Parity Permutation Pattern Matching

Virginia Ardévol Martínez[1[0000−0002−3703−2335]], Florian Sikora[1[0000−0003−2670−6258]], and Stéphane Vialette[2[0000−0003−2308−6970]]

[1] Université Paris-Dauphine, PSL University, CNRS, LAMSADE, 75016 Paris, France {virginia.ardevol-martinez,florian.sikora}@dauphine.fr
[2] LIGM, CNRS, Univ Gustave Eiffel, F77454 Marne-la-Vallée, France stephane.vialette@univ-eiffel.fr

Abstract. Given two permutations, a pattern σ and a text π, Parity Permutation Pattern Matching asks whether there exists a parity and order preserving embedding of σ into π. While it is known that Permutation Pattern Matching is in FPT, we show that adding the parity constraint to the problem makes it W[1]-hard, even for alternating permutations or for 4321-avoiding patterns. However, it remains in FPT if the text avoids a fixed permutation, thanks to a recent meta-theorem on twin-width. On the other hand, as for the classical version, Parity Permutation Pattern Matching remains polynomial-time solvable when both permutations are separable, or if both are 321-avoiding, but NP-hard if the pattern is 321-avoiding and the text is 4321-avoiding.

Keywords: Permutation Pattern Matching · Fixed Parameter Tractability · Parameterized hardness · NP-hardness

1 Introduction

Permutations are one of the most fundamental objects in discrete mathematics, and in concrete, deciding if a permutation contains another permutation as a pattern is one of the most natural decision problems related to them. More precisely, in the well-known problem Permutation Pattern Matching (PPM), given two permutations σ and π, the task is to determine if σ is a pattern of π, or equivalently, if π contains a subsequence which is order-isomorphic to σ. For example, if $\pi = 3\,1\,5\,4\,2$, it contains $\sigma = 2\,3\,1$, as $3\,5\,2$ is a subsequence of π with the same relative order as σ, but π does not contain $\sigma = 1\,2\,3$, as there are no 3 increasing elements in π. In the latter case, we say that π *avoids* $1\,2\,3$. The notion of avoidance allows to define classes of permutations as sets of permutations that avoid certain patterns, for example, 321-avoiding permutations, or $(2413, 3142)$-avoiding permutations, which are known as *separable permutations*.

PPM was proven to be NP-complete by Bose, Buss, and Lubiw in 1998 [6]. This motivated the search for exact exponential time algorithms [1,4,13,8]. However, some special cases, such as Longest Increasing Subequence, or the cases where both σ and π are separable or 321-avoiding, are known to be polynomial time solvable [10,6,16,2]. In fact, it was shown in [17] that PPM is always polynomial-time solvable if the pattern avoids any fixed permutation $\tau \in \{1, 12, 21, 132, 231, 312, 213\}$, and NP-complete otherwise. This result was then extended in [18].

Its parameterized complexity was open for a long time, with a series of partial results, but a breakthrough result of Guillemot and Marx showed that it is fixed parameter tractable when parameterized by the size of the pattern σ, using a new *width measure* structure theory of permutations [15]. They showed that the problem can be solved in time $2^{\mathcal{O}(k^2 \log k)}n$, and later on, Fox improved the running time of the algorithm by removing a factor $\log k$ from the exponent [12].

This led to the question of whether a graph-theoretic generalization of their permutation parameter could exist, that was answered positively in [5], by introducing the notion of *twin-width*, which has proven huge success recently. They showed that graphs of bounded twin-width define a very natural class with respect to computational complexity, as FO model checking becomes linear in them.

Pattern matching for permutations, together with its many variants, has been widely studied in the literature (the best general reference is [19], see also [7]).Here we introduce a natural variation of PPM, which we call Parity Permutation Pattern Matching, and that incorporates the additional constraint that the elements of σ have to map to elements of π with the same parity, i.e., even (resp. odd) elements of σ have to be mapped to even (resp. odd) elements of π. For one thing, pattern avoidance with additional constraints [3,9], including parity restrictions [20,14], has emerged as a promising research area. For another, Parity Permutation Pattern Matching aims at providing concrete use cases of the 2-colored extension of PPM introduced in [16]. We show that, surprisingly, it does not fit into the twin-width framework, and this increases the complexity of the problem, as it becomes W[1]-hard parameterized by the length of the pattern.

