

Partitioning permutations into monotone subsequences

David Wörn*

Centre for Mathematical Sciences
University of Cambridge
Cambridge, U.K.

warnd@chalmers.se

Submitted: Feb 25, 2021; Accepted: Jul 10, 2021; Published: Sep 24, 2021

© The author. Released under the CC BY-ND license (International 4.0).

Abstract

A permutation is *k-coverable* if it can be partitioned into k monotone subsequences. Barber conjectured that, for any given permutation, if every subsequence of length $\binom{k+2}{2}$ is k -coverable then the permutation itself is k -coverable. This conjecture, if true, would be best possible.

Our aim in this paper is to disprove this conjecture for all $k \geq 3$. In fact, we show that for any k there are permutations such that every subsequence of length at most $(k/6)^{2.46}$ is k -coverable while the permutation itself is not.

Mathematics Subject Classifications: 05D05, 05A05

1 Introduction

By *permutation of length n* , we mean a permutation π of the set $[n] := \{1, \dots, n\}$ written as a sequence $\pi(1) \pi(2) \dots \pi(n)$. We say a permutation π *contains* another permutation τ , or τ is a pattern of π , if π has a subsequence ordered like τ .

A (possibly empty) subsequence of a permutation is *monotone* if it is either increasing or decreasing. A permutation is k -coverable if it can be partitioned into k monotone subsequences. If a permutation π is k -coverable, then so is every pattern of π . Barber [1] conjectured that, conversely, if every pattern of π of length at most $\binom{k+2}{2}$ is k -coverable, then so is π itself. This has been verified for $k \leq 2$.

We say π is *k-critical* if every proper pattern of π is k -coverable while π itself is not. Barber's conjecture is equivalent to the assertion that critical permutations have length

*Supported by a Trinity College summer studentship.

at most $\binom{k+2}{2}$. A simple inductive argument¹ shows that any k -critical permutation has length at least $\binom{k+2}{2}$. So Barber's conjecture, if true, would be the best possible.

In this paper, we show that the conjecture is false for every $k \geq 3$. In fact the length of a k -critical permutation need not even be $O(k^2)$ – we show that it can be as big as $(k/6)^{2.46}$.

Kézdy, Snevily, and Wang [2] showed that k -critical permutations cannot be arbitrarily long. In fact they showed that the length of a k -critical permutation is at most $k^{O(k^6)}$. Feder and Hell [3] improved this to show that the length is at most $k^{O(k^2)}$. We present a version of their argument. This exponential upper bound is currently the best known.

Wagner [4] showed that the problem of recognising k -coverable permutations (with k given as input) is NP-complete. Thus, unless $\text{NP} = \text{coNP}$, there is generally no simple reason why a permutation is not k -coverable.

Although this is a paper about permutations, the results all have graph-theoretic analogues. Given a permutation π , we can construct a graph – called a permutation graph – with vertex set $[n]$ where $i < j$ are adjacent if $\pi(i) > \pi(j)$. Thus the problem of partitioning a permutation into increasing and decreasing subsequences is a special case of the problem of partitioning the vertex set of a graph into independent sets and cliques sets. This is called *graph cocolouring* [5]. One special property of permutation graphs is that they are perfect. As we will mention in the text, some of our results generalise to perfect graphs, and some to general graphs.

The plan of the paper is as follows. In Section 2 we make some simple remarks, and then in Section 3 we give our first examples that disprove Barber's conjecture. In Section 4, which is really the heart of the paper, we show how to construct long k -critical permutations for general k . In Section 5, we show that Barber's conjecture is true in the special case of separable permutations. In Section 6, we discuss upper bounds. And finally in Section 7 we mention some open problems.

2 Initial remarks

For completeness, we start with the following easy fact. It is equivalent to the assertion that permutation graphs are perfect, and is a special case of Dilworth's theorem.

Lemma 1. *Let π be a permutation, and suppose the length of the longest increasing subsequence of π is s . Then π can be partitioned into s decreasing subsequences.*

Similarly, if the length of the longest decreasing subsequence of π is r , then π can be partitioned into r increasing subsequences.

Proof. By symmetry, it suffices to prove the first part.

For $1 \leq t \leq s$, let D_t be the set of $i \in [n]$ such that the longest increasing subsequence of π ending at $\pi(i)$ has length t . The D_t 's certainly partition $[n]$. Moreover, if $i < j$ and $\pi(i) < \pi(j)$, then we can get an increasing subsequence ending at $\pi(j)$ by appending $\pi(j)$ to an increasing subsequence ending at $\pi(i)$. So each D_t must be decreasing, as needed. \square

¹See comments following Lemma 2.

Note that this result is sharp: we cannot hope to partition π into less than s decreasing subsequences.

We will now use the previous lemma to say something about partitions into monotone subsequences. For a permutation π and nonnegative integers $r, s \in \mathbb{N}$, we say π is (r, s) -coverable if π can be partitioned into r increasing sequences and s decreasing sequences. Thus π is k -coverable iff π is (r, s) -coverable for some $r + s = k$. We say π is (r, s) -critical if every proper pattern is (r, s) -coverable while π itself is not.

We write $D(\pi) \subseteq \mathbb{N}^2$ for the set of pairs (r, s) such that π is *not* (r, s) -coverable.

