We construct a bijection Φ explictly from $\boldsymbol{e}(\mathcal{P}) \to \boldsymbol{e}_{\{2,4,6,\ldots,2n\}}(\mathcal{Q})$, and give its inverse Ψ. First, given a linear extension ρ of $\mathcal{P}$, let $\Phi(\rho) \in \boldsymbol{e}_{\{2,4,6,\ldots,2n\}}(\mathcal{Q})$ be defined by

$$\Phi(\rho)[x_i] = 2\rho[x_i] - 1$$

and

$$\Phi(\rho)[x'_i] = 2\rho[x_i].$$

Note that $\Phi(\rho)[x_i] < \Phi(\rho)[x'_i]$ for all i, and if $x_i < x_j \in \mathcal{P}$, then

$$\Phi(\rho)[x_i] = 2\rho[x_i] - 1 < 2\rho[x_j] - 1 < 2\rho[x_j] = \Phi(\rho)[x'_j].$$

Thus $\Phi(\rho)$ is indeed a linear extension of $\mathcal{Q}$. Since $\Phi : \boldsymbol{e}(\mathcal{P}) \to \boldsymbol{e}_{\{2,4,6,\ldots,2n\}}(\mathcal{Q})$ is injective, we have $\boldsymbol{e}(\mathcal{P}) \leqslant \boldsymbol{e}_{\{2,4,6,\ldots,2n\}}(\mathcal{Q})$.

We next give the inverse map Ψ. Given a linear extension $\eta \in \boldsymbol{e}_{\{2,4,6,\ldots,2n\}}(\mathcal{Q})$ we construct a linear extension $\Psi(\eta) \in \boldsymbol{e}(\mathcal{P})$. For every $x_i \in \mathcal{P}$, we set

$$\Psi(\eta)[x_i] = \eta[x'_i]/2.$$

For every $\rho \in \boldsymbol{e}(\mathcal{P})$ we have $\Psi(\Phi(\rho)) = \rho$, by construction. To complete the proof, we must show that $\Phi(\Psi(\eta)) = \eta$ for every $\eta \in \boldsymbol{e}_{\{2,4,6,\ldots,2n\}}(\mathcal{Q})$, and that $\Psi(\eta) \in \boldsymbol{e}(\mathcal{P})$ for every $\eta \in \boldsymbol{e}_{\{2,4,6,\ldots,2n\}}(\mathcal{Q})$, that is, we must show that $\Psi(\eta)$ is a linear extension of $\mathcal{P}$.

Consider $\eta \in \boldsymbol{e}_{\{2,4,6,\ldots,2n\}}(\mathcal{Q})$. We will show that $\eta[x_i] = \eta[x'_i] - 1$ for all i with $1 \leqslant i \leqslant n$. Since

$$\eta[X'] = \{2, 4, 6, \ldots, 2n\}$$

we must have

$$\eta[X] = \{1, 3, 5, \ldots, 2n - 1\}.$$

For some $x_i' \in X'$ and for some $x_j \in X$, we have $\eta[x_i'] = 2n$ and $\eta[x_j] = 2n - 1$. For the relation $x_j \prec x_j'$ to be satisfied, we must have $\eta[x_j] < \eta[x_j']$. We must therefore have $x_j' = 2n$, that is, $i = j$ and $\eta[x_i] = \eta[x_i'] - 1$ for this value of i.

Suppose that for some m we have that $\eta[x_i] = \eta[x_i'] - 1$ for all x_i with $\eta[x_i] > 2m$. We have just proved this statement holds for the case $m = n - 1$. We proceed by induction on $n - m$. There exist $x_j' \in X'$ and $x_k \in X$ with $\eta[x_j'] = 2m$ and $\eta[x_k] = 2m - 1$. Note that we cannot have $\eta[x_k'] > 2m$, since then we would have $\eta[x_k] = \eta[x_k'] - 1 > 2m - 1$ by the induction hypothesis. But $\eta[x_k'] > \eta[x_k] = 2m - 1$, so that we must have $\eta[x_k'] = 2m$ and $\eta[x_j] = \eta[x_j'] - 1$. Thus by induction $\eta[x_i] = \eta[x_i'] - 1$ for all i with $1 \leqslant i \leqslant n$, as desired.

Applying this result, we have

$$\Phi(\Psi(\eta))[x_i'] = 2\Psi(\eta)[x_i] = \eta[x_i']$$

and

$$\Phi(\Psi(\eta))[x_i] = 2\Psi(\eta)[x_i] - 1 = \eta[x_i'] - 1 = \eta[x_i],$$

so that $\Phi(\Psi(\eta)) = \eta$, as desired.