In fact, the approach used by Guillemot and Marx [15] to prove that PPM is FPT is based on a result that states that given a permutation π, there exists a polynomial time algorithm that either finds an $r \times r$-grid of π or determines that the permutation has bounded width (and returns the merge sequence of the decomposition, which is used to solve the PPM problem in FPT time). This win-win approach works because, if π contains an $r \times r$-grid, it's not hard to see that it contains every possible pattern σ. However, this cannot be generalized to Parity PPM, as here we have no information on the parity of the elements of the grid, and thus, it is not guaranteed that every pattern maps via a parity respecting embedding into the grid.

Structure of the paper The paper is organized as follows. Section 2 briefly introduces the necessary concepts and definitions. In Section 3, we study the parameterized complexity of Parity PPM, showing that it is harder than PPM in general, but that it remains in FPT for some cases, namely when the *twin-width* of the host permutation is bounded. Finally, in Section 4, we show that concerning the classical P vs NP questions, Parity PPM is similar to PPM. A summary of the complexity of the problems is given in Table 1.

	PPM	PARITY PPM
General case	NP-hard, FPT	W[1]-hard
Separable permutations	P	P
321-av σ and 321-av π	P	P
321-av σ and 4321-av π	NP-hard	NP-hard
4321-av σ	FPT	W[1]-hard
Alternating π and σ	FPT	W[1]-hard
π is fixed pattern avoiding	FPT	FPT

Table 1. Summary of known results (for PPM) and our results (for PARITY PPM).

Due to space constraints, some proofs (marked with ($\star$)) are deferred to the full version of this paper.

2 Preliminaries

Let $[n] = \{1, \dots, n\}$. A permutation of length n is a bijection $f : [n] \longrightarrow [n]$. Given two permutations $\sigma \in S_k$ and $\pi \in S_n$, we say that π (the text, or the host) *contains* σ (the pattern) if there is an embedding from σ into π, i.e., an injective function f such that for every pair of elements x and y of σ, their images $f(x)$ and $f(y)$ of π are in the same relative order as x and y. Otherwise, we say that π *avoids* σ. If π contains σ, we write $\sigma \leq \pi$.

A *permutation class* is a set $\mathcal{C}$ of permutations such that for every permutation $\pi \in \mathcal{C}$, every pattern of π is also contained in $\mathcal{C}$. Every permutation class can be defined by the minimal set of permutations that do not lie inside it, and we define this as $\mathcal{C} = \mathrm{Av}(B)$, where B is the minimal set of avoided permutations.

In this manner, we can define the class $\mathrm{Av}(4321)$, which is the set of permutations that avoid 4321, $\mathrm{Av}(321)$, which is the set of permutations that avoid 321, and $\mathrm{Av}(2413, 3142)$, i.e., the class of permutations that avoid both 2413 and 3142. As we mentioned in the introduction, the latter is known as the class of *separable permutations*, and it can also be characterized as the set of permutations that have a separating tree. In other words, a permutation is separable if there exists an ordered binary tree $\mathcal{T}$ in which the elements of the permutation appear in the leaves and such that the descendants of a tree node form a contiguous subset of these elements.

Furthermore, we define the set of *alternating permutations* as the set of permutations $\sigma \in S_n$ such that $\sigma_1 > \sigma_2 < \sigma_3 > \dots$.

The problem of determining whether a fixed pattern is contained in a permutation has been well studied in the literature, and it is referred to as PERMUTATION PATTERN MATCHING. Here, we study a natural variation of PPM, PARITY PERMUTATION PATTERN MATCHING, which we define formally below.

Definition 1. *Given two permutations, a* pattern $\sigma \in S_k$ *and a* text $\pi \in S_n$*, the problem* PERMUTATION PATTERN MATCHING *asks whether* π *contains* σ*.*

Definition 2. *An injective function f from σ to π is a parity respecting embedding if for all elements x and y of σ, $f(x)$ and $f(y)$ are in the same relative order as x and y, and for every element x of σ, $f(x)$ has the same parity as x.*

We say that an occurrence of a pattern σ in a permutation π respects parity if there is a parity respecting embedding of σ into π. Furthermore, if there is an occurrence of σ in π which respects parity, we say that π parity contains σ, and we write $\sigma \leq_P \pi$. Otherwise, we say that π parity avoids σ.