Lemma 2. *For any permutation π of length n , we have $|D(\pi)| \leq n$.*

For example, if $n < \binom{k+2}{2}$, then $D(\pi)$ cannot contain all of $\{(r, s) \in \mathbb{N}^2 \mid r + s \leq k\}$, so π must be k -coverable. Similarly, if $n < (r + 1)(s + 1)$, then $D(\pi)$ cannot contain all of $\{0 \dots r\} \times \{0 \dots s\}$, so π must be (r, s) -coverable.

Proof. If $n = 0$, then $D(\pi) = \emptyset$. So assume $n \geq 1$. Let I be a longest increasing subsequence of π , and say it has length k . Let τ be the permutation of length $n - k$ ordered like $\pi \setminus I$. Note that if τ is (r, s) -coverable, then π is $(r + 1, s)$ -coverable. Also, by Lemma 1, π is $(0, k)$ -coverable.

From this and the inductive hypothesis we get that $|D(\pi)| \leq |D(\tau)| + k \leq n - k + k = n$, as desired. \square

We remark that the graph-theoretic analogue of Lemma 2 is true for perfect graphs by the same proof, but wildly false for general graphs. Indeed, the largest clique or independent set in a random graph on n vertices is of order $\log n$, and then the analogue of $D(\pi)$ has size on the order of $(n/\log n)^2 \gg n$.

Let us write $C(k)$ for the length of the longest k -critical permutation, and $C(r, s)$ for the length of the longest (r, s) -critical permutation. We will later see that $C(k)$ and $C(r, s)$ are always finite.

We have $C(r, s) \geq (r + 1)(s + 1)$ by Lemma 2. Lemma 1 shows that we get equality $C(r, s) = (r + 1)(s + 1)$ when $r = 0$ or $s = 0$. One can also show that $C(1, 1) = 4 = (1 + 1)(1 + 1)$ as follows. A permutation is $(1, 1)$ -coverable iff its graph is a *split* graph – that is, a graph whose vertex set can be partitioned into a clique and an independent set. A graph is split iff it does not contain a 5-cycle, a 4-cycle, or the complement of a 4-cycle as an induced subgraph [8]. Finally, no permutation graph contains an induced 5-cycle, since it is not perfect, so $C(1, 1) = 4$.

Similarly, we have $C(k) \geq \binom{k+2}{2}$ by Lemma 2. It is easy to see that $C(0) = 1 = \binom{0+2}{2}$ and $C(1) = 3 = \binom{1+2}{2}$. Jørgensen [6] characterised the graph analogues of 2-critical permutations, and from this characterisation we get $C(2) = 6 = \binom{2+2}{2}$. (This again uses the fact that permutation graphs do not contain induced long odd cycles.)

As we will see, we never get equality in $C(k) \geq \binom{k+2}{2}$ or $C(r, s) \geq (r + 1)(s + 1)$ except for the cases listed above.

3 Specific examples

We now present some specific examples of critical permutations that are a bit longer than expected. We will later use these to construct infinite families of critical permutations, but the exact structure of these starting examples will be irrelevant. These examples and many more were found by computer search.

One particularly beautiful 3-critical permutation, which we will call π_{12} , shows that $C(3) \geq 12$. In numbers, it is 10 5 1 7 11 4 9 2 6 12 8 3, but it is more easily appreciated by drawing the points $(i, \pi(i))$ in the plane, as pictured.




Figure 1: π_{12}

It is clear from the picture that π_{12} has four-fold rotational symmetry. This makes it feasible to check by hand that π_{12} really is 3-critical. We note that π_{12} is not $(0, 3)$ - or $(3, 0)$ -coverable, since it has increasing and decreasing subsequences of length 4. It is not $(1, 2)$ -coverable, since its subsequence 5 1 7 6 12 8 is ordered like 2 1 4 3 6 5 which is not $(1, 2)$ -coverable. Hence by symmetry it is not $(2, 1)$ -coverable either. So it is not 3-coverable. On the other hand, if we remove the term $1 = \pi_{12}(3)$, then we can cover the remaining 11 terms by 5 6 12 and two decreasing subsequences. If we instead remove $5 = \pi_{12}(2)$ or $7 = \pi_{12}(4)$, then we use 1 4 6 12 and two decreasing subsequences. By symmetry, any pattern of length 11 is 3-coverable, so π_{12} is 3-critical.




Figure 2: π_9

Our second example, which we will call π_9 , shows that $C(2, 1) \geq 9$. In numbers, it is 5 2 7 1 6 3 9 8 4. Verifying that π_9 is $(2, 1)$ -critical by hand would be labourious, but it is easy for a computer program.

Our third and final example, which we will call π_{15} , shows that $C(2, 2) \geq 15$. In numbers, it is 12 14 5 10 3 9 1 7 15 13 11 4 2 8 6. This permutation has a clear asymmetry to it: the longest decreasing subsequence has length 6, but the longest increasing subsequence has length 3. The latter property will be useful later.




Figure 3: π_{15}

4 Combining critical permutations

We will now explain the main idea of this paper: combining critical permutations to get new critical permutations. The next lemma is the simplest example of this.