Finally, for every pair $x_i < x_j \in \mathcal{P}$, we have $x_i \prec x_j' \in \mathcal{Q}$. Then $\eta[x_i] < \eta[x_j']$, so that $\eta[x_i] + 1 = \eta[x_i'] \leqslant \eta[x_j']$. Of course η is a bijection so we have $\eta[x_i'] < \eta[x_j']$. Thus $\Psi(\eta)[x_i] < \Psi(\eta)[x_j]$, and $\Psi(\eta)$ is a linear extension of $\mathcal{P}$. We conclude that Φ and Ψ are inverse maps, which completes the proof. $\qquad\square$

The above lemma should be compared with the following result:

Lemma 5. $\boldsymbol{e}(\mathcal{Q}_p) \equiv (-1)^n \boldsymbol{e}_{\{2,4,6,\ldots,2n\}}(\mathcal{Q}) \bmod p$.

Proof. Throughout the proof of this lemma, we will consider colorings of a set of integers. A *coloring* of a set is a function from that set to some list of acceptable colors.

Let $A \in \binom{[2n]}{n}$, and write $A = \{a_1, \ldots, a_n\}$, with $a_1 < a_2 < \cdots < a_n$. A coloring of the set $[pn] = \{1, 2, \ldots, pn\}$ is called *A-compatible* if the following conditions are satisfied:

1. there is a sequence of $2n$ integers $b_1 < \cdots < b_{2n}$ colored black,

2. there are another n colors $C_1, \ldots, C_n$, and $p - 2$ integers are colored with each of these colors,

3. all of the elements colored with C_k lie before b_{a_k}.

Let $f_p(A)$ be the number of A-compatible colorings of $[pn]$. Given a linear extension $\eta \in \boldsymbol{e}(\mathcal{Q})$, we write $\boldsymbol{e}_\eta(\mathcal{Q}_p)$ for the number of linear extensions of $\mathcal{Q}_p$ which preserve the ordering on $X \cup X'$ given by η. When $\eta \in \boldsymbol{e}_A(\mathcal{Q})$, we claim that

$$\boldsymbol{e}_\eta(\mathcal{Q}_p) = f_p(A)\big((p-2)!\big)^n.$$

Given a linear extension $\eta \in \boldsymbol{e}_A(\mathcal{Q})$ and a coloring in $f_p(A)$, we can construct $\big((p-2)!\big)^n$ linear extensions $\rho \in \boldsymbol{e}_\eta(\mathcal{Q}_p)$ as follows. Let

$$\rho[x_i] = b_{\eta[x_i]}$$

for all $x_i \in X$, and similarly let
$$\rho[x_i'] = b_{\eta[x_i']}$$
for all $x_i' \in X'$. Thus $\rho[x_i] < \rho[x_j]$ if and only if $\eta[x_i] < \eta[x_j]$. Note that for every integer k, with $1 \leqslant k \leqslant n$, there is some i with $\eta[x_i'] = a_k$ and $\rho[x_i'] = b_{a_k}$. For the $p-2$ elements $x_{ij} \prec x_i'$, assign to $\rho[x_{ij}]$ some permutation of the integers with color C_k. This gives $(p-2)!$ choices for each k, so the total number of linear extensions ρ preserving the ordering η for a fixed coloring is $\left((p-2)!\right)^n$, as desired. Reversing this procedure gives a linear extension for every choice of a linear extension $\eta \in \boldsymbol{e}_A(\mathcal{Q})$ and an A-compatible coloring.

We then have, by Wilson's theorem:
$$\boldsymbol{e}(\mathcal{Q}_p) = \sum_{A \in \binom{[2n]}{n}} \sum_{\eta \in \boldsymbol{e}_A(\mathcal{Q})} \boldsymbol{e}_\eta(\mathcal{Q}_p)$$
$$= ((p-2)!)^n \sum_{A \in \binom{[2n]}{n}} \boldsymbol{e}_A(\mathcal{Q}) f_p(A)$$
$$\equiv \sum_{A \in \binom{[2n]}{n}} \boldsymbol{e}_A(\mathcal{Q}) f_p(A) \bmod p.$$

