Combinatorics, Probability and Computing (2001) 10, 159–177. Printed in the United Kingdom

Counting Partial Orders with a Fixed Number of Comparable Pairs

HANS JÜRGEN PRÖMEL,[1]† ANGELIKA STEGER[2]
and ANUSCH TARAZ[1]

[1] Institut für Informatik, Humboldt–Universität zu Berlin,
10099 Berlin, Germany
(e-mail: proemel@informatik.hu-berlin.de,
taraz@informatik.hu-berlin.de)

[2] Institut für Informatik, Technische Universität München,
80290 München, Germany
(e-mail: steger@in.tum.de)

Received 17 May 1999; revised 3 March 2000

In 1978, Dhar suggested a model of a lattice gas whose states are partial orders. In this context he raised the question of determining the number of partial orders with a fixed number of comparable pairs. Dhar conjectured that in order to find a good approximation to this number, it should suffice to enumerate families of layer posets. In this paper we prove this conjecture and thereby prepare the ground for a complete answer to the question.

1. Introduction and results

Let $\mathscr{P}_n$ be the set of all labelled partial orders with point set $[n] = \{1, \dots, n\}$. A trivial lower bound on $|\mathscr{P}_n|$ is given by

$$|\mathscr{P}_n| \geqslant 2^{\frac{n^2}{4}},$$

since we can fix two antichains X and Y, each on $n/2$ points, and decide independently for each of the $n^2/4$ pairs $(x, y) \in X \times Y$ whether or not $x < y$ should hold.

Upper bounds are much harder to obtain. In 1970, Kleitman and Rothschild [3] first gave the following bound:

$$|\mathscr{P}_n| \leqslant 2^{n^2/4+O(n^{3/2}\log_2 n)}. \tag{1.1}$$

† Research supported in part by DFG-project Pr 296/4-2.

A few years later [4], they were able to compute the exponent much more precisely:

$$\log_2 |\mathscr{P}_n| = \frac{n^2}{4} + \frac{3n}{2} + O(\log_2 n). \tag{1.2}$$

The underlying principle of the proofs of these results can be stated in rough terms as follows. Find a subclass $\mathscr{Q}_n \subseteq \mathscr{P}_n$ that on the one hand has a nice structure and can therefore be enumerated easily. On the other hand it should be so large that $|\mathscr{Q}_n|$ is a good approximation for $|\mathscr{P}_n|$. Here Kleitman and Rothschild chose $\mathscr{Q}_n$ so that it contained only 3-*layer posets* – these are posets whose point set can be partitioned into three antichains X_1, X_2, X_3 such that no point in X_1 is above any element of X_2, no point in X_2 is above any element of X_3, and every point in X_1 is below every point in X_3. One of the particularly appealing features of this technique is that it also proves that the proportion of posets in $\mathscr{P}_n$ that are 3-layer posets tends to one as n tends to infinity – in other words, *almost all posets are 3-layer posets* [4].

The central purpose of this paper is the investigation of the *number of partial orders with a fixed number of comparable pairs*. More precisely, for $0 < d < \frac{1}{2}$ denote by $\mathscr{P}_{n,d}$ those posets in $\mathscr{P}_n$ with $[dn^2]$ comparable pairs (where $[dn^2]$ denotes the nearest integer to dn^2) and let

$$\mathrm{c}(d) := \lim_{n\to\infty} \frac{\log_2 |\mathscr{P}_{n,d}|}{n^2}, \quad \text{in other words,} \quad |\mathscr{P}_{n,d}| = 2^{\mathrm{c}(d)n^2 + o(n^2)},$$

provided the limit exists. Recall from (1.1) that, for any d,

$$\mathrm{c}(d) \leqslant \frac{1}{4}. \tag{1.3}$$

In 1978, Dhar [1] raised the question of determining $\mathrm{c}(d)$ and suggested that partial orders can represent the states of a certain model of lattice gas with energy proportional to the number of comparable pairs in the order. In this context, $\mathrm{c}(d)$ would correspond to the entropy function of the lattice gas.

Results due to Dhar [1, 2] as well as Kleitman and Rothschild [5] show that, in the whole range $0 < d < \frac{1}{2}$, the function $\mathrm{c}(d)$ is continuous and that

$$\text{for } 0 < d \leqslant \frac{1}{8}, \qquad \mathrm{c}(d) = \frac{1}{4} \cdot H(4 \cdot d), \tag{1.4}$$

$$\text{for } \frac{1}{8} \leqslant d \leqslant \frac{3}{16}, \qquad \mathrm{c}(d) \equiv \frac{1}{4}, \tag{1.5}$$

where $H(x) = -x \log_2 x - (1-x)\log_2(1-x)$. The problem has remained open for larger values of d. Here Dhar conjectured that for each d there is a family of k-layer posets that is large enough to 'dominate' the set $\mathscr{P}_{n,d}$ and thus determine $\mathrm{c}(d)$ (see below for the formal definitions). In other words, this family would have a significance for $\mathscr{P}_{n,d}$ similar to the one that the 3-layer posets had for $\mathscr{P}_n$. The aim of this paper is to prove this conjecture and thereby prepare the ground for a complete solution of the problem.

We first extend the definition of a 3-layer poset to a k-layer poset in a natural way. A poset $P = (X, P)$ is a k-layer poset, if there exists a partition of its point set $X = X_1 \cup \ldots \cup X_k$

into k disjoint antichains (the so-called *layers*) such that

$$x < y \text{ with } x \in X_i \text{ and } y \in X_j \quad \Longrightarrow \quad i < j,$$
$$\text{for every } i, j \text{ with } j > i+1: \quad x \in X_i,\ y \in X_j \quad \Longrightarrow \quad x < y.$$

For some constants $\lambda_1, \ldots, \lambda_k$ with $0 < \lambda_i < 1$ and $\sum_i \lambda_i = 1$ and a constant $0 \leqslant p \leqslant 1$, we say that a poset $P \in \mathscr{P}_n$ has *configuration* $Q = (\lambda_1, \ldots, \lambda_k; p)$, if it belongs to the set $\mathscr{P}_{n,Q} \subseteq \mathscr{P}_n$, which is defined as the set containing all k-layer posets in $\mathscr{P}_n$ that have $p\,|X_i||X_{i+1}|$ comparable pairs between X_i and X_{i+1} (for all $i \in [n-1]$) and satisfy $|X_i| = \lambda_i n$ (for all $i \in [n]$). For the sake of a more legible introduction, let us assume for now that all the real numbers $p\,|X_i||X_{i+1}|$ and $\lambda_i n$ happen to be integers. Of course, we will need to fix this inaccuracy (and will do so at the end of this introduction), but given that we are only aiming at a very rough approximation of $|\mathscr{P}_{n,d}|$, namely, the coefficient $\mathrm{c}(d)$ of the leading term in the logarithm, it should be clear that this is by no means critical.

Obviously, any two posets P and P' with the same configuration must have the same number of comparable pairs, which means that for every Q there exists a d such that, *for every* n,

$$\mathscr{P}_{n,Q} \subset \mathscr{P}_{n,d}. \tag{1.6}$$

The main result of this paper states that, on the other hand, for each d we can find a configuration Q such that (1.6) holds almost with equality and thereby proves the conjecture of Dhar mentioned above.

Theorem 1.1. *For every* $0 < d < \frac{1}{2}$ *there exists a configuration* $Q = (\lambda_1, \ldots, \lambda_k;\ p)$ *with* $\mathscr{P}_{n,Q} \subset \mathscr{P}_{n,d}$ *such that*

$$\lim_{n\to\infty} \frac{\log_2 |\mathscr{P}_{n,Q}|}{n^2} = \lim_{n\to\infty} \frac{\log_2 |\mathscr{P}_{n,d}|}{n^2} = \mathrm{c}(d),$$

in other words

$$2^{o(n^2)}\,|\mathscr{P}_{n,Q}| = |\mathscr{P}_{n,d}|.$$

The following two observations will be helpful when it comes to actually constructing the configuration Q mentioned in the theorem.

Lemma 1.2. *For* $\frac{1}{8} \leqslant d \leqslant \frac{1}{2}$, *the configuration* $Q = (\lambda_1, \ldots, \lambda_k;\ p)$ *must be chosen so that* $p \geqslant \frac{1}{2}$.