Definition 3. PARITY PPM *is the problem of determining whether given a pattern σ and a text π, there exists a parity respecting embedding of σ into π.*

As a remark, note that if instead of considering the problem PPM with the constraint that even (resp. odd) elements have to map to even (resp. odd) elements, we require that elements in even (resp. odd) indices (positions) map to elements in even (resp. odd) indices, the problem is equivalent. Indeed, σ parity avoids π if and only if σ^{-1} parity index avoids π^{-1}.

For example, the parity+order preserving embedding of $\sigma = 2\,4\,1\,3$ into $\pi = 4\,\mathbf{2}\,7\,\mathbf{6}\,3\,\mathbf{1}\,\mathbf{5}$ yields the parity-index+order-preserving embedding of $\sigma^{-1} = 3\,1\,4\,2$ into $\pi^{-1} = \mathbf{6}\,\mathbf{2}\,5\,1\,\mathbf{7}\,\mathbf{4}\,3$ (occurrences are depicted with bold integers).

To see this, assume that there is a parity respecting embedding of σ into π. Denote by $P_\sigma(i)$ the position in σ of the element with value i and by f the parity respecting injective map between σ and π associated to the embedding. Since $\sigma^{-1} = P_\sigma(1) \dots P_\sigma(k)$, and f respects parity, if $f(i) = j$, both i and j have the same parity, and thus, the indices in the inverses will also have the same parity (by definition, odd elements are placed in odd indices in the inverses, and vice versa). Furthermore, since f is an an embedding, for $i < j$, $\sigma_i < \sigma_j$ if and only if $f(\sigma_i) < f(\sigma_j)$. Thus, $P(\sigma_i)$ is to the left of $P(\sigma_j)$ in both σ^{-1} and π^{-1}, and by assumption, we also have $i < j$, so f induces a parity index respecting embedding between the inverses.

In this paper, we focus mainly on the parameterized complexity of the above-mentioned problem. Parameterized complexity allows the classification of NP-hard problems on a finer scale than in the classical setting. Fixed parameter tractable (FPT) algorithms are those with running time $O(f(k)\cdot poly(n))$, where n is the size of the input and f is a computable function that depends only on some well-chosen parameter k. On the other hand, problems for which we believe that there does not exist an algorithm with that running time belong to the W-hierarchy. We refer to [11] for more background on the topic.

3 Parameterized complexity

We already saw in the introduction that PPM is in FPT in general, and why the win-win approach of Guillemot and Marx for the parameterized algorithm for PPM doesn't work for PARITY PPM. We show that this intuition is indeed true, proving that the problem is W[1]-hard. In fact, we prove something stronger, which is that PARITY PPM is W[1]-hard even when restricted to alternating permutations or when the pattern is 4321-avoiding. Note that both results are independent from each other, as alternating permutations and 4321-avoiding permutations are not comparable, but they both imply the $\mathsf{W}[1]$-hardness of the general case.

However, the twin-width framework (on which the parameterized algorithm of Guillemot and Marx is an initial step) will be useful to prove that Parity PPM remains in FPT when the text avoids a fixed pattern.

3.1 Parameterized hardness for alternating permutations

Theorem 4. Parity PPM *is* $\mathsf{W}[1]$*-hard parameterized by the length* k *of the pattern, even for alternating permutations* σ *and* π*.*

Proof. We reduce from k-Clique in general graphs, which is known to be $\mathsf{W}[1]$-hard parameterized by the size of the clique k [11]. Given as input a graph G and a parameter k, k-Clique asks whether G contains a clique of size k. For our reduction, given a graph $G = (V, E)$, with $|V| = n$ and $|E| = m$, and a parameter k, we construct a permutation σ that depends only on the parameter k, and a permutation π, that depends on G, such that there exists a clique of size k in the graph G if and only if there is a parity respecting embedding of σ into π.

Construction We explain the construction of π for a general graph G. The high-level idea is to construct different gadgets to represent the vertices and the edges of the graph, and to somehow *link* each edge gadget to the corresponding vertex gadgets, that is, we *link* the gadget associated to edge (u, v) with the gadgets associated to vertices u and v by placing elements of value greater than the minimum element of each vertex gadget and smaller than the maximum element of each vertex gadget between the elements of the edge gadget.