In an A-compatible coloring of $\{1, 2, \ldots, pn\}$, there are $a_k - 1 + k(p-2)$ terms to the left of b_{a_k} colored either black or one of the colors $C_1, \ldots, C_k$. Among these terms, we can choose the position of the elements colored C_k arbitrarily. This gives
$$f_p(A) = \prod_{k=1}^n \binom{a_k - 1 + k(p-2)}{p-2}.$$

For $A = \{2, 4, 6, \ldots, 2n\}$, we have $a_k = 2k$, so this becomes
$$f_p\left(\{2, 4, 6, \ldots, 2n\}\right) = \prod_{k=1}^n \binom{kp-1}{p-2} \equiv (-1)^n \bmod p,$$

by Lucas's theorem. For every other A with $\boldsymbol{e}_A(\mathcal{Q}) \neq 0$, we have $f_p(A) \equiv 0 \bmod p$. Indeed, suppose $\boldsymbol{e}_A(\mathcal{Q}) \neq 0$ and consider some $\eta \in \boldsymbol{e}_A(\mathcal{Q})$. Then $\eta[x_i'] = 2n$ for some i, since $\eta[x_i] = 2n$ contradicts $x_i \prec x_i'$. Thus $a_n = 2n$. We proceed by induction on $n - k$. Suppose that
$$(a_{k+1}, \ldots, a_n) = (2k+2, \ldots, 2n).$$
Then for every integer $j > a_k$ with $j \notin (2k+2, \ldots, 2n)$, we have $\eta[x_i] = j$, for some $x_i \in X$. The relation $x_i \prec x_i'$ gives $\eta[x_i'] > \eta[x_i] > a_k$, so that $\eta[x_i'] \in (2k+2, \ldots, 2n)$. If we had $a_k < 2k$, we would then have at least $(n-k+2)$ possible values of $j > a_k$, but only $(n-k+1)$ possible values for $\eta[x_i']$ to take in $(2k+2, \ldots, 2n)$. Thus either $a_k = 2k$ or $a_k = 2k+1$.

If $a_k = 2k+1$, then
$$\binom{a_k - 1 + kp - 2k}{p-2} = \binom{kp}{p-2}$$

will divide $f_p(A)$. By another application of Lucas's theorem, $\binom{kp}{p-2} \equiv 0 \mod p$, so we have $f_p(A) \equiv 0 \mod p$ unless $a_k = 2k$. This completes the induction, and we conclude $\boldsymbol{e}_A(\mathcal{P}) f_p(A) \equiv 0 \mod p$ unless $A = \{2, 4, \ldots, 2n\}$. $\qquad\square$

We need the following number theory result:

Proposition 6 (see e.g. [BW91, p. 4]). *For $k \geqslant 4$, the product of primes between k and k^2 is at least $2^k k!$.*

Proof of Theorem 1. We make an argument based on the Chinese Remainder Theorem similar to that in [BW91]. Since $\boldsymbol{e}(\mathcal{P}) \leqslant n!$, the proposition above together with the Chinese Remainder Theorem shows that computing the residue of $\boldsymbol{e}(\mathcal{P}) \mod p$ for the primes p with $n \leqslant p \leqslant n^2$ is sufficient to determine $\boldsymbol{e}(\mathcal{P})$. The lemmas above show that we can compute the residue of $\boldsymbol{e}(\mathcal{P}) \mod p$ by computing $\boldsymbol{e}(\mathcal{Q}_p)$. Since #LE is #P-complete, so is #H2LE. $\qquad\square$

3 Incidence posets

3.1 Counting incidence posets

Given a graph $G = (V, E)$, we construct its incidence poset I_G, with elements corresponding to vertices *and* edges of G, with $x < y$ in $\mathcal{P}$ if and only if $x \in E$, $y \in V$ and y is an endpoint of x. We write $\boldsymbol{e}(G)$ for the number of linear extensions of I_G.

Our approach here is similar to our approach in Section 2. We produce, given a poset $\mathcal{P}$ and a prime $p > |\mathcal{P}|$, a graph $G_p(\mathcal{P})$ with:

$$\boldsymbol{e}\left(G_p(\mathcal{P})\right) \equiv (-1)^{|\mathcal{P}|} \cdot 8\,\boldsymbol{e}(\mathcal{P}) \mod p.$$

Let $G = (V, E)$ be a graph, with $V = \{x_1, \ldots, x_n\}$, and $\sigma \in S_n$ a permutation. Denote by $\boldsymbol{e}_\sigma(G)$ the number of linear extensions of I_G, which satisfy the following condition: when restricted to V, induce the permutation σ, so that $x_{\sigma^{-1}(1)} \leqslant x_{\sigma^{-1}(2)} \leqslant \cdots \leqslant x_{\sigma^{-1}(n)}$. We have:

$$\boldsymbol{e}(G) = \sum_{\sigma \in S_n} \boldsymbol{e}_\sigma(G).$$

Informally, to compute $\boldsymbol{e}_\sigma(G)$ we visit the vertices of G in the order dictated by σ, accounting for the new edges we meet at each step.