Kleitman and Rothschild [5] observed that

$$\text{for } 0 < d \leqslant \frac{1}{8}: \quad Q = \left(\frac{1}{2}, \frac{1}{2};\ 4d\right) \tag{1.7}$$

satisfies the requirements of Theorem 1.1. We shall re-prove this statement when we prove Theorem 1.1 for $0 < d \leqslant \frac{1}{8}$. Our methods used to prove the above results do not seem strong enough to give results about *almost all posets in* $\mathscr{P}_{n,d}$. We conjecture, however, that indeed almost all posets in $\mathscr{P}_{n,d}$ lie in some $\mathscr{P}_{n,Q}$.

Before we come to the proofs, we indicate how Theorem 1.1 can be used to compute $\mathrm{c}(d)$. Consider an arbitrary configuration $Q = (\lambda_1, \ldots, \lambda_k; p)$. Clearly, the number of comparable pairs in every poset in $\mathscr{P}_{n,Q}$ is given by $\mathrm{d}(Q)n^2$, where

$$\mathrm{d}(Q) := p \sum_{i=1}^{k-1} \lambda_i \lambda_{i+1} + \sum_{i=1}^{k-2} \sum_{j \geqslant i+2}^{k} \lambda_i \lambda_j. \tag{1.8}$$

On the other hand, the only degree of freedom one has when constructing a poset in $\mathscr{P}_{n,Q}$ lies in the placement of the comparable pairs between successive layers. Thus we let $\mathrm{c}(Q) := H(p) \sum_{i=1}^{k-1} \lambda_i \lambda_{i+1}$ and arrive at the following estimate for $|\mathscr{P}_{n,Q}|$:

$$\prod_{i=1}^{k-1} \binom{\lambda_i \lambda_{i+1} n^2}{p \cdot \lambda_i \lambda_{i+1} n^2} = 2^{\sum_{i=1}^{k-1} H(p) \lambda_i \lambda_{i+1} n^2 \; + o(n^2)} = 2^{\mathrm{c}(Q)n^2 + o(n^2)}. \tag{1.9}$$

(Actually we did have more freedom: since we are considering labelled posets we also had the choice of assigning points to classes. But this merely gives a factor of $O(n!) = O(2^{n \log_2 n})$.)

This now puts us in the following position: in order to determine $\mathrm{c}(d)$ for some fixed d, it suffices to determine the configuration Q whose existence is proved in Theorem 1.1, since we then know that $\mathrm{c}(d) = \mathrm{c}(Q)$. To find Q, one can use the fact that there cannot be another configuration Q' with $\mathrm{d}(Q') = \mathrm{d}(Q)$ and $\mathrm{c}(Q') > \mathrm{c}(Q)$. Hence Q must be the solution to the following maximization problem:

$$\text{Choose } k,\ \lambda_1, \ldots, \lambda_k, \text{ and } p \text{ such as to maximize } \ H(p) \sum_{i=1}^{k-1} \lambda_i \lambda_{i+1},$$

$$\text{subject to} \quad p \sum_{i=1}^{k-1} \lambda_i \lambda_{i+1} + \sum_{i=1}^{k-2} \sum_{j \geqslant i+2}^{k} \lambda_i \lambda_j = d,$$

$$\sum_{i=1}^{k} \lambda_i = 1, \qquad 0 < \lambda_i < 1, \qquad 0 \leqslant p \leqslant 1.$$

However, the solution of this problem is technically quite involved and we therefore defer it to a separate paper [6], where – based on the results presented here – we determine $\mathrm{c}(d)$ in the complete interval $0 < d < \frac{1}{2}$.

Let us say a few words about the underlying idea of the proof of Theorem 1.1. We first show that every poset is very close to one with a certain 'partitionable' structure. Here the main tool will be Szemerédi's Regularity Lemma, or rather an analogue of the latter for partial orders (Lemma 3.1), which might be of independent interest. Then we prove in a second step that it suffices to consider the case where the partition classes are arranged in a 'linear' way, *i.e.*, where they form a layer poset. For this step we shall use and prove the following elementary lemma, which may find further applications, too.

Lemma 1.3. *For every poset $P \in \mathscr{P}_n$ with height k there exists a k-layer poset $P' \in \mathscr{P}_n$ that has*

(i) *at least as many comparable pairs as P,* *and*
(ii) *at least as many cover relations as P.*

This paper is organized as follows. We conclude this introduction with a few words on how to round real numbers when defining the set $\mathscr{P}_{n,Q}$ and with some remarks concerning notation and terminology. Section 2 contains the proof of Lemma 1.3, our first auxiliary result. In Section 3 we then prove Theorem 1.1, using the second auxiliary result, Lemma 3.1, whose proof can be found in Section 4.

Consider an arbitrary configuration $Q = (\lambda_1, \ldots, \lambda_k; p)$ with $0 < \lambda_i < 1$, $\sum_i \lambda_i = 1$, and $0 \leqslant p \leqslant 1$. When we defined the set $\mathscr{P}_{n,Q}$ we assumed that $\lambda_i n$ and $p\lambda_i\lambda_{i+1}n^2$ were integers. Here we demonstrate that this assumption can be made without loss of generality. More precisely, we will show that for every n it is possible to choose $\lambda'_1, \ldots, \lambda'_k$ and $p'_{1,2}, \ldots, p'_{k-1,k}$ such that

$$\lambda'_i n \in \mathbb{N}, \qquad p'_{i,i+1}\lambda'_i\lambda'_{i+1}n^2 \in \mathbb{N}, \qquad \sum_i \lambda'_i n = n,$$

$$\sum_i p'_{i,i+1}\lambda'_i\lambda'_{i+1}n^2 + \sum_{j\geqslant i+2} \lambda'_i\lambda'_j n^2 = \left[\sum_i p\lambda_i\lambda_{i+1}n^2 + \sum_{j\geqslant i+2}\lambda_i\lambda_j n^2\right], \tag{1.10}$$

$$|\lambda_i - \lambda'_i| = O(1/n) \quad \text{and} \quad |p - p'_{i,i+1}| = O(1/n) \tag{1.11}$$

will hold. Now we redefine $\mathscr{P}_{n,Q}$ to be the set of all k-layer posets $P \in \mathscr{P}_n$ (with layers $X_1, \ldots, X_k$) that satisfy $|X_i| = \lambda'_i n$ and have exactly $p'_{i,i+1}|X_i||X_{i+1}|$ comparable pairs between X_i and X_{i+1}. By (1.10), posets in $\mathscr{P}_{n,Q}$ have $[\mathrm{d}(Q)n^2]$ comparable pairs, where $\mathrm{d}(Q)$ is still defined via the λ_i and p as in (1.8). Obviously, λ'_i and $p'_{i,i+1}$ will now depend on n, but, as observed in (1.11), they will always be very close to λ_i and p. Hence the estimate for $|\mathscr{P}_{n,Q}|$ from (1.9) remains true for the old definition of $\mathrm{c}(Q)$ via the λ_i and p together with the new definition of $\mathscr{P}_{n,Q}$ via the λ'_i and $p'_{i,i+1}$.

To see that it is possible to choose λ'_i and $p'_{i,i+1}$ as above, choose some integers n_i that satisfy $\lfloor \lambda_i n \rfloor \leqslant n_i \leqslant \lceil \lambda_i n \rceil$ and $\sum n_i = n$, and let $\lambda'_i := n_i/n$. This already implies $|\lambda_i - \lambda'_i| = O(1/n)$, and hence

$$\left| \sum_{j\geqslant i+2} \lambda'_i\lambda'_j n^2 \; - \sum_{j\geqslant i+2} \lambda_i\lambda_j n^2 \right| = O(n).$$

In other words, by rounding the λ_i we obtain a linear error in the number of comparable pairs between X_i and X_j (where $j \geqslant i+2$), which we need to balance in order to guarantee (1.10). The balancing can be done by slightly varying the number of comparable pairs between X_i and X_{i+1}: choose $p'_{i,i+1}$ so that $p'_{i,i+1}\lambda'_i\lambda'_{i+1}n^2$ is an integer, $|p'_{i,i+1}\lambda'_i\lambda'_j n^2 - p\lambda_i\lambda_j n^2| = O(n)$ and (1.10) is satisfied.