We define the following gadgets (see also Figure 1):

- A *vertex gadget* $\pi[V]$, which is a direct sum of n decreasing permutations, all order-isomorphic to 21 and composed of odd elements. It contains $2n$ elements and starts at element $8m + 3$.
- An *edge gadget* $\pi[E]$, which is a direct sum of m permutations, all order-isomorphic to 435261 and formed by odd elements. It contains $6m$ elements and starts at element 3.
- The *separator gadget* is composed of the four even integers $4(n + 3m) + 4$, $8m + 2$, $4(n + 3m) + 2$ and 2 (and hence is order-isomorphic to 4231). The *separator gadget* lies between the *vertex gadget* and the *edge gadget*.
- Let Even be the $2(n + 3m) - 2$ even integers between 4 and $4(n + 3m)$ that do not appear in the *vertex gadget*, the *separator gadget* or the *edge gadget*. The *even garbage gadget* is the alternating sequence composed of the even integers of Even. It is constructed recursively from left to right as follows: place (and remove from Even) the maximum of Even, place (and remove from Even) the minimum of Even and recurse. It is placed to the right of the edge gadget.

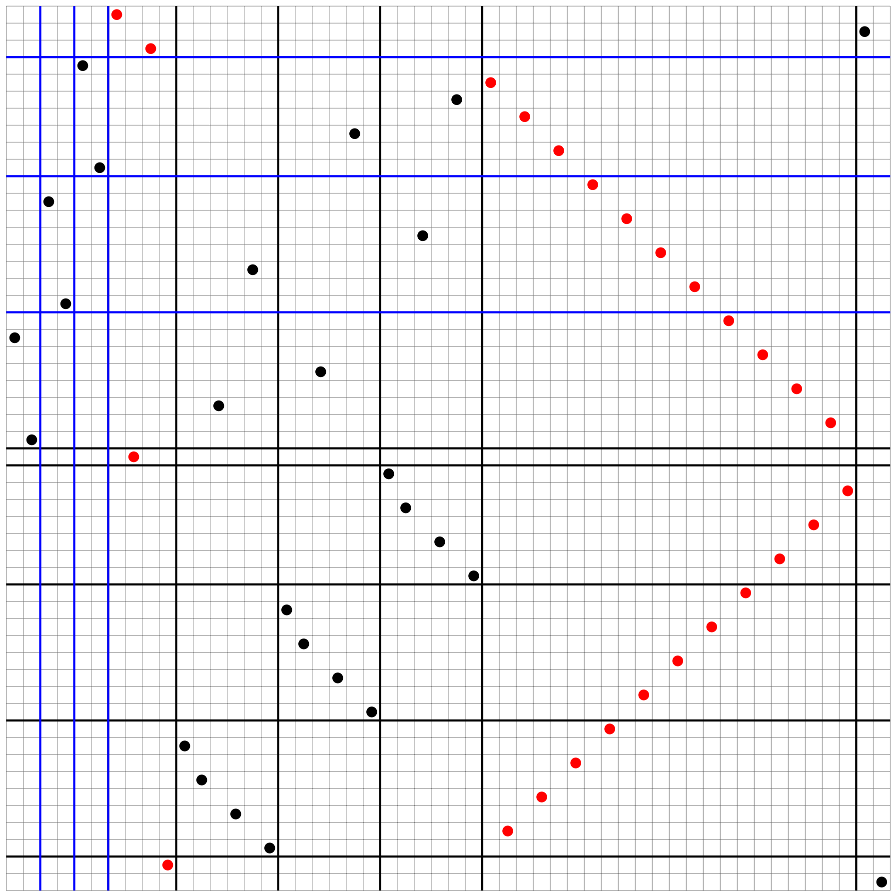

Fig. 1. Illustration of the construction introduced in the proof of Theorem 4. The permutation in the figure corresponds to a clique of size 3 with vertices v_1, v_2, v_3 and edges (v_1, v_2), (v_1, v_3), (v_2, v_3). The odd elements are represented in black while even elements are colored in red. Furthermore, blue lines delimit the vertex boxes.

– Let Odd be the two odd elements $4(n+3m)+3$ and 1 that do not appear in the *vertex gadget*, the *separator gadget* or the *edge gadget*. The *odd garbage gadget* is the decreasing sequence composed of the two odd integers of Odd. It is constructed as the *even garbage gadget* and placed directly to its right.