Our basic tools are the direct and skew sums of permutations. Given permutations π , σ of lengths n and m , their direct sum $\pi \oplus \sigma$ of length $n + m$ is given by the sequence $\pi(1) \dots \pi(n), \sigma(1) + n \dots \sigma(m) + n$. Their skew sum $\pi \ominus \sigma$ is similarly given by $\pi(1) + m \dots \pi(n) + m, \sigma(1) \dots \sigma(m)$. For example, $(132) \oplus (21) = 13254$ and $(132) \ominus (21) = 35421$.

Lemma 3. *For any $r_1, r_2, s \in \mathbb{N}$, we have $C(r_1 + r_2 + 1, s) \geq C(r_1, s) + C(r_2, s)$. Similarly, for $r, s_1, s_2 \in \mathbb{N}$, we have $C(r, s_1 + s_2 + 1) \geq C(r, s_1) + C(r, s_2)$*

Note that we would get equality above if we replaced $C(r, s)$ with $(r + 1)(s + 1)$ throughout.

Proof. By symmetry, it suffices to prove the first part.

Given an (r_1, s) -critical permutation π and an (r_2, s) -critical permutation σ , we claim that $\pi \ominus \sigma$ is $(r_1 + r_2 + 1, s)$ -critical. The lemma then follows by taking π, σ as long as possible.

First, suppose $\pi \ominus \sigma$ were $(r_1 + r_2 + 1, s)$ -coverable and consider such a covering. Each increasing subsequence of $\pi \ominus \sigma$ lies entirely in either π or σ . This covering must have at least $r_1 + 1$ increasing subsequences lying in π , since π is not (r_1, s) -coverable, and similarly at least $r_2 + 1$ increasing subsequences lying entirely in σ . This is a contradiction as $(r_1 + 1) + (r_2 + 1) > r_1 + r_2 + 1$.

Now consider a proper pattern of $\pi \ominus \sigma$. Without loss of generality it is of the form $\tau \ominus \sigma$ where τ is a proper pattern of π . By assumption, τ is (r_1, s) -coverable. It remains to check that σ is $(r_2 + 1, s)$ -coverable. Pick some term in σ . By assumption, σ minus this term is (r_2, s) -coverable, so by putting in the term as an increasing subsequence of length 1, σ is $(r_2 + 1, s)$ -coverable, as needed. \square

Starting from the fact that $C(2, 1) \geq 9$, repeated application of Lemma 3 gives $C(3r - 1, 2s - 1) \geq 9(r + 1)(s + 1)$, and so $C(r, s) \geq 1.5(r + 1)(s + 1) - O(r + s)$. One similarly obtains $C(r, s) > (r + 1)(s + 1)$ for $r \geq 2, s \geq 1$.

We now look more closely at direct sums, and in particular at $D(\pi \oplus \sigma)$. First note two properties of $D(\pi) \subseteq \mathbb{N}^2$: it is finite, and downward closed, in the sense that if $r' \leq r, s' \leq s$, and $(r, s) \in D(\pi)$, then $(r', s') \in D(\pi)$. We will call a subset of \mathbb{N}^2 with these two properties a *downset*. Our prototypical downset is the ‘triangle’ $T(k) := \{(r, s) \mid r + s \leq k\}$. Note that π is k -coverable iff $T(k) \not\subseteq D(\pi)$.

Given downsets A, B , we let $A \oplus B$ be the downset consisting of (r, s) such that whenever $s = s_1 + s_2$, either $(r, s_1) \in A$ or $(r, s_2) \in B$. We can think of this as merging A and B column by column. We have $|A \oplus B| = |A| + |B|$. Similarly we let $A \ominus B$ consist of (r, s) such that whenever $r = r_1 + r_2$, either $(r_1, s) \in A$ or $(r_2, s) \in B$. Here A and B are merged row by row, and again $|A \ominus B| = |A| + |B|$. The following lemma motivates these definitions.

Lemma 4. *For any permutations π, σ , we have $D(\pi \oplus \sigma) = D(\pi) \oplus D(\sigma)$ and $D(\pi \ominus \sigma) = D(\pi) \ominus D(\sigma)$.*

Proof. For the first part, simply note that $\pi \oplus \sigma$ is (r, s) -coverable iff, for some s_1, s_2 with $s = s_1 + s_2$, π is (r, s_1) -coverable and σ is (r, s_2) -coverable. The second part is proved analogously. \square

Given a permutation π and a downset A , we say π is A -coverable if π is (r, s) -coverable for some $(r, s) \in A$. As before, we say π is A -critical if every proper pattern of π is A -coverable while π itself is not. Thus π is k -coverable iff π is $T(k)$ -coverable, and k -critical iff π is $T(k)$ -critical.

If π is A -critical and moreover $A = D(\pi)$, then we say π is A -minimal. For example, π is $T(k)$ -minimal iff π is k -critical and also (r, s) -coverable whenever $r + s = k + 1$. By inspection, π_{12} is $T(3)$ -minimal.

Let $M(k)$ be the length of the longest $T(k)$ -minimal permutation, so that $M(k) \leq C(k)$ and $M(3) \geq 12$. We introduced the notion of A -minimal permutations in order for the following lemma to work.