For a *partially ordered set* $P = (X, P)$ (often abbreviated as *poset*) and two points $x, y \in X$, we write $x \leqslant y$ if $(x, y) \in P$, and $x < y$ if $x \leqslant y$ and $x \neq y$. If $x < y$ then we say that x, y form a *comparable pair*, or, in abuse of notation, a *relation*. If neither $x \leqslant y$ nor $y \leqslant x$ then we say that x and y are *incomparable* and write $x \| y$. We denote by $\mathrm{inc}(x)$ the set of all points that are incomparable to x. Moreover we say that x is *covered* by y (also y *covers* x, or (x, y) is a *cover relation*) if $x < y$ and there is no point z for which $x < z$ and $z < y$ holds. In this case we write $x <: y$. On the other hand, if $x < y$ but (x, y) is *not* a cover relation, we write $x \ll y$ and call it a *forced relation*.

A subset $\{x_1, \ldots, x_k\} \subseteq X$ is called a *chain* if all pairs x_i, x_j are comparable. It is called an *antichain* if all pairs are incomparable. In the case of a chain we write $[x_1, \ldots, x_k]$ if $x_1 < \cdots < x_k$. If the complete point set X is a chain, P is called a *linear* order.

A point x is called *maximal* (respectively, *minimal*) if there is no point y with $x < y$ (respectively, $y < x$). A chain is called *maximal* if it cannot be extended to a larger chain. It is called *maximum* if no other chain contains more points. The *height* of a poset is the number of points in a maximum chain.

With a poset $P = (X, P)$ we associate the *comparability digraph* G and the *cover graph* G'. The vertex-sets of both graphs are given by X, the edges (x, y) in G are formed by the comparable pairs $x < y$ in P, while the edges $\{x, y\}$ in G' are formed by the cover relations $x <: y$ in P.

2. Proof of Lemma 1.3

For a poset P, denote by $\sigma(P)$ the number of comparable pairs, by $\sigma_\infty(P)$ the number of incomparable pairs, and by $\sigma_1(P)$ the number of cover relations in P. Every pair is counted only once.

Proof of Lemma 1.3. Let $C_1, \ldots, C_\ell$ be a chain decomposition of P that is obtained by recursively removing maximum chains from P. Hence we have that $|C_1| \geqslant \cdots \geqslant |C_\ell|$ and furthermore that C_i is a maximal chain in $P - C_1 - \cdots - C_{i-1}$ for all $i \in [\ell]$. Let $c_i := |C_i|$. Note that $c_1 = \text{height}(P) = k$. The underlying idea of the proof is to glue the chains C_i together again, but in such a way as to control carefully the parameters σ and σ_1.

We first give bounds on $\sigma_1(P)$ and $\sigma_\infty(P)$. Within each chain C_i there can be at most $c_i - 1$ cover relations. Denote the number of cover relations between two chains C_i and C_j with $i < j$ by $\sigma_1(C_i, C_j)$. Hence

$$\sigma_1(P) \leqslant \sum_{i=1}^{\ell}(c_i - 1) + \sum_{i<j} \sigma_1(C_i, C_j), \tag{2.1}$$

and we claim that

$$\sigma_1(C_i, C_j) \leqslant \begin{cases} 2c_j - 2, & \text{if } c_i = c_j, \\ 2c_j - 1, & \text{if } c_i = c_j + 1, \\ 2c_j, & \text{always.} \end{cases} \tag{2.2}$$

The best way to see this might be to view this as a bipartite graph with vertex-sets C_i, C_j where an edge represents a cover relation. Since each point in one chain *can cover* at most one point and *can be covered* by at most one point from the other chain, the graph has maximum degree at most 2. Let $C_i = [x_{c_i}, \ldots, x_1]$ and $C_j = [y_{c_j}, \ldots, y_1]$. Then x_1 cannot be covered by any element in C_j and x_{c_i} cannot cover any element in C_j (otherwise C_i would not be maximal), so they have degree at most one. Hence the sum of the degrees in C_i is bounded from above by $2c_i - 2$ (which settles the first case of the claim) and the sum of the degrees in C_j is bounded from above by $2c_j$ (which settles the third case). For the second case, where $c_i = c_j + 1$, the only possibility that might contradict our claim would be if $\sigma_1(C_i, C_j) = 2c_j = 2c_i - 2$, implying that all points in C_i *and* C_j indeed have degree 2 except for x_1 and x_{c_i}, which have degree 1. Now if x_1 did cover y_i for any

$i > 1$, then the point y_1 could not be covered, hence $y_1 <: x_1$. Similarly y_1 cannot cover any point other than x_2, for otherwise x_2 could not be covered. Thus $x_2 <: y_1$. But now $[x_{c_i}, \ldots, x_3, x_2, y_1, x_1]$ contradicts the maximality of C_i. This completes the proof of (2.2).

To give a lower bound on $\sigma_\infty(P)$ note that any incomparable pair in $C_i \cup C_j$ must obviously have one point in C_i and one point in C_j. Denote the number of such pairs by $\sigma_\infty(C_i, C_j)$. We claim that

$$\sigma_\infty(C_i, C_j) \geqslant c_j. \tag{2.3}$$

Suppose that a point $y \in C_j$ were comparable to all points $x \in C_i$. This would imply the existence of some index t with $0 \leqslant t \leqslant c_i$ such that $x_{t+1} < y < x_t$. But then $[x_{c_i}, \ldots, x_{t+1}, y, x_t, \ldots, x_1]$ contradicts the maximality of C_i. (The cases $t = 0$ and $t = c_i$ then correspond to $[x_{c_i}, \ldots, x_1, y]$ and $[y, x_{c_i}, \ldots, x_1]$.) Therefore every point $y \in C_j$ must be incomparable to at least one point $x \in C_i$, hence in total $\sigma_\infty(C_i, C_j) \geqslant c_j$, which proves (2.3).

If we now succeed in constructing a layer poset P' by taking the chains $C_1, \ldots, C_\ell$ as 'building blocks' (which means that, again, C_i is a maximum chain in $P' - C_1 - \cdots - C_{i-1}$) and combining them in such a way that (2.1), (2.2), and (2.3) hold for P' with equality, then

$$\sigma_1(P') \geqslant \sigma_1(P)$$

and

$$\sigma_\infty(P') \leqslant \sigma_\infty(P).$$

Observe that this would immediately imply (i) and (ii) as stated in the lemma.

To construct P' now renumber the points in the chains C_i so that

$$\begin{aligned} C_i &= [x^i_{c_i}, \ldots, x^i_4, x^i_2, x^i_1, x^i_3, \ldots, x^i_{c_i-1}] \quad \text{if } c_i \text{ is even}, \\ C_i &= [x^i_{c_i-1}, \ldots, x^i_4, x^i_2, x^i_1, x^i_3, \ldots, x^i_{c_i}] \quad \text{if } c_i \text{ is odd}. \end{aligned}$$

For $s = 1, \ldots, c_1$ we now let the sets

$$A_s := \{x^i_s : \quad \text{for all } i \in [\ell] \text{ where } s \leqslant c_i\}$$

form antichains. (They constitute the layers in P'.) Now add all cover relations in

$$\ldots, A_4 \times A_2,\ A_2 \times A_1,\ A_1 \times A_3,\ A_3 \times A_5, \ldots.$$

(An alternative description of the same construction is to say that P' is a c_1-layer poset with configuration $(\ldots, a_4, a_2, a_1, a_3, \ldots; 1)$ where $a_s := |A_s|/n$.) For an illustration of this construction see Figure 1, where only the cover relations inside the chains C_i and those involving x^3_1 are shown.