Formally, define,

$$\forall v_i \in V,\ \pi[v_i] = \boxed{8m+2+2\sum_{j<i}(deg(v_j)+2)+2\sum_i(deg(v_i))+3} \quad \boxed{8m+2+2\sum_{j<i}(deg(v_j)+2)+1} \tag{1}$$

$$\forall e_k = (i,j) \in E,\ \pi[e_k] = \boxed{6k+1}\boxed{6k-1}\boxed{8m+2+2\sum_{j'<i}(deg(v'_j)+2)+2\sum_{(i,j')/j'<j}(1)+3}$$
$$\boxed{6k-3}\boxed{8m+2+2\sum_{j'<j}(deg(v'_j)+2)+2\sum_{(i',j)/i'<i}(1)+3}\boxed{6k-5} \quad (2)$$

$$\pi = \boxed{\pi[v_1]}\dots\boxed{\pi[v_n]}\boxed{4(n+3m)+4}\boxed{4(n+3m)+2}\boxed{8m+2}\boxed{2}$$
$$\boxed{\pi[e_1]},\dots\boxed{\pi[e_m]},\boxed{\text{EVEN}}\boxed{\text{ODD}} \quad (3)$$

(boxes are used for readability purposes only).

The permutation σ is constructed as the permutation π but considering K_k as the graph G.

Clearly, this construction can be carried out in polynomial time and σ depends only on the parameter k, i.e., the new parameter $|\sigma|$ is a function of k. Furthermore, both σ and π are alternating permutations. We claim that there exists a clique of size k in the graph G if and only if there is a parity respecting embedding of σ into π.

Notation Before proving this reduction, we need to define some notation for the elements of the permutations.

Let us denote by w_i and w'_i ($i \in \{1,2,3,4\}$) the four even elements of the separator gadget, placed in between the vertex gadget and the edge gadget of σ and π, respectively.

For each vertex v_i, with $i \in \{1,\dots,n\}$, we will refer to the decreasing subsequence of length two associated to it, $\sigma[v_i]$, as the vertex box associated to v_i.

For each edge e_i, with $i \in \{1,\dots,m\}$, we will refer to the decreasing subsequence of length four associated to it (i.e., the elements of $\sigma[e_i]$ which correspond to 4321 in the permutation 435261) as the edge box of e_i.

We will denote the elements of the vertex box associated to vertex v_i as $v_{i,1}$ and $v_{i,2}$, from left to right (i.e., $v_{i,1} > v_{i,2}$), and the elements of the edge box associated to edge e_i as $e_{i,1}$, $e_{i,2}$, $e_{i,3}$ and $e_{i,4}$, again from left to right. On the other hand, for each edge, we denote the two elements placed in between $e_{i,2}$ and $e_{i,3}$, and between $e_{i,3}$ and $e_{i,4}$, as $h_{i,1}$ and $h_{i,2}$, respectively, where here $h_{i,1} < h_{i,2}$ (these are the elements that correspond to the subsequence 56 in $\sigma[e_i]$).

Finally, the even elements to the right of the edge gadget placed below w_2 are referred to as $w_{i,1}$, $w_{i,2}$, $w_{i,3}$ and $w_{i,4}$, for every edge $i \in \{1,\dots,m\}$, where $w_{i,1}$ is the element $e_{i,4}+1$, $w_{i,2}$ is $e_{i,3}+1$, $w_{i,3}$ is $e_{i,2}+1$, and $w_{i,4}$ is $e_{i,1}+1$. Note that $w_{i,4}$ is not defined for the last edge. On the other hand, the even elements to the right of the edge gadget placed above w_2 are denoted as $x_{i,t}$, for every vertex $i \in \{1,...,n\}$ and every edge incident to v_i, $t \in \{1,...,m_i\}$ ($x_{i,t} = h_{x,y}+1$ for some pair x,y).

Furthermore, we denote by $x_{i,0}$ and x_{i,m_i+1} the even elements in the extremes such that $x_{i,0} = v_{i,2}+1$ and $x_{i,m_i+1} = v_{i,1}+1$. Again, note that x_{n,m_n+1} is not defined.

For the elements of π, we follow an analogous notation denoting the elements by $v'_{i,1}$, $e'_{i,1}$, etc.

Direct