Lemma 5. *If π is A -minimal and σ is B -minimal, then $\pi \oplus \sigma$ is $A \oplus B$ -minimal, and $\pi \ominus \sigma$ is $A \ominus B$ -minimal.*

Proof. We have $D(\pi \oplus \sigma) = A \oplus B$ by Lemma 4. Now consider a proper pattern of $\pi \oplus \sigma$, without loss of generality of the form $\tau \oplus \sigma$ with τ a proper pattern of π . By assumption, $D(\tau)$ is a proper subset of A . So $D(\tau \oplus \sigma) = D(\tau) \oplus D(\sigma) = D(\tau) \oplus B$ is a proper subset of $A \oplus B$, as needed. \square

The following lemma lets us prove lower bounds on $M(k)$ and hence on $C(k)$.

Lemma 6. *For any $k \in \mathbb{N}$, we have*

$$\begin{aligned} M(k+1) &\geq M(k) + k + 2, \\ M(2k+2) &\geq 3M(k) + M(k+1), \text{ and} \\ M(2k+3) &\geq M(k) + 3M(k+1). \end{aligned}$$

Note that all three inequalities become equalities if we replace $M(k)$ with $\binom{k+2}{k}$ throughout. For the first inequality, this is immediate. The latter two are illustrated in Figure 4.




Figure 4: Dividing one triangle into four

Proof. For the first inequality, we claim that if π is $T(k)$ -minimal, then $\pi \ominus (1 \dots k+2)$ is $T(k+1)$ -minimal. This follows from Lemma 5 together with the fact that $T(k+1) = T(k) \ominus \{0\} \times \{0 \dots k+1\}$.

Now suppose π is $T(k)$ -minimal and σ is $T(k+1)$ -minimal. We claim that $(\pi \oplus \pi) \ominus (\pi \oplus \sigma)$ is $T(2k+2)$ -minimal and $(\pi \oplus \sigma) \ominus (\sigma \oplus \sigma)$ is $T(2k+3)$ -minimal. By Lemma 5, it suffices to prove that $T(2k+2) = (T(k) \oplus T(k)) \ominus (T(k) \oplus T(k+1))$ and $T(2k+3) = (T(k) \oplus T(k+1)) \ominus (T(k+1) \oplus T(k+1))$.

One can prove the first identity as follows. Let us write $V(A)$ for the sequence obtained from a downset by counting the number of elements in each column, and $H(A)$ for the sequence similarly obtained by counting the number of elements in each row. Then $V(T(k) \oplus T(k)) = 2(k+1), 2k, \dots, 2$ and so $H(T(k) \oplus T(k)) = k+1, k+1, k, k, \dots, 1, 1$. Similarly one finds that $V(T(k) \oplus T(k+1)) = 2k+3, 2k+1, \dots, 1$ and $H(T(k) \oplus T(k+1)) = k+2, k+1, k+1, \dots, 1, 1$. Adding these sequences up one gets $2k+3, 2k+2, \dots, 2, 1 = H(T(k+2))$ as needed. The second identity is similar. \square

We can now show that $C(k)$ is bigger than the lower bound $\binom{k+2}{2}$ by at least a constant factor.

Corollary 7. For any $k \geq 3$, we have $M(k) > 1.07 \binom{k+2}{2}$, and hence $C(k) > 1.07 \binom{k+2}{2}$.

Proof. Starting from $M(3) \geq 12$, the first case of Lemma 6 gives $M(4) \geq 17$, $M(5) \geq 23$, and $M(6) \geq 30$. The third case gives $M(7) \geq 42$. So the desired inequality holds for $3 \leq k \leq 7$. Hence it holds for $k \geq 8$ as well, using the second two cases of Lemma 6 and induction. \square

For the next construction, we need a condition stronger than ‘ (r, s) -critical’ but weaker than ‘ $\{0 \dots r\} \times \{0 \dots s\}$ -minimal’. We say π is (r, s) -sharp if π is (r, s) -critical and $(0, s+1)$ -coverable. We write $S(r, s)$ for the length of the longest (r, s) -sharp permutation, so that $S(r, s) \leq C(r, s)$. By inspection, π_{15} is $(2, 2)$ -sharp, so $S(2, 2) \geq 15$.

We will also need a new way of combining permutations. Given permutations π of length n and σ of length m , their *tensor product* $\pi \otimes \sigma$ has length nm and is given explicitly by $(\pi \otimes \sigma)((i-1)m+j) = (\pi(i)-1)m + \sigma(j)$. We think of this as n copies of σ , ordered like π . For example, $(132) \otimes (21) = 216543$.

Lemma 8. For any positive integers a, b, c, d , we have

$$S(ac-1, bd-1) \geq S(a-1, b-1)S(c-1, d-1).$$

Note that, as before, the inequality becomes an equality if we replace $S(r, s)$ with $(r+1)(s+1)$.

Proof. We will prove that if π is $(a-1, b-1)$ -sharp and σ is $(c-1, d-1)$ -sharp, then $\pi \otimes \sigma$ is $(ac-1, bd-1)$ -sharp.

There are three things to check. In each case, we use the fact that an increasing subsequence of $\pi \otimes \sigma$ is obtained by picking an increasing subsequence I of π and then an increasing subsequence of σ for every copy of σ corresponding to a point in I .