Observe that, for any two chains C_i and C_j with $i < j$, the only point in $C_i \cup C_j$ that is incomparable to $x^j_s \in C_j$ is x^i_s, hence (2.3) holds with equality. Moreover it is easy to see that equality also holds in (2.1) and (2.2). Finally, by the construction of P', it is clear that height$(P') = c_1 = k$. □

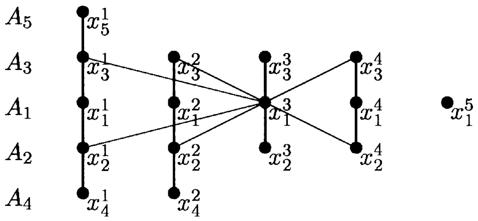

Figure 1 Construction of P'

3. Proof of Theorem 1.1

In order to prove the main theorem, we need a slightly more general concept of a configuration than the one used in the introduction. By a *k-configuration* Q we now mean a weighted poset with point set $\{x_1, \ldots, x_k\}$, where every point x_i carries weight λ_i (where $0 < \lambda_i < 1$ and $\sum_i \lambda_i = 1$) and every relation (x_i, x_j) carries weight $0 \leqslant p_{i,j} \leqslant 1$. Forced relations (x_i, x_j) must all have weight $p_{i,j} \equiv 1$.

We say that a poset $P = (X, P) \in \mathscr{P}_n$ has k-configuration Q, if there exists a partition of its point set $X = X_1 \cup \ldots \cup X_k$ into k antichains such that, for $x \in X_i$ and $y \in X_j$, one can only have $x < y$ in P if $x_i < x_j$ in Q. On the other hand, if $x_i < x_j$ in Q then there must be exactly $p_{i,j}|X_i||X_j|$ comparable pairs $x < y$ with $x \in X_i$ and $y \in X_j$ in P. Furthermore we require that the partition classes satisfy $|X_i| = \lambda_i \cdot n$ for all i. Again, $\mathscr{P}_{n,Q}$ denotes the set of all posets in $\mathscr{P}_n$ that have configuration Q. (Obviously the same remarks concerning the rounding of real numbers as in the introduction apply, so we do not repeat them here.)

A poset $P = (X, P) \in \mathscr{P}_n$ will be called *k-partitionable* if it has a k-configuration. Obviously, every poset is n-partitionable (in which case its configuration is just the poset itself), but we will be interested in partitionable posets with a constant number of classes.

If the number of points in a k-configuration Q is clear or irrelevant, we will simply speak of a configuration. A configuration Q is called *linear* if Q is a linear order. It is called *p-uniform* if there exists a $0 \leqslant p \leqslant 1$ such that $p_{i,j} \equiv p$ for all cover relations (x_i, x_j) in Q. The unique (up to isomorphism) *complete* poset P *induced* by a configuration Q is obtained by letting $p_{i,j} \equiv 1$ for all relations (x_i, x_j) in Q.

Comparing this with the terminology used in the introduction, a poset is a k-layer poset if it has a p-uniform linear k-configuration.

For the proof of Theorem 1.1, the following lemma makes the breakthrough by showing that every poset is 'close' to a k-partitionable poset (for some constant k).

Lemma 3.1. *For every $\epsilon > 0$ and every $0 < d < \frac{1}{2}$, there exist two constants k_0, n_0 such that, for every poset $P \in \mathscr{P}_{n,d}$ with $n \geqslant n_0$, there is a k-partitionable poset $P' \in \mathscr{P}_{n,d}$ with $k \leqslant k_0$ that differs from P in at most ϵn^2 relations and in which the partition classes differ in size by at most one.*

The proof of this lemma is based on Szemerédi's Regularity Lemma, and shows in addition to the above properties that the partition is ϵ-regular in the usual sense (see

Section 4). It thus seems to be the natural translation of the Regularity Lemma to partial orders and may well find further applications. However, the proof of Lemma 3.1 requires some work and is different in nature from the other proofs in this section, so we defer it to the last section.

Denote by $\mathscr{P}^k_{n,d}$ the family of all k-partitionable posets from $\mathscr{P}_{n,d}$. Then Lemma 3.1 states that we can enumerate the set $\mathscr{P}_{n,d}$ in the following way: for every ϵ there exist two constants k_0, n_0 such that, if $n \geqslant n_0$,

$$\mathscr{P}_{n,d} = \bigcup_{k=1}^{k_0} \bigcup_{P \in \mathscr{P}^k_{n,d}} \Gamma_\epsilon(P),$$

where $\Gamma_\epsilon(P)$ denotes all those posets in $\mathscr{P}_{n,d}$ that can be constructed from P by changing at most ϵn^2 relations: these are no more than

$$\binom{n^2}{\epsilon \cdot n^2} = 2^{(1+o(1))H(\epsilon)n^2},$$

where, as before, $H(\epsilon)$ denotes the entropy function. Thus the following corollary holds.

Corollary 3.2. *For every $\epsilon > 0$ there exists a constant k_0 such that*

$$|\mathscr{P}_{n,d}| = 2^{(1+o(1))H(\epsilon)n^2} \cdot \sum_{k=1}^{k_0} |\mathscr{P}^k_{n,d}|. \qquad \square$$

For a given configuration we would like to count the number of different posets that have this configuration. The following is no more than a generalization of the discussion in the introduction. Let P be a k-partitionable poset with partition $X_1, \dots, X_k$ and let Q be its configuration with point set $\{x_1, \dots, x_k\}$. When counting the number of posets with configuration Q it is clear that the degree of freedom we have lies in where we place the $p_{i,j}|X_i||X_j|$ relations between X_i and X_j when $x_i <: x_j$ in Q. Hence the number is approximately

$$\prod_{x_i <: x_j} \binom{\lambda_i \lambda_j n^2}{p_{i,j} \cdot \lambda_i \lambda_j n^2} = 2^{\sum_{i,j} H(p_{i,j})\lambda_i\lambda_j n^2 \; + o(n^2)}.$$

where the sum is taken over all pairs i, j with $x_i <: x_j$. (Again, we wasted a factor of $O(n!) = O(2^{n \log_2 n})$ since we did not assign points to classes.) Let

$$\mathrm{c}(Q) := \sum_{x_i <: x_j} H(p_{i,j})\lambda_i\lambda_j.$$

Observe that Q determines the total number of relations in P. It must have dn^2 relations where

$$d = \mathrm{d}(Q) := \sum_{x_i \ll x_j} \lambda_i\lambda_j + \sum_{x_i <: x_j} p_{i,j}\lambda_i\lambda_j.$$

We will refer to these parameters as the c-*value* and the d-*value* of the configuration Q. A configuration is called *d-significant* if it has d-value d and if there is no other configuration (possibly with a different number of partition classes) that has the same d-value and a higher c-value. Sometimes we only say that a configuration Q is significant – obviously this means that it is $\mathrm{d}(Q)$-significant.

Notice that since k is a constant and independent of n, there are not actually all that many different k-configurations: there are less than $2^{k^2/2}$ posets, for each X_i there is a choice of at most n values to determine $|X_i|$. Finally, for each pair (X_i, X_j) there is a choice of less than n^2 values to determine the number of relations between X_i and X_j. Therefore in total there are $2^{o(n^2)}$ different k-configurations. Hence, if Q is a d-significant configuration then

$$|\mathscr{P}^k_{n,d}| = \sum_{Q'} 2^{c(Q')n^2} = 2^{c(Q)\cdot n^2 + o(n^2)},$$

where the sum is taken over all k-configurations Q' with $\mathrm{d}(Q') = d$. Together with Corollary 3.2 this now proves the following lemma.

Lemma 3.3. *For any $0 < d < \frac{1}{2}$ let Q be a d-significant configuration. Then*

$$|\mathscr{P}_{n,d}| = |\mathscr{P}_{n,Q}|\ 2^{o(n^2)} = 2^{c(Q)\cdot n^2 + o(n^2)}. \qquad \square$$

Comparing our present position as stated in Lemma 3.3 with our aim as stated in Theorem 1.1, it now suffices to show that for each d-significant configuration Q there exists a p-uniform linear configuration Q' with $\mathrm{d}(Q') = \mathrm{d}(Q)$ and $\mathrm{c}(Q') \geqslant \mathrm{c}(Q)$.