Formally, given a permutation $\sigma \in S_n$, we produce the sequence $\{t_1, \ldots, t_n\}$, where t_i is the number of edges in E with $x_{\sigma^{-1}(i)}$ as an endpoint, and no endpoint $x_{\sigma^{-1}(j)}$ for $j < i$. Let $\{u_1, \ldots, u_n\}$ be the sequence of partial sums of the t_i's, so that

$$u_k = t_1 + \ldots + t_k.$$

Note that u_k is the total number of edges incident to the set of vertices $x_{\sigma^{-1}(1)}, \ldots, x_{\sigma^{-1}(k)}$.

Let $|E| = m$. Then we call a coloring of the set $\{1, 2, \ldots, m+n\}$ (G, σ)-compatible if the following conditions are satisfied:

1. there is a sequence of n integers $b_1 < \cdots < b_n$ colored black,

2. there are another n colors $C_1, \ldots, C_n$, and t_k integers are colored with the color C_k,

3. all of the elements colored with C_k lie before b_k.

Let $f(G, \sigma)$ be the number of (G, σ)-compatible colorings. In such a coloring, there are $u_k + k - 1$ numbers to the left of b_k colored either black or one of the colors $C_1, \ldots, C_k$. Among these terms, we can choose the position of the elements colored C_k arbitrarily. This gives:

$$f(G, \sigma) = \prod_{k=1}^{n} \binom{u_k + k - 1}{t_k}.$$

A (G, σ)-compatible coloring corresponds to a collection of linear extensions of I_G counted by $e_\sigma(G)$. The values assigned to the t_k new edges at $x_{\sigma^{-1}(k)}$ are given by the numbers colored with C_k, and these values can be assigned in $(t_k)!$ ways, so that we have:

$$e(G) = \sum_{\sigma \in S_n} f(G, \sigma) \prod_{k=1}^{n} (t_k)! = \sum_{\sigma \in S_n} \prod_{k=1}^{n} (t_k)! \binom{u_k + k - 1}{t_k}. \tag{1}$$

In particular, when we are counting modulo p we can restrict our attention to permutations σ, which have corresponding sequences $\{t_1, \ldots, t_n\}$ with $t_i < p$ for all i. Informally, we want to visit each vertex of G in the order given by σ, deleting the edges incident to each vertex after we visit it, and ensure that no vertex has at least p edges by the time we visit it.

Now we give the actual construction of $G_p(\mathcal{P})$. The first step is to construct a gadget J_p, which is a graph defined as follows. Start with the complete bipartite graph $K_{p-1,p-1}$ on $2p - 2$ vertices. Call these vertices $y_1, \ldots, y_{p-1}$ and $z_1, \ldots, z_{p-1}$ and add an additional $p - 2$ edges from z_{p-1} to z_i for $1 \leqslant i < p - 1$. Note that each of the y_i's has degree $p - 1$ and the z_i's have degree $\geqslant p$ (see Figure 4). We need the following:

Lemma 7. $e(J_p) \equiv -8 \mod p$.

We defer the proof of this lemma to the end of this section.

To construct $G_p(\mathcal{P})$, add below J_p the Hasse diagram of $\mathcal{P}$ (treated as an undirected graph). For each element $x \in \mathcal{P}$, let v_x be the number of elements in $\mathcal{P}$ that cover x. Add $p - 1 - v_x$ edges from x to the degree $p - 1$ vertices y_i of J_p in an arbitrarily way (see Figure 5).

Theorem 2 follows immediately from the following:

Lemma 8. $e(G_p(\mathcal{P})) \equiv (-1)^{|\mathcal{P}|+1} \cdot 8\,e(\mathcal{P}) \mod p$

Proof. Every maximal element of $\mathcal{P}$ has $v_x = 0$, and so is connected to each of the y_i's in J_p. Since $\mathcal{P}$ has at least one maximal element, every element of J_p has degree $\geqslant p$. Thus every σ which visits a vertex in J_p before visiting every maximal element of $\mathcal{P}$ has
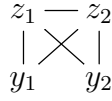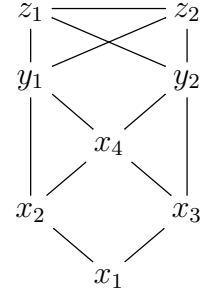
Figure 4: J_p for $p = 3$.