Firstly, we claim that $\pi \otimes \sigma$ is $(0, bd)$ -coverable. Simply pick b decreasing subsequences to cover π and use each one d times to cover the corresponding copies of σ .

Secondly, we claim that $\pi \otimes \sigma$ is not $(ac-1, bd-1)$ -coverable. Suppose we have such a covering Γ . Let C be the set of $i \in [n]$ such that at least c increasing subsequences in Γ pass through the i th copy of σ , and similarly let D be the set of $i \in [n]$ such that at least d decreasing subsequences in Γ pass through the i th copy of σ . Because σ is not $(c-1, d-1)$ -coverable, we have $C \cup D = [n]$. Because π is not $(a-1, b-1)$ -coverable, by Lemma 1 either C has a decreasing subsequence of length a , or D has an increasing subsequence of length b . In the first case, we get a contradiction as we used less than ac increasing subsequences to cover each point of C c times. The second case is analogous.

Finally, we claim that any proper subsequence of $\pi \otimes \sigma$ is $(ac-1, bd-1)$ -coverable. Suppose we remove a term from the i th copy of σ . First pick a $(0, b)$ -covering of π , and make $d-1$ copies of each of these b subsequences. Then pick an $(a-1, b-1)$ -covering of π minus its i th term, making 1 copy of each decreasing sequence and c copies of each increasing sequence. Now each copy of σ except the i th one has either $d-1+1 = d$ decreasing sequences through it, or c increasing sequences and $d-1$ decreasing sequences; either way it can be fully covered. So far we have used $(a-1)c = ac - c$ increasing

sequences and $b(d-1) + b - 1 = bd - 1$ decreasing sequences, so we have $c - 1$ increasing sequences left. This is precisely enough to cover the the i th copy since it already has $d - 1$ decreasing sequences through it and it is missing one term. \square

Starting from $S(3 - 1, 3 - 1) \geq 15$, repeated application of Lemma 8 gives $S(3^n - 1, 3^n - 1) \geq 15^n$, and hence $C(3^n - 1, 3^n - 1) \geq 15^n$. The following lemma lets us use this to get a lower bound on $C(k)$.

Lemma 9. *For any r, s, k , we have $C(r + s) \geq C(r, s)$ and $C(k + 1) > C(k)$.*

Proof. We first show that if τ is (r, s) -critical, then there exists an $(r + s)$ -critical permutation containing τ as a pattern. It follows that $C(r + s) \geq C(r, s)$.

Let L be a skew sum of r copies of $1 \dots N$, and let R be a direct sum of s copies of $N \dots 1$, where $N > r + s$. For any permutation σ , we claim that $(L \oplus \sigma) \ominus R$ is $r + s$ -coverable iff σ is (r, s) -coverable.

If σ is (r, s) -coverable, then we extend the r increasing subsequences to cover L and the s decreasing subsequences to cover R . Conversely, suppose $(L \oplus \sigma) \ominus R$ is $(r + s)$ -coverable. Each of copy of $1 \dots N$ in L must have at least one increasing subsequence through it, since even $r + s$ decreasing subsequences would not be enough. Moreover, the r copies cannot share any increasing subsequences, by construction. So we need at least r increasing subsequences, and similarly at least s decreasing subsequences. So $(L \oplus \sigma) \ominus R$ must in fact be (r, s) -coverable, and hence σ is (r, s) -coverable.

Thus if τ is (r, s) -critical, then $(L \oplus \tau) \ominus R$ contains an $(r + s)$ -critical pattern π , and π must contain τ , as needed.

For the second part, we show that if τ is k -critical, then there exists a $(k + 1)$ -critical permutation containing τ as a proper pattern. The key is to observe that for any permutation σ , the permutation $(1 \dots k + 2) \ominus \sigma$ is $k + 1$ -coverable iff σ is k -coverable. Indeed a $k + 1$ -covering of $(1 \dots k + 2) \ominus \sigma$ must use one increasing sequence to cover $1 \dots k + 2$, and then we are left to k -cover σ .

Thus as before we get a $k + 1$ -critical permutation π containing τ as a pattern. We cannot have $\pi = \tau$ since τ itself is $k + 1$ -coverable, so we get a strict inequality. \square

Corollary 10. *For any n, k , we have*

$$C(2 \cdot 3^n - 2) \geq 15^n \text{ and} \\ C(k) \geq (k/6)^{2.46}$$

Proof. The first inequality follows from $C(2 \cdot 3^n - 2) \geq C(3^n - 1, 3^n - 1)$. For the second inequality, pick n such that $k/6 \leq 3^n \leq k/2$. Now $C(k) \geq C(2 \cdot 3^n - 2) \geq 15^n = (3^n)^{\log_3 15} \geq (k/6)^{2.46}$, as needed. \square

We remark that \oplus , \ominus , and \otimes can be defined on arbitrary graphs. Thus all results in this section remain true in the more general setting of graph cocolouring, with the same proofs, except for Lemma 8 which only remains true for perfect graphs. If we instead consider the problem of partitioning the vertex set of a graph into two sets A and B such that A has no independent set of size $r + 1$ and B has no clique of size $s + 1$, then all the results in this section remain true, including Lemma 8, with similar proofs.

5 Separable permutations

We were able to construct long critical permutations using skew sums, direct sums, and tensor products, but only given some base cases. We now explain this phenomenon.