Lemma 3.4. *Any significant configuration Q must be p-uniform for some $p \in [0, 1]$.*

Proof. Assume without loss of generality that we have the cover relations $x_1 <: x_3$ with weight p_1 and $x_2 <: x_4$ with weight p_2. Suppose that $p_1 \neq p_2$. Let

$$p' := \frac{p_1\lambda_1\lambda_3 + p_2\lambda_2\lambda_4}{\lambda_1\lambda_3 + \lambda_2\lambda_4}$$

and consider the configuration Q' derived from Q by replacing both p_1 and p_2 by p'. In Q as well as in Q', the cover relations $x_1 <: x_3$ and $x_2 <: x_4$ together contribute

$$p_1\lambda_1\lambda_3 + p_2\lambda_2\lambda_4$$

to $\mathrm{d}(Q)$ and $\mathrm{d}(Q')$ respectively, so Q and Q' have the same d-value. But because of the concavity of $H(x)$ we have that

$$H(p_1)\lambda_1\lambda_3 + H(p_2)\lambda_2\lambda_4 < H\left(\frac{p_1\lambda_1\lambda_3 + p_2\lambda_2\lambda_4}{\lambda_1\lambda_3 + \lambda_2\lambda_4}\right)(\lambda_1\lambda_3 + \lambda_2\lambda_4) = H(p')(\lambda_1\lambda_3 + \lambda_2\lambda_4),$$

which means that the c-value of Q is smaller than that of Q'. $\square$

In other words, in a significant configuration Q all cover relations carry the same weight, which we will call the *density* of Q and denote by $p = \mathrm{p}(Q)$.

Recall that we denote by $\sigma(P)$ the number of comparable pairs in P and by $\sigma_1(P)$ the number of cover relations in P. Now, similarly for configuration Q, let

$$\sigma_1(Q) := \sum_{x_i <: x_j} \lambda_i\lambda_j, \quad \sigma_2(Q) := \sum_{x_i \ll x_j} \lambda_i\lambda_j, \quad \sigma(Q) := \sigma_1(Q) + \sigma_2(Q),$$

$$\sigma_0(Q) := \sum_{i=1}^{k} (\lambda_i)^2, \quad \sigma_\infty(Q) := \sum_{x_i \| x_j} \lambda_i\lambda_j.$$

Notice that for a poset $P \in \mathscr{P}_{n,Q}$, a pair $x_i \ll x_j$ in Q contributes $\lambda_i\lambda_j$ to $\sigma(Q)$ and $\lambda_i\lambda_j n^2$ to $\sigma(P)$. Similarly, a pair $x_i <: x_j$ in Q contributes $\lambda_i\lambda_j$ to $\sigma_1(Q)$ and $p_{i,j}\lambda_i\lambda_j n^2$ to $\sigma_1(P)$. Thus we have

$$\sigma(P) \leqslant \sigma(Q) \cdot n^2, \qquad \sigma_1(P) \leqslant \sigma_1(Q) \cdot n^2,$$

and equality holds if and only if for all $x_i < x_j$ in Q *all* cover relations between the two partition classes X_i and X_j exist in P, *i.e.*, if P is the complete poset induced by Q. The next step is a corollary derived from Lemma 1.3.

Corollary 3.5. *For every k-configuration Q there exists a linear k'-configuration Q' with $k' \leqslant k$ such that*

$$\sigma_1(Q') \geqslant \sigma_1(Q) \qquad \textit{and} \qquad \sigma(Q') \geqslant \sigma(Q).$$

Proof of Corollary 3.5. Denote by P the complete k-partitionable poset on n points induced by Q. Apply Lemma 1.3 to P and obtain a k'-layer poset P'. Let Q' be the (linear) configuration of P'. Obviously

$$\sigma_1(Q') \geqslant \frac{\sigma_1(P')}{n^2} \geqslant \frac{\sigma_1(P)}{n^2} = \sigma_1(Q),$$

and

$$\sigma(Q') \geqslant \frac{\sigma(P')}{n^2} \geqslant \frac{\sigma(P)}{n^2} = \sigma(Q). \qquad \square$$

Now we use the new terminology to simplify the expressions for $\mathrm{c}(Q)$ and $\mathrm{d}(Q)$. We then prove Theorem 1.1 for $d \in (0, \frac{1}{8}]$. Note that since $\sum_{i=1}^{k} \lambda_i = 1$ we have

$$2 \cdot (\sigma_1(Q) + \sigma_2(Q) + \sigma_\infty(Q)) + \sigma_0(Q) = 1. \tag{3.1}$$

For a significant configuration we can (using Lemma 3.4) now write

$$\mathrm{c}(Q) = H(p) \cdot \sigma_1(Q) \quad \text{and} \tag{3.2}$$

$$\mathrm{d}(Q) = p \cdot \sigma_1(Q) + \sigma_2(Q) = \frac{1}{2} - \frac{1}{2}\sigma_0(Q) - (1-p)\sigma_1(Q) - \sigma_\infty(Q), \tag{3.3}$$

where $p = \mathrm{p}(Q)$. Observe that for every configuration Q we must have

$$\sigma_1(Q) \leqslant \frac{1}{4}. \tag{3.4}$$

This can be easily established as follows. Consider the cover graph of Q with weight λ_i on the vertex x_i. Denote by γ_i the sum of the weights of all neighbours of x_i. We propose the following process. As long as there are two non-adjacent vertices x_i, x_j with positive weights λ_i, λ_j, take the vertex with smaller neighbourhood weight (say $\gamma_i \leqslant \gamma_j$) and shift its weight completely to the other vertex: $\lambda_i := 0$, $\lambda_j := \lambda_j + \lambda_i$. We will check the following two observations.

(1) During this process $\sigma_1(Q)$ does not decrease.
(2) After the process there will be only two vertices with positive weight.

Hence at the end $\sigma_1 \leqslant \frac{1}{4}$, which would prove the proposition.

To check (1), simply observe that during one step of the process the loss in σ_1 is $\lambda_i \cdot \gamma_i$ and the win is $\lambda_i \cdot \gamma_j$; hence in total we do not lose anything. As for (2), since the graph is triangle-free, if there are at least three vertices with positive weight, we will always find two non-adjacent ones.

Proof of Theorem 1.1 for $0 < d \leqslant \frac{1}{8}$. By Lemma 3.3 it suffices to prove that, for an arbitrary d-significant configuration Q' (which, by Lemma 3.4, must be p-uniform), there exists a linear configuration Q with $\mathrm{d}(Q) = \mathrm{d}(Q')$ and $\mathrm{c}(Q) \geqslant \mathrm{c}(Q')$. We claim that choosing Q as in (1.7) will succeed. Q has point set $\{x_1, x_2\}$ with

$$x_1 <: x_2, \qquad \lambda_1 := \lambda_2 := \frac{1}{2},$$
$$p := \mathrm{p}(Q) := 4d \;=\; 4p'\sigma_1(Q') + 4\sigma_2(Q'),$$

where $p' = \mathrm{p}(Q')$. Obviously $\mathrm{d}(Q) = \frac{1}{4}p = d = \mathrm{d}(Q')$ and our aim is to show that $\mathrm{c}(Q) \geqslant \mathrm{c}(Q')$. Observe that $d \leqslant \frac{1}{8}$ implies that $p \leqslant \frac{1}{2}$, hence for any $p'' \leqslant p$ we have $H(p) \geqslant H(p'')$. Furthermore, by the concavity of $H(x)$ we know that, for any $0 < \alpha \leqslant 1$, we have that $H(\alpha \cdot x) \geqslant \alpha \cdot H(x)$. Equipped with these facts we abbreviate $\sigma_1' := \sigma_1(Q')$ and $\sigma_2' := \sigma_2(Q')$ and obtain

$$\begin{aligned} c(Q) &= H(p) \cdot \lambda_1\lambda_2 = H(4\sigma_2' + 4p'\sigma_1') \cdot \frac{1}{4} \\ &\geqslant H(p' \cdot 4\sigma_1') \cdot \frac{1}{4} \geqslant H(p')\sigma_1' \;=\; \mathrm{c}(Q'), \end{aligned}$$

where we first applied the definition of p, λ_1, and λ_2, then used $p \leqslant \frac{1}{2}$, and finally relied on $4\sigma_1' \leqslant 1$, which is guaranteed by observation (3.4). □