Figure 5: $G_p(\mathcal{P})$ for $\mathcal{P}$ as in Figure 1 and $p = 3$.

a term $t_i \geqslant p$, so that $\boldsymbol{e}_\sigma(G_p(\mathcal{P})) \equiv 0 \mod p$. Likewise, of these permutations, every permutation σ that visits an element of $\mathcal{P}$ before visiting all of its immediate predecessors has $\boldsymbol{e}_\sigma(G_p(\mathcal{P})) \equiv 0 \mod p$.

Thus we can restrict our count of $\boldsymbol{e}(G_p(\mathcal{P}))$ modulo p to permutations that have as their first n terms a linear extension of $\mathcal{P}$. For these permutations, we have $t_1 = t_2 = \ldots = t_n = p - 1$, so that $(t_k)! \equiv -1 \mod p$ by Wilson's theorem, and

$$\binom{u_k + k - 1}{t_k} = \binom{kp - 1}{p - 1} \equiv 1 \mod p.$$

Furthermore, for every $k > n$, we have $t_1 + \ldots + t_k = np - n + (t_{n+1} + \cdots + t_k) + k - 1$, so that

$$\binom{u_k + k - 1}{t_k} \equiv \binom{u_k - u_n + (k - n) - 1}{t_k} \mod p.$$

Now comparing the expressions for $\boldsymbol{e}(G_p(\mathcal{P}))$ and $\boldsymbol{e}(J_p)$ given by (1), we have

$$\boldsymbol{e}\left(G_p(\mathcal{P})\right) \equiv (-1)^{|\mathcal{P}|}\boldsymbol{e}(\mathcal{P})\boldsymbol{e}(J_p) \mod p,$$

and Lemma 7 completes the proof. $\qquad\square$

Proof of Theorem 2. Using the same Chinese Remainder Theorem argument we used in Section 2, the two lemmas above show that computing $\boldsymbol{e}(G_p(P))$ for the primes between $|\mathcal{P}|$ and $|\mathcal{P}|^2$ is sufficient to determine $\boldsymbol{e}(\mathcal{P})$. Since #LE is #P-Complete, so is #IPLE. $\qquad\square$

3.2   Proof of Lemma 7

Note that the values t_k and $u_k + k - 1$ in (1) are both independent of the order in which the previous $k - 1$ vertices are visited. They can be computed solely by identifying the vertex $x_{\sigma^{-1}(k)}$ and the collection of vertices $\{x_{\sigma^{-1}(i)}\}_{i<k}$. This motivates the following
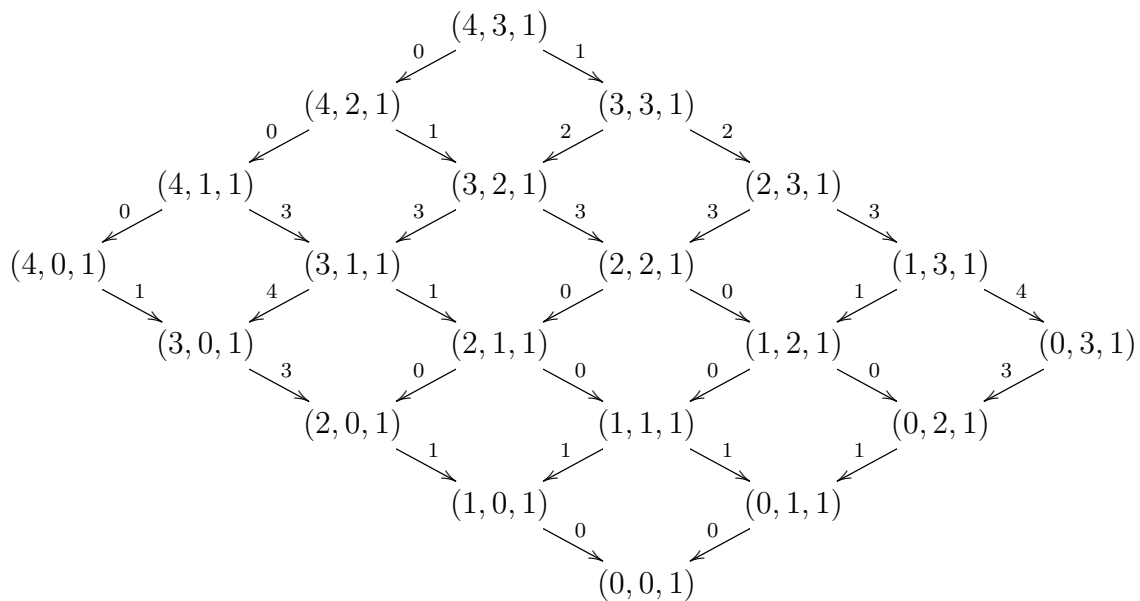
Figure 6: The $c = 1$ half of the directed graph $\mathcal{G}'$, with weights, for $p = 5$.