A permutation is said to be *separable* if it can be obtained from the length-one permutation (1) by repeatedly taking skew sums and direct sums. For example, 2 1 4 3 is separable but 3 1 4 2 is not. It can be shown that $\pi \otimes \sigma$ is separable if π and σ are, by inducting on π and using the right-distributivity $(\pi \oplus \tau) \otimes \sigma = (\pi \otimes \sigma) \oplus (\tau \otimes \sigma)$, $(\pi \ominus \tau) \otimes \sigma = (\pi \otimes \sigma) \ominus (\tau \otimes \sigma)$.

We will classify k -critical and (r, s) -critical separable permutations, showing that they all have lengths $\binom{k+2}{2}$ and $(r+1)(s+1)$, respectively. This explains why our lemmas about \oplus , \ominus , and \otimes alone are not enough to get non-trivial lower bounds on $C(k)$ and $C(r, s)$. The proof crucially relies on our more general notion of being critical with regard to a downset.

Lemma 11. *Given a downset A , if a direct sum $\pi \oplus \sigma$ of permutations is A -critical, then $A = B \oplus C$ for some B, C such that π is B -critical and σ is C -critical.*

Similarly, if a skew sum $\pi \ominus \sigma$ is A -critical, then $A = B \ominus C$ for some B, C such that π is B -critical and σ is C -critical.

Proof. By symmetry, it suffices to prove the first part.

Since $\pi \oplus \sigma$ is not A -coverable, we have $A \subseteq D(\pi \oplus \sigma) = D(\pi) \oplus D(\sigma)$, using Lemma 4. We claim that there exist downsets B, C such that $B \subseteq D(\pi)$, $C \subseteq D(\sigma)$, and $A = B \oplus C$. Suppose the pair (B, C) minimises $|B| + |C|$ subject to the constraints $B \subseteq D(\pi)$, $C \subseteq D(\sigma)$ and $A \subseteq B \oplus C$ (there is such a pair since $(D(\pi), D(\sigma))$ satisfies the constraints). Arguing by contradiction, suppose A is a proper subset of $B \oplus C$ and let $i \in \mathbb{N}$ be maximal such that $V(A)_i < V(B \oplus C)_i$ (using notation from the proof of Lemma 6). Now $V(B)_{i+1} + V(C)_{i+1} = V(B \oplus C)_{i+1} = V(A)_{i+1} \leq V(A)_i < V(B \oplus C)_i = V(B)_i + V(C)_i$, so without loss of generality we may assume $V(B)_{i+1} < V(B)_i$. Thus we get a new downset B' by removing an element from the i th column of B . By construction, we still have $A \subseteq B' \oplus C$, contradicting the minimality of (B, C) .

Now we claim that π is B -critical. We know that π is not B -coverable. Moreover, if $B \subseteq D(\tau)$ for some proper pattern τ of π , then $A = B \oplus C \subseteq D(\tau \oplus \sigma)$ contradicting the assumption that every proper pattern of $\pi \oplus \sigma$ is A -coverable. So π is B -critical. Similarly σ is C -critical, as needed. \square

Corollary 12. *If π is separable and A -critical, then π has length $|A|$. In particular if π is separable and k -critical, then π has length $\binom{k+2}{2}$. If π is separable and (r, s) -critical, then π has length $(r+1)(s+1)$.*

Proof. The first claim follows from Lemma 11 and induction on π . The other claims are special cases of the first. \square

We can say a bit more here. Namely, if π is a separable permutation, then $|D(\pi)|$ is the length of π , and π is A -critical iff $A = D(\pi)$. This gives a procedure to generate all

separable A -critical permutations: consider all ways to write A as $B \oplus C$ or $B \ominus C$, and for each one recursively find the corresponding separable critical permutations.

We remark that if $A \oplus B$ is of the form $\{0 \dots r\} \times \{0 \dots s\}$ for some r, s , then so are A and B . Thus if we were only interested in (r, s) -critical separable permutations, then we would not need to consider general downsets. However, in order to understand k -critical separable permutations, we must consider general downsets.

A graph built from the single-vertex graph using the graph analogues of \oplus and \ominus is called a *cograph*. It is immediate that a graph is a cograph iff it corresponds to a separable permutation. Thus Corollary 12 can be thought of as characterising cographs that are critical with regard to k -cocolouring, (r, s) -cocolouring, or more generally A -cocolouring.

6 Upper bounds

We now prove that $C(r, s)$ and $C(k)$ are always finite. The following combinatorial lemma will be key. It can be deduced from the sunflower lemma [7], but we will get a slightly better bound using an ad-hoc argument. The argument we present is essentially the same as the argument used to prove Theorem 3.1 in [3], and similar to the proof of Theorem 1 in [9], but we give a new presentation of it which shows that it is in some sense sharp.

Given two Boolean functions $P : A \rightarrow [2]$, $Q : B \rightarrow [2]$, possibly with different domains, we will say P and Q are d -close if there are at most d points $k \in A \cap B$ where $P(k) \neq Q(k)$.