Hence we can from now on assume that $d \geqslant \frac{1}{8}$. In the following lemmas we will often start from a configuration Q and build a new configuration Q', possibly with a different number of partition classes, different weights and relations. Often we will shift weight ϵ from one point x_i to another point x_j, *i.e.*, $\lambda(x_i) := \lambda(x_i) - \epsilon$ and $\lambda(x_j) := \lambda(x_j) + \epsilon$. When doing so, we will sometimes refer to the original weights as λ_i, and to the new weights as $\lambda(x_i)$. Since we will be moving from one linear configuration to another, $\sigma_\infty(Q) = 0$ will always hold. Therefore (3.1) now stands as

$$2\sigma(Q) + \sigma_0(Q) = 1. \tag{3.5}$$

Another trivial observation: If $\lambda_i \geqslant \lambda_j$ then shifting any weight $0 \leqslant \epsilon \leqslant \lambda_i - \lambda_j$ from x_i to x_j will not increase $\sigma_0(Q)$:

$$\sigma_0(Q') - \sigma_0(Q) = (\lambda_i - \epsilon)^2 + (\lambda_j + \epsilon)^2 - \lambda_i^2 - \lambda_j^2 = -2\epsilon(\lambda_i - \lambda_j - \epsilon) \leqslant 0. \tag{3.6}$$

Lemma 3.6. *For every linear configuration Q and for every $0 < s \leqslant \sigma_1(Q)$ there exists a linear configuration Q' such that*

$$\sigma(Q') \geqslant \sigma(Q) \qquad \text{and} \qquad \sigma_1(Q') = s.$$

Proof. Assume w.l.o.g. that $s < \sigma_1(Q)$, for otherwise $Q' := Q$ does the job. Let $Q = [x_k, \ldots, x_1]$. We will shift weights several times, so denote by λ_i the original weights in Q. Let ϵ be such that $\epsilon < \lambda_i - 2\epsilon k$ for all $i \in [k]$. In a first round we add $2k$ new points $y_1, \ldots, y_{2k}$ and obtain a new configuration $[x_k, \ldots, x_1, y_1, \ldots, y_{2k}]$. To the new points we assign weight $\lambda(y_j) := \epsilon$ for all $j \in [2k]$ and reduce the weight of x_1 by $2\epsilon k$. Using observation (3.6) it is clear that $\sigma_0(Q') \leqslant \sigma_0(Q)$. Hence by (3.5) $\sigma(Q') \geqslant \sigma(Q)$.

In a second round, for all $i \in [k]$ consecutively, shift weight ϵ_i from x_i to y_{2i}, where $0 \leqslant \epsilon_i \leqslant \lambda(x_i) - \epsilon$. As before, use observation (3.6) to see that $\sigma_0(Q') \leqslant \sigma_0(Q)$ and therefore, by (3.5), $\sigma(Q') \geqslant \sigma(Q)$.

So, no matter how we choose ϵ and all the ϵ_i (provided they satisfy the above inequalities), the first assertion of the lemma is guaranteed. For a particular choice of ϵ make ϵ_i as large as possible, namely $\epsilon_1 := \lambda_1 - 2\epsilon k - \epsilon$ and $\epsilon_i := \lambda_i - \epsilon$ for $2 \leqslant i \leqslant k$. Then we have $\lambda(x_i) = \lambda(y_{2i-1}) = \epsilon$ for all $i \in [k]$, $\lambda(y_2) = \lambda_1 - 2\epsilon k$ and $\lambda(y_{2i}) = \lambda_i$ and hence an upper bound on $\sigma_1(Q')$ is given by

$$\sigma_1(Q') \leqslant \epsilon^2 \cdot k + \sum_{i=1}^{k} 2\epsilon \cdot \lambda_i \leqslant \epsilon \cdot (k+2).$$

This means that for a given s it is possible to choose ϵ sufficiently small that the above two-round process can force $\sigma_1(Q')$ to become arbitrarily small. To ensure that the process produces $\sigma_1(Q') = s$, we choose ϵ so small that after the first round we still have $\sigma_1(Q') > s$ and $\epsilon(k+2) < s$. Then continuously increase the ϵ_i until at some point the second round must produce a Q' with $\sigma_1(Q') = s$. □

Corollary 3.7. *For every configuration Q there exists a linear configuration Q'' satisfying*

$$\sigma_1(Q'') = \sigma_1(Q), \qquad \sigma_2(Q'') \geqslant \sigma_2(Q).$$

Proof. Apply Corollary 3.5 to Q and obtain a linear configuration Q' with

$$\sigma_1(Q') \geqslant \sigma_1(Q) \qquad \text{and} \qquad \sigma(Q') \geqslant \sigma(Q).$$

Now apply Lemma 3.6 to Q', setting $s := \sigma_1(Q) \leqslant \sigma_1(Q')$. We obtain a linear configuration Q'' with

$$\sigma_1(Q'') = s = \sigma_1(Q), \qquad \sigma(Q'') \geqslant \sigma(Q') \geqslant \sigma(Q),$$

and therefore

$$\sigma_2(Q'') \geqslant \sigma_2(Q),$$

as we were required to prove. □

Lemma 3.8. *For every linear configuration Q and for every $0 \leqslant s \leqslant \sigma_2(Q)$ there exists a linear configuration Q' satisfying*

$$\sigma_1(Q') \geqslant \sigma_1(Q), \qquad \sigma_2(Q') = s.$$

Proof. Let $Q = [x_k, \ldots, x_1]$. Again denote by $\lambda_i := \lambda(x_i)$ the original weights. Start with x_1 and shift an increasing amount ϵ_1 of weight to x_3, until $\lambda(x_1) = 0$ and hence $\lambda(x_3) = \lambda_3 + \lambda_1$. Then move on to x_2, shifting weight ϵ_2 to x_4 until $\lambda(x_2) = 0$. Continue until the final step, where weight ϵ_{k-2} is shifted from x_{k-2} to x_k.

Notice that, whenever weight ϵ_i is shifted from x_i to x_{i+2}, we can be sure that x_i is the maximum of the chain and that the only point covered by x_i is x_{i+1}, which in turn also covers x_{i+2}. So if Q' denotes the new configuration, we have $\sigma_1(Q') \geqslant \sigma_1(Q)$ at any moment of the process.

Observe that, if the process runs until the very end, Q' has only two points x_{k-1} and x_k, and hence $\sigma_2(Q') = 0$. But since this process is continuous it must at one point produce a Q' with $\sigma_2(Q') = s$ for any $0 \leqslant s \leqslant \sigma_2(Q)$. □

Now we come back to the d- and c-value of a p-uniform linear configuration Q. Recall that they are given by $\mathrm{d}(Q) = p \cdot \sigma_1(Q) + \sigma_2(Q)$ and $\mathrm{c}(Q) = H(p) \cdot \sigma_1(Q)$, where $p = \mathrm{p}(Q)$.

Corollary 3.9. *For every linear configuration Q and for every d with $\frac{1}{8} \leqslant d \leqslant \mathrm{d}(Q)$ there exists a linear configuration Q' such that*

$$\mathrm{c}(Q') \geqslant \mathrm{c}(Q), \qquad \mathrm{d}(Q') = d.$$

Proof. Apply Lemma 3.8 to Q, with s slowly decreasing from $\sigma_2(Q)$ to 0. This means that on the one hand $\sigma_1(Q)$ does not decrease and thus $\mathrm{c}(Q)$ does not; and on the other hand it means that simultaneously $\sigma_2(Q)$ steadily approaches 0. Having arrived there, denote the new configuration by Q' and observe that $\mathrm{c}(Q') \geqslant \mathrm{c}(Q)$ and $\mathrm{d}(Q') = p\sigma_1(Q')$, where $p = \mathrm{p}(Q)$. If $p > 1/2$ then let p approach $1/2$, thereby increasing $\mathrm{c}(Q') = H(p) \cdot \sigma_1(Q')$ and forcing $\mathrm{d}(Q')$ to approach $\frac{1}{2}\sigma_1(Q')$. Owing to (3.4) we know that

$$\frac{1}{2} \cdot \sigma_1(Q') \leqslant \frac{1}{8} \leqslant d.$$

Hence this process must reach the point where $\mathrm{d}(Q') = d$ while maintaining at all times $c(Q') \geqslant c(Q)$. □

The corollary above now allows us to prove the remaining part of our main theorem.