construction. Recall that the induced subgraphs of a graph G are those formed by deleting some vertices together with all incident edges. Take a directed graph $\mathcal{G}$ whose vertices are the induced subgraphs of J_p and whose edges point from each subgraph to those obtained from it by deleting a single vertex. Attach to each edge the weight

$$(t_k)! \binom{u_k + k - 1}{t_k} = (t_k)! \binom{u_k + k - 1}{u_k - u_{k-1}}. \tag{2}$$

Then $\boldsymbol{e}(J_p)$ is equal to the sum of all weighted paths in $\mathcal{G}$ from J_p to the empty subgraph.

Let $J_p(a, b, c)$ be an induced subgraph of J_p with a of the y_i's, b of the z_i's, for $1 \leqslant i < p-1$, and $c = 1$ if $z_{p-1} \in J_p(a, b, c)$, $c = 0$ otherwise, for $0 \leqslant a \leqslant p-1$ and $0 \leqslant b \leqslant p-2$. Since the y_i's, and the z_i's, except for z_{p-1}, are indistinguishable, these subgraphs $J_p(a, b, c)$ are *all* of the induced subgraphs of J_p, up to isomorphism.

We can thus reduce our graph of subgraphs $\mathcal{G}$ to the graph $\mathcal{G}'$ containing only these $2p^2 - 2p$ vertices. We re-weight the edges from $J_p(a, b, c)$ where a, b or c is reduced by one, by multiplying by a, b or c, respectively. This accounts for the a, b or c choices of vertex to remove. Write $\ell(a, b, c)$ for the value of $u_{k-1} + k - 1$ upon reaching $J_p(a, b, c)$, that is, $\ell(a, b, c)$ is the number of vertices and edges that must be deleted from $J_p(p-1, p-2, 1)$ to give $J_p(a, b, c)$. Then (2) gives the weight of the edge from $J_p(a, b, c)$ to $J_p(a-1, b, c)$ in terms of a, b, c and ℓ:

$$a(b+c)! \binom{\ell(a, b, c) + b + c}{b + c} = a(b+c)! \binom{\ell(a-1, b, c) - 1}{\ell(a-1, b, c) - \ell(a, b, c) - 1}. \tag{3}$$

The equations for the edges from $J_p(a, b, c)$ to $J_p(a, b-1, c)$ and $J_p(a, b, c-1)$ are the same up to a cyclic permutation of (a, b, c). The total number of edges in J_p is $(p-1)^2 + (p-2) = p^2 - p - 1$. The number of edges in $J_p(a, b, c)$ is $ab + ac + bc$, and we reach $J_p(a, b, c)$ by deleting $(p - 1 - a) + (p - 2 - b) + (1 - c)$ vertices. We then calculate:

$$\ell(a, b, c) = p^2 - p - 1 - (ab + (a + b)c) + (p - 1 - a) + (p - 2 - b) + (1 - c)$$
$$\equiv (a + 2)(p - b - 2) + (c - 1)(a + b + 2) \mod p.$$

Lemma 9. *When $c = 1$, $(a + 2)(p - b - 2) > p$ and $(p - a - 2)(b + 2) > p$, every path in $\mathcal{G}'$ that visits $J_p(a, b, c)$ has weight zero modulo p.*

Proof. We argue by induction on $(2p - 3) - (a + b)$, that is, on the distance in $\mathcal{G}'$ from $J_p(p - 1, p - 2, 1)$ to $J_p(a, b, c)$. When $a = p - 1$, $b = p - 2$, $c = 1$, the conditions of the lemma are not met, and the statement is true vacuously.

Now suppose that a, b, c satisfy the conditions in this lemma. Then a path that visits $J_p(a, b, c)$ must come from either $J_p(a + 1, b, c)$ or $J_p(a, b + 1, c)$. If the values $a + 1, b, c$ satisfy the conditions in this lemma, we can then apply the induction hypothesis to show that every path through $J_p(a + 1, b, c)$ has weight 0 modulo p. In particular, a path that includes the edge from $J_p(a + 1, b, c)$ to $J_p(a, b, c)$ has weight 0 modulo p.