Lemma 13. *Fix $r, d \in \mathbb{N}$ with $r \geq 2$ and let U be any nonempty set. Suppose for every $i \in U$, we have a Boolean function $P_i : U \setminus \{i\} \rightarrow [2]$ such that the P_i are pairwise d -close.*

Then there exists a least integer $N = N(r, d)$ such that if $|U| > N$, then we can find a Boolean function $P : U \rightarrow [2]$ such that for any $S \subseteq U$ of size at most $r + 1$, there exists $i \in U \setminus S$ such that $P(k) = P_i(k)$ for all $k \in S$.

In fact $N(r, d) \leq \min(4r^{d+1}, (4r)^{d/2+1})$.

Proof. We first explain why $N(r, d) \leq 4r^{d+1}$.

Fix an arbitrary $i \in U$. The idea is to extend P_i to all of U and then try change it a bit to get the desired P . We show that this works when U is sufficiently large.

Let $p \in [2]$ be arbitrary. For $T \subseteq U \setminus \{i\}$, define $P_T : U \rightarrow [2]$ by²

$$P_T(j) = \begin{cases} p & \text{if } j = i \\ P_i(j) & \text{if } j \notin T \cup \{i\} \\ 3 - P_i(j) & \text{if } j \in T. \end{cases}$$

We may assume that this P_T does not satisfy the desired properties, i.e. that there exists $S_T \subseteq U$ of size at most $r + 1$ such that no P_k agrees with P_T on all of S_T . We must have $i \in S_T$ or $S_T \cap T \neq \emptyset$ since P_i agrees with itself. For $j \in S_T \setminus (T \cup \{i\})$, say $T \cup \{j\}$ is a *child* of T .

²To be precise, we should write $P_T = P_{p,T}$, $S_T = S_{p,T}$, but we suppress this dependency for now.

Let \mathcal{F}_p be the minimal collection of subsets of $U \setminus \{i\}$ such that $\emptyset \in \mathcal{F}_p$, and if $T \in \mathcal{F}_p$ has size at most d , then all its children are also in \mathcal{F}_p . Let $\mathcal{T}_p = \bigcup_{T \in \mathcal{F}_p} T$, and let $\mathcal{T} = \{i\} \cup \mathcal{T}_1 \cup \mathcal{T}_2$.

We claim that $|\mathcal{T}| \leq 4r^{d+1}$. It suffices to prove that $|\mathcal{T}_p| < 2r^{d+1}$. We have $|\mathcal{T}_p| \leq |\mathcal{F}_p|$, since each new set in \mathcal{F}_p introduces at most one new element. By construction, \mathcal{F}_p has 1 element of size 0, at most r elements of size 1, etc, and at most r^{d+1} elements of size $d+1$. So $|\mathcal{F}_p| \leq 1 + r + \dots + r^{d+1} < 2r^{d+1}$, as desired.

Finally, suppose \mathcal{T} is not all of U . Pick $j \in U \setminus \mathcal{T}$. Consider P_j . Let $p = P_j(i)$. Let $T \in \mathcal{F}_p$ be maximal such that $P_j(k) \neq P_i(k)$ for all $k \in T$. Since P_i and P_j are d -close, T has size at most d . By definition of S_T , we know that $P_j(k) \neq P_T(k)$ for some $k \in S_T$. Now this contradicts maximality of T , since the child $T \cup \{k\}$ of T is also in \mathcal{F}_p . This finishes the proof that $N(r, h) \leq 4r^{d+1}$.

Let us now prove that $N(r, h) \leq (4r)^{d/2+1}$. We may assume that, for some i, j , the functions P_i and P_j differ in exactly d points, since otherwise we could take d smaller. Fix such i and j and write D for the set where they differ. Given $Q : D \rightarrow [2]$, we have that Q differs with P_i or P_j in at least $d/2$ points. Say Q differs with P_i in at least $d/2$ points. As before, for $p \in [2]$ and $T \subseteq U \setminus (D \cup \{i\})$, we define $P_T : U \rightarrow [2]$ by $P_T(k) = Q(k)$ for $k \in D$, $P_T(i) = p$, and $P_T(k) = P_i(k)$ iff $k \notin T$ for $k \in U \setminus (D \cup \{i\})$. The rest of the proof proceeds as before: for every Q and p we construct a tree of sets T where this time the maximum size allowed is $d/2 + 1$ instead of $d+1$. The set $\mathcal{T} \subseteq U$ we end up with has size at most $d+2 + 2^{d+1}(r + r^2 + \dots + r^{\lfloor d/2 \rfloor + 1}) \leq 2^{d+1} \cdot 2r^{d/2+1} = (4r)^{d/2+1}$, as needed. \square

In the other direction, we have $N(r, d) \geq r^{d/2}$. To see this, let U be a rooted tree with $d/2+1$ layers where every non-leaf node has r children. For $i \in U$, define $P_i : U \setminus \{i\} \rightarrow [2]$ by taking $P_i(j) = 2$ iff j lies on the path from the root to i . The P_i are close since each one takes the value 2 at most $d/2$ times. Suppose there were a $P : U \rightarrow [2]$ with the desired properties. Then P must take the value 2 at the root, since all P_i 's do, and 1 at all leaves. We get a contradiction by looking at the deepest vertex where P takes the value 2.

We now explain how Lemma 13 relates to permutations.