Proof of Theorem 1.1 for $d > \frac{1}{8}$. Consider an arbitrary d-significant configuration Q. By Lemma 3.3 it suffices to show that there exists a linear configuration Q' with $\mathrm{d}(Q') = \mathrm{d}(Q)$ and $\mathrm{c}(Q') \geqslant \mathrm{c}(Q)$. By Corollary 3.7 there exists a linear configuration Q'' with

$$\sigma_1(Q'') = \sigma_1(Q) \quad \text{and} \quad \sigma_2(Q'') \geqslant \sigma_2(Q).$$

Setting $\mathrm{p}(Q'') := \mathrm{p}(Q)$, this immediately implies that

$$\mathrm{c}(Q'') = \mathrm{c}(Q) \quad \text{and} \quad \mathrm{d}(Q'') \geqslant \mathrm{d}(Q).$$

Now apply Corollary 3.9 to Q'' and d. Hence there must be a linear configuration Q' with

$$\mathrm{d}(Q') = d = \mathrm{d}(Q) \quad \text{and} \quad \mathrm{c}(Q') \geqslant \mathrm{c}(Q'') = \mathrm{c}(Q). \qquad \square$$

We conclude this section with the proof of Lemma 1.2.

Proof of Lemma 1.2. Suppose to the contrary that there were a significant linear configuration Q with $\mathrm{d}(Q) \geqslant \frac{1}{8}$ and $p := \mathrm{p}(Q) < \frac{1}{2}$. Then increase p very slightly, thereby causing $H(p)$ to increase and thus both $\mathrm{d}(Q)$ and $\mathrm{c}(Q)$ must increase. Call the new configuration Q'' and set $d := \mathrm{d}(Q)$ so that

$$\frac{1}{8} \leqslant d < \mathrm{d}(Q'').$$

Applying Corollary 3.9 to Q'' and d, there must be a linear configuration Q' with

$$\mathrm{d}(Q') = d = \mathrm{d}(Q) \quad \text{and} \quad \mathrm{c}(Q') \geqslant \mathrm{c}(Q'') > \mathrm{c}(Q).$$

Hence Q cannot be significant. $\square$

4. Proof of Lemma 3.1

We start with a simple lemma which states that one can add or remove relations to a partitionable poset without forcing or destroying other relations, and still maintain a partitionable poset.

Lemma 4.1. *For any two constants d, d' in the interval $(0, \frac{1}{2})$ and any $k \in \mathbb{N}$, there exists a $\bar{k} = \bar{k}(d, d', k) \in \mathbb{N}$ such that the following holds. For every k-partitionable poset $P \in \mathscr{P}_{n,d}$ whose partition classes differ in size by at most one, there exists a k'-partitionable poset $P' \in \mathscr{P}_{n,d'}$ with $k \leqslant k' \leqslant \bar{k}$. The new poset P' differs from P in exactly $|d' - d| \cdot n^2$ relations and, again, its partition classes differ in size by at most one.*

Proof. Let Q be the configuration of P. If $d' < d$, then first remove relations between X_i and X_j whenever $x_i <: x_j$ in Q. If in all such pairs no relations are left, then this will turn previously forced relations in Q into cover relations and so the process can continue until there are no relations at all.

If $d' > d$, then we will have to add relations and there are three ways to do so:

(i) whenever $x_i <: x_j$ in Q simply add new relations between X_i and X_j. If this is not enough, then

(ii) whenever $x_i \| x_j$ in Q, add new relations between X_i and X_j. Here some care is needed to avoid the forcing of other relations: let x_i be a point in Q with $\mathrm{inc}(x_i) \neq \emptyset$ and choose x_j to be a point that is maximal within $\mathrm{inc}(x_i)$. Again, if this is not enough, then

(iii) split all partition classes X_i into two parts X_i^- and X_i^+ so that X_i^- and X_i^+ differ in size by at most one, maintain all previous relations and add new relations $x < y$ where $x \in X_i^-$ and $y \in X_i^+$.

Repeating and combining these steps produces a k'-partitionable poset $P' \in \mathscr{P}_{n,d'}$ where d' can be arbitrarily close to $\frac{1}{2}$. Observe also that, in order to obtain a density of at most $d' < \frac{1}{2}$, we can bound k' by a constant that depends only on d, d' and k, but not on n. □

For the proof of Lemma 3.1 we need Szemerédi's Regularity Lemma and some related definitions. Let $G = (V, E)$ be a graph and consider two disjoint subsets $A, B \subset V$. Denote by $E(A, B)$ the set of those edges in E that have one endpoint in A and one endpoint in B. Then the *density* $d(A, B)$ is defined as

$$d(A, B) = \frac{|E(A, B)|}{|A| \cdot |B|}.$$

For $\epsilon \in (0, 1)$ a pair A, B is called *ϵ-regular* if, for every $X \subseteq A$ and $Y \subseteq B$ satisfying

$$|X| > \epsilon|A| \quad \text{and} \quad |Y| > \epsilon|B|,$$

it is true that

$$|d(X, Y) - d(A, B)| < \epsilon.$$

To say it roughly, the Regularity Lemma [7] guarantees that, for every graph, one can find a partition of its vertex-set into classes of almost the same size such that almost all pairs are regular. Natural modifications to the original proof easily give the following variant, for which we need a few more definitions.

Let $G_1, \ldots, G_r$ be spanning subgraphs of a graph $G = (V, E)$. In this setting a partition $V = X_1 \cup \ldots \cup X_k$ is called *ϵ-regular* if the classes X_i differ in size by at most 1 and all but at most ϵk^2 pairs are ϵ-regular *for all* G_i. Such a partition is said to *refine* another partition $V = V_1 \cup \ldots \cup V_{k'}$ if for each $1 \leqslant i \leqslant k$ there exists a $1 \leqslant j \leqslant k'$ so that $X_i \subseteq V_j$.

Theorem 4.2 (Szemerédi's Regularity Lemma). *For every $\tilde{\epsilon} > 0$ and $\ell, r \geqslant 1$ there exist two positive integers $\tilde{n}_0 = \tilde{n}_0(\tilde{\epsilon}, \ell, r)$ and $\tilde{k}_0 = \tilde{k}_0(\tilde{\epsilon}, \ell, r)$ such that the following is true. If $G = (X, E)$ is a graph with $|X| \geqslant \tilde{n}_0$ and $X = X'_1 \cup \ldots \cup X'_\ell$ is a partition where the classes differ in size by at most one, and if $G_1, \ldots, G_r$ are spanning subgraphs of G, then there exists an $\tilde{\epsilon}$-regular partition $X = X_1 \cup \ldots \cup X_k$ with $\ell \leqslant k \leqslant \tilde{k}_0$ that refines the previous partition.* □

Proof of Lemma 3.1. Let $G = (X, E)$ be the comparability digraph of P.

Set $\tilde{\epsilon} := \epsilon/12$. Choose an integer k' so that $1/k' < \tilde{\epsilon}$ and take an arbitrary partition $X = X'_1 \cup \ldots \cup X'_{k'}$ satisfying

$$\left\lfloor \frac{n}{k'} \right\rfloor \leqslant |X'_i| \leqslant \left\lceil \frac{n}{k'} \right\rceil \qquad \text{for all} \quad 1 \leqslant i \leqslant k'.$$

Colour the edges of G in the following way. An edge (x, y) is coloured in *blue* if $x <: y$ in P. Otherwise it is coloured in *red*. Note that for every red edge (x, y) there must be a directed path $x, z_1, \ldots, z_k, y$ with $k \geqslant 1$ of blue edges.

Since we want to turn P into a partitionable poset, we will have to remove edges from G. Obviously red edges cannot be removed without destroying the transitivity. Thus by removing a family $\mathscr{F}$ of edges we *always* mean removing all blue edges in $\mathscr{F}$ and putting a spell on the red edges in $\mathscr{F}$: if later red edges in $\mathscr{F}$ turn blue, we will remove them as well. Notice that removing a blue edge results in a digraph which is the comparability digraph of a poset with one relation less.

We start by removing all edges of G that lie inside a class X'_i.

By the (ordered) pair (X'_i, X'_j) we denote the bipartite graph on the vertices $X'_i \cup X'_j$ that contains all edges that in G leave X'_i and enter X'_j. (Note that in (X'_i, X'_j) the edges are undirected.) Let

$$E_< := \bigcup_{i<j} E(X'_i, X'_j), \qquad E_> := \bigcup_{i>j} E(X'_i, X'_j),$$

and consider the two graphs $G_< := (X, E_<)$ and $G_> := (X, E_>)$.