On the other hand, suppose that $a + 1, b, c$ do not satisfy the conditions in this lemma. Then $(a + 3)(p - b - 2) > (a + 2)(p - b - 2) > p$, so we must have $(p - a - 3)(b + 2) \leqslant p$. Note that if a or b is greater than or equal to $p - 2$, either $(a + 2)(p - b - 2) \leqslant 0$ or $(p - a - 2)(b + 2) \leqslant 0$. We thus have $a, b < p - 2$, so that $(p - a - 3)(b + 2) = p$ is impossible.

However, when $(p - a - 3)(b + 2) < p$, since $b < p - 2$, we have $(p - a - 3)(b - 2) >_p (p - a - 2)(b + 2)$. Thus, $\ell(a + 1, b, c) >_p \ell(a, b, c)$, and so by (3), the edge from $J_p(a + 1, b, c)$ to $J_p(a, b, c)$ has weight 0 modulo p. The argument for the edge from $J_p(a, b + 1, c)$ to $J_p(a, b, c)$ is the same by symmetry. $\square$

Lemma 10. *Given a, b with $(b + 2)(p - a - 2) \leqslant p$ the edge from $J_p(a, b, 1)$ to $J_p(a, b, 0)$ has weight 0 unless $a = p - 3$ and $b = 0$, $a = p - 2$ and $b = 0$ or 1, or $a = p - 1$ with b arbitrary. Similarly, given a, b with $(a + 2)(p - b - 2) \leqslant p$, the edge from $J_p(a, b, 1)$ to $J_p(a, b, 0)$ has weight 0 unless $b = p - 3$ and $a = 0$, or $b = p - 2$ and $a = 0$ or 1.*

Proof. We give the proof of the first statement, since the proof of the second is essentially identical. Permuting (a, b, c) in (3) to find the weight of the edge from $J_p(a, b, 1)$ to $J_p(a, b, 0)$, we note that we must have $a + b < p$ and $a + b <_p a + b + \ell(a, b, 1)$. Since $\ell(a, b, 1) \equiv (b + 2)(p - a - 2) \mod p$, this gives:

$$a + b + (b + 2)(p - a - 2) < p.$$

This implies that

$$p < a + 1 + \frac{3}{b + 1} \leqslant a + 4.$$

We conclude that $a > p - 4$, and the rest of the lemma follows by elementary case analysis. $\square$

Proof of Lemma 7. Note that the edges from $J_p(p-1, p-2, 1)$ to $J_p(p-1, p-3, 1)$ and $J_p(p-1, p-2, 0)$ have weight 0 modulo p. Combining this with the previous two lemmas, we conclude that every path in $\mathcal{G}'$ has weight 0 modulo p unless it visits either $J_p(p-2, 1, 1)$ or $J_p(1, p-2, 1)$. We now complete the desired calculation, through repeated applications of (3), symmetry, and Wilson's theorem:

$$
\begin{aligned}
\boldsymbol{e}(J_p(p-1, p-2, 1)) &\equiv (p-1)(p-1)! \, \boldsymbol{e}\big(J_p(p-2, p-2, 1)\big) \\
&\equiv (p-2)! \, (-1)^{p-3} \Big[\boldsymbol{e}(J_p(p-2, 1, 1)) + \boldsymbol{e}(J_p(p-2, 0, 1)) \Big] \\
&\equiv 2 \, \boldsymbol{e}\big(J_p(p-2, 1, 1)\big) \\
&\equiv 2 \, (p-1)! \Big[\boldsymbol{e}\big(J_p(p-2, 1, 0)\big) + \boldsymbol{e}\big(J_p(p-2, 0, 1)\big) \Big] \\
&\equiv -4 \, \boldsymbol{e}\big(J_p(p-2, 0, 1)\big) \\
&\equiv -4 \, (p-2)! \, \boldsymbol{e}\big(J_p(p-2, 0, 0)\big) - 4 \, (p-2) \, \boldsymbol{e}\big(J_p(p-3, 0, 1)\big) \\
&\equiv -4 \, \boldsymbol{e}\big(J_p(p-2, 0, 0)\big) + 8 \binom{p-1}{2} \boldsymbol{e}\big(J_p(p-3, 0, 0)\big) \\
&\equiv -4 \, (p-2)! + 4 \, (p-1)(p-2)(p-3)! \\
&\equiv -8 \mod p.
\end{aligned}
$$