Corollary 14. *For any r, s with $r \geq s$ and $r \geq 2$, we have $C(r, s) \leq N(r, 2rs)$. Hence $C(r, s) \leq \min(4r^{2rs+1}, (4r)^{rs+1})$.*

Proof. Let π be an (r, s) -critical permutation of length n . We wish to show that $n \leq N(r, 2rs)$. For any $i \in [n]$, π minus its i th term is (r, s) -coverable, so pick an (r, s) -covering. Define $P_i : [n] \setminus \{i\} \rightarrow [2]$ by $P_i(j) = 1$ if j is in an increasing subsequence and 2 if it is in a decreasing subsequence. Given $i, j \in [n]$, suppose there are $2rs + 1$ points $k \in [n] \setminus \{i, j\}$ where $P_i(k) \neq P_j(k)$. Then we can find a set $S \subseteq [n] \setminus \{i, j\}$ of size $rs + 1$ such that, without loss of generality, $P_i(k) = 1$ and $P_j(k) = 2$ for all $k \in S$. Now S is coverable by r increasing sequences, since P_i comes from an (r, s) -covering, and also by s decreasing sequences, since P_j comes from an (r, s) -covering. This contradicts $|S| > rs$ since some two points must end up in the same increasing sequence and the same decreasing sequence. So the P_i are all $(2rs)$ -close.

We now argue by contradiction. Suppose $n > N(r, 2rs)$ and let $P : [n] \rightarrow [2]$ be the Boolean function guaranteed by Lemma 13. By construction, there can be no decreasing subsequence of length $r + 1$ in $\{k \in [n] \mid P(k) = 1\}$ and also no increasing subsequence of length $s + 1$ in $\{k \in [n] \mid P(k) = 2\}$. Thus by Lemma 1, π is (r, s) -coverable, contradicting the assumption that it is (r, s) -critical. \square

We remark that the above argument applies not only to permutations, but more generally to cocolouring perfect graphs.

Now that we have upper bounds on $C(r, s)$, finding upper bounds on $C(k)$ is easy.

Proposition 15. *For any k , we have $C(k) \leq \sum_{r+s=k} C(r, s)$. Hence, for $k \geq 3$, we have $C(k) \leq (4k)^{k^2/4+2}$.*

Proof. For the first part, suppose π is k -critical. Given $r + s = k$, π cannot be (r, s) -coverable, so it contains some (r, s) -critical pattern $\tau_{r,s}$. Now $\tau_{r,s}$ has length at most $C(r, s)$, so π has a pattern σ of length at most $\sum_{r+s=k} C(r, s)$ which contains each $\tau_{r,s}$ as a pattern. This σ cannot be k -coverable, since it is not (r, s) -coverable for any r, s , so it must be all of π , as needed.

The second part follows from the first together with Corollary 14 and the bound $rs \leq k^2/4$. \square

Together with Lemma 9, this shows that $C(k)$ and $C(r, s)$ have roughly the same order of growth.

7 Open questions

This work leaves many questions unanswered. The main question is whether $C(r, s)$ and $C(k)$ grow polynomially or exponentially.

We also do not know any values of $C(r, s)$ or $C(k)$ exactly except for the small values mentioned in Section 2. In particular, the values of $C(2, 1)$ and $C(3)$ are not known. There is no $(2, 1)$ -critical permutation of length 10, 11, or 12, by exhaustive search, so it seems reasonable to believe that $C(2, 1) = 9$; in fact this is conjectured already in [2]. Careful analysis of the proof of Lemma 13 gives an upper bound $C(2, 1) \leq 94$, which is an order of magnitude bigger than the conjectured value.

Similarly, there is no 3-critical permutation of length 13, so we conjecture that $C(3) = 12$. Careful analysis of Proposition 15 gives an upper bound $C(3) \leq 192$.

Acknowledgements

Many thanks are due to Gheehyun Nahm for suggesting Lemma 8, to Imre Leader for helpful discussions, and to two anonymous reviewers for their thoughtful comments.

References

- [1] B. Barber. *British Combinatorial Conference 2019, open problem session*.
- [2] André E Kézdy, Hunter S. Snevily, and Chi Wang. Partitioning permutations into increasing and decreasing subsequences. *Journal of Combinatorial Theory, Series A*, 73(2):353–359, 1996.
- [3] Tomás Feder and Pavol Hell. Matrix partitions of perfect graphs. *Discrete Mathematics*, 306(19-20):2450–2460, 2006.
- [4] Klaus Wagner. Monotonic coverings of finite sets. *Elektronische Informationsverarbeitung und Kybernetik*, 20(12):633–639, 1984.
- [5] Linda Lesniak and H. Joseph Straight. The cochromatic number of a graph. *Ars Combinatoria*, 3:39–46, 1977.
- [6] Leif K Jørgensen. Critical 3-cochromatic graphs. *Graphs and Combinatorics*, 11(3):263–266, 1995.
- [7] Ryan Alweiss, Shachar Lovett, Kewen Wu, and Jiapeng Zhang. Improved bounds for the sunflower lemma. *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, 624–630, 2020.
- [8] Alan Gibbons. *Algorithmic graph theory*, 1985.
- [9] András Gyárfás. Generalized split graphs and Ramsey numbers. *Journal of Combinatorial Theory, Series A*, 81(2):255–261, 1998