Now apply Theorem 4.2 with parameters $\tilde{\epsilon}$, $\ell := k'$ and $r := 2$, a first partition $X = X'_1 \cup \ldots \cup X'_{k'}$ and the two spanning graphs $G_<$ and $G_>$. Thus we obtain two integers $\tilde{n}_0$ and $\tilde{k}_0$. Choose the constants n_0, k_0 in the statement of Lemma 3.1 so that $n_0 \geqslant \tilde{n}_0$ as well as $k_0 \geqslant \bar{k}(d \pm \epsilon/2, d, \tilde{k}_0)$. (The latter will enable us to apply Lemma 4.1 at the very end of our proof.) We are then guaranteed a partition

$$X = X_1 \cup \ldots \cup X_k$$

with $k \leqslant \tilde{k}_0$ that refines $X = X'_1 \cup \ldots \cup X'_{k'}$ and has the property that, for all but at most $\tilde{\epsilon}k^2$ pairs $i < j$, both (X_i, X_j) and (X_j, X_i) are $\tilde{\epsilon}$-regular (and the X_i differ in size by at most one).

Consider the following properties of an arbitrary pair A, B.

(i) All but at most $\tilde{\epsilon}|A|$ vertices in A have degree at least $2\tilde{\epsilon}|B|$, and analogously with the roles of A and B exchanged.

(ii) For every set $A' \subseteq A$ with $|A'| > \tilde{\epsilon}|A|$ the set of neighbours $\Gamma(A')$ must have cardinality at least $(1 - \tilde{\epsilon})|B|$, and analogously with the roles of A and B exchanged.

Any $\tilde{\epsilon}$-regular pair with density at least $3\tilde{\epsilon}$ satisfies property (i). For if not, then denote those vertices with degree less than $2\tilde{\epsilon}|B|$ by A'; then $d(A', B) < 2\tilde{\epsilon}$ would contradict the regularity. Any $\tilde{\epsilon}$-regular pair with density at least $\tilde{\epsilon}$ satisfies property (ii); if not, then the pair $A', B \setminus \Gamma(A')$ would again contradict the regularity.

The following third property is obviously not possessed by every regular pair.

(iii) All but at most $\tilde{\epsilon}|A|$ vertices in A have degree at least $(1 - \tilde{\epsilon})|B|$, and analogously with the roles of A and B exchanged.

Now call a pair (X_i, X_j) *good* if it has properties (i) and (ii). Call it *bad* otherwise. Observe that (iii) implies both (i) and (ii), so a pair satisfying (iii) will always be good. Remove those edges in G that lie in a bad pair (X_i, X_j) or (X_j, X_i), where $i < j$. Denote by P' the poset that is obtained in this way. Since all $\tilde{\epsilon}$-regular pairs with density at least $3\tilde{\epsilon}$ are good, observe that up to now at most $5\tilde{\epsilon}n^2$ edges (that is, relations in P) have been removed, namely

at most $k'(\frac{n}{k'})^2 \leqslant \tilde{\epsilon}n^2$ edges inside the X_i',
at most $k^2 \cdot 3\tilde{\epsilon} \cdot (\frac{n}{k})^2 = 3\tilde{\epsilon}n^2$ edges inside pairs (X_i, X_j) with density less than $3\tilde{\epsilon}$,
at most $\tilde{\epsilon}k^2(\frac{n}{k})^2 = \tilde{\epsilon}n^2$ edges inside irregular pairs (X_i, X_j).

Consider the digraph R with vertex-set $\{X_1, \dots, X_k\}$ and edges (X_i, X_j) if the pair (X_i, X_j) is good. We claim the following.

Claim 1. If $Y_1, Y_2, \dots, Y_l$ is a dipath in R with $l \geqslant 3$, then (Y_1, Y_l) has property (iii).

Claim 2. R is acyclic.

Claim 3. All the edges left in G lie in good pairs.

These claims are easily verified as follows.

Proof of Claim 1. Recall that G is still a comparability digraph, that is, $(x, y) \in E$ and $(y, z) \in E$ imply that $(x, z) \in E$. We assume that (Y_1, Y_2) and (Y_2, Y_3) are good pairs and prove that this implies that (Y_1, Y_3) has property (iii). Then Claim 1 follows by induction on l. As (Y_1, Y_2) has property (i), all but at most $\tilde{\epsilon}|Y_1|$ vertices in Y_1 have degree at least $2\tilde{\epsilon}|Y_2| > \tilde{\epsilon}|Y_2|$. Now, since (Y_2, Y_3) has property (ii), we are done. □

Proof of Claim 2. By Claim 1 we know that any cycle $Y_1, Y_2, \dots, Y_{l-1}, Y_1$ implies that all but at most $\tilde{\epsilon}|Y_1|$ vertices in Y_1 have at least $(1 - \tilde{\epsilon})|Y_1|$ neighbours in Y_1. This in turn implies a directed cycle $y_1, y_2, \dots, y_{l-1}, y_1$ in G, which is impossible, since G is the comparability digraph of P'. □

Proof of Claim 3. This is obviously true for blue edges. For red edges it might seem a little surprising at first glance, since so far we have never bothered to remove red edges. Nevertheless it is true: consider a red edge (y_1, y_l) with $y_1 \in Y_1$ and $y_l \in Y_l$. Then there must be a path $y_1, y_2, \dots, y_l$ of blue edges. Since blue edges can only be found in good pairs, there must be a directed path $Y_1, \dots, Y_l$ in R. Since R is acyclic, we must have $Y_1 \neq Y_l$, and Claim 1 implies that (Y_1, Y_l) is also good. □

By Claims 1 and 2, R is a comparability digraph and we denote by Q the corresponding poset with point set $\{x_1, \dots, x_k\}$. Then Claims 1 and 3 assert that P' will have configuration Q if we complete all pairs (X_i, X_j) that satisfy property (iii) at the cost of at most

$$\binom{k}{2}\left(\tilde{\epsilon}\frac{n}{k} \cdot \frac{n}{k} + \tilde{\epsilon}\frac{n}{k} \cdot \frac{n}{k}\right) < \tilde{\epsilon}n^2$$

new relations (and then choose the weights λ_i and $p_{i,j}$ in Q accordingly). Note that inserting these new relations does not violate transitivity: any new edge (y, y') lies in a pair with property (iii), and if, together with another (new or old) edge (y', y''), it requires the edge (y, y'') to exist, then, since (y', y'') lies in a good pair, we know by Claim 1 that (y, y'') lies in a pair with property (iii); hence it either already exists or will be inserted anyway in the completion process.

In total we have changed less than $6\tilde{\epsilon}n^2 = \frac{\epsilon}{2}n^2$ edges and the new poset now has configuration Q. In order to satisfy the requirements of Lemma 3.1 we have to make sure that it has the same number of relations as in the beginning, which means that we might have to add or remove at most $\frac{\epsilon}{2}n^2$ relations. This can be done as described by Lemma 4.1. □

Acknowledgement

We would like to thank an anonymous referee, whose careful work and suggestions improved the presentation of this paper.

References

[1] Dhar, D. (1978) Entropy and phase transitions in partially ordered sets. *J. Math. Phys.* **19** 1711–1713.
[2] Dhar, D. (1980) Asymptotic enumeration of partially ordered sets. *Pacific J. Math.* **90** 299–305.
[3] Kleitman, D. J. and Rothschild, B. L. (1970) The number of finite topologies. *Proc. Amer. Math. Soc.* **25** 276–282.
[4] Kleitman, D. J. and Rothschild, B. L. (1975) Asymptotic enumeration of partial orders on a finite set. *Trans. Amer. Math. Soc.* **205** 205–220.
[5] Kleitman, D. J. and Rothschild, B. L. (1979) A phase transition on partial orders. *Physica* **96A** 254–259.
[6] Prömel, H. J., Steger, A. and Taraz, A. (1999) Phase transitions in the evolution of partial orders. *J. Combin. Theory Ser. A*, to appear.
[7] Szemerédi, E. (1978) Regular partitions of graphs. In *Problèmes en Combinatoire et Théorie des Graphes*, Proc. Colloque Inter. CNRS (J.-C. Bermond, J.-C. Fournier, M. Las Vergnas, D. Sotteau, eds), CNRS, Paris, pp. 399–401.