This completes the proof. $\qquad\square$

4 Polytope of modes

Motivated by probabilistic applications, Montúfar and Rauh [MR16] recently defined the *polytope of modes* $\mathbf{M}(G, X)$, for every simple graph $G = (V, E)$ and independent subset of vertices $X \subset V$. The polytope $\mathbf{M}(G, X)$ consists of all functions $p : V \to [0, 1]$ satisfying

$$
\sum_{v \in V} p(v) = 1
$$

and

$$
p(x) \geqslant p(y)
$$

for every pair (x, y) with $x \in X$ and $(x, y) \in E$. From the perspective of probability, the functions p are probability distributions, and the points x are *modes* of the distribution. Montúfar and Rauh proved that

$$
\operatorname{vol} \mathbf{M}(G, X) = \frac{\operatorname{vol}(\Delta^n)}{n!} \, \boldsymbol{e}\big(P_{G,X}\big),
$$

where $n = |V|$, $\operatorname{vol}(\Delta^n) = \sqrt{n}/(n-1)!$, and $P_{G,X}$ is a poset constructed from G and X [MR16, Prop. 3] (see also [Sta97] for a strongly related *order polytope*). The poset $P_{G,X}$ is formed by taking the elements of V together with the relation $x < y$ for every pair (x, y) with $x \in X$ and $(x, y) \in E$. This poset has height 2, with vertices in X

on one level and $V \setminus X$ on the other. The authors then discuss the problem of computing $e(P_{G,X})$.

The following result follows easily from our Theorem 2. Curiously, we learned about this problem after the paper had been written.

Proposition 11. *For every incidence poset I_G of a simple graph $G = (V, E)$, there exists some graph H and some independent set $X \subseteq H$ with $P_{H,X} = I_G$.*

Proof. The desired graph H is the *medial graph* defined as a graph on the set of vertices $V \cup E$. The edges of this graph are pairs (v, e) where v is incident to e in G. This graph is bipartite, so V is an independent set, and we take $X = V$. Now we note that the poset $P_{H,X}$ consists of the set $V \cup E$ together with the relation $v < e$ whenever v is incident to e, so that $P_{H,X} = I_G$, as desired. $\square$

Corollary 12. *The problem of computing $e(P_{G,X})$ is #P-complete.*

Proof. By the proposition, computing $e(P_{H,X})$ allows us to compute $e(I_G)$ for any graph G. Applying Theorem 2 completes the proof. $\square$

References

[BjW91] A. Björner and M. Wachs, Permutation statistics and linear extensions of posets, *J. Combin. Theory* **A 58** (1991), 85–114.

[BW91] G. Brightwell and P. Winkler, Counting linear extensions, *Order* **8** (1991), 225–247; extended abstract in *Proc. 23rd STOC* (1991), 175–181.

[DL97] M. Deza and M. Laurent, *Geometry of cuts and metrics*, Springer, Berlin, 1997.

[Hub14] M. Huber, Near-linear time simulation of linear extensions of a height-2 poset with bounded interaction, *Chicago J. Theoret. Comput. Sci.* (2014), Art. 3, 16 pp.

[LS17] J. Lee and D. Skipper, Volume computations for sparse boolean quadric relaxations; arXiv:1703.02444.

[LSS18] J. Lee, D. Skipper and E. Speakman, Algorithmic and modeling insights via volumetric comparison of polyhedral relaxations, *Math. Program.* **170** (2018), no. 1, Ser. B, 121–140.

[MR16] G. Montúfar and J. Rauh, Mode poset probability polytopes, *J. Algebr. Stat.* **7** (2016), 1–13.

[MM11]   C. Moore and S. Mertens, *The nature of computation*, Oxford Univ. Press, Oxford, 2011.

[Pap94]   C. H. Papadimitriou, *Computational Complexity*, Addison-Wesley, Reading, MA, 1994.

[Sta97]   R. P. Stanley, *Enumerative Combinatorics*, Vol. 1 and 2, Cambridge Univ. Press, Cambridge, MA, 1997 and 1999.

[Tro95]   W. T. Trotter, Partially ordered sets, in *Handbook of combinatorics*, Vol. 1, Elsevier, Amsterdam, 1995, 433–480.

[TW14]   W. T. Trotter and R. Wang, Incidence posets and cover graphs, *Order* **31** (2014), 279–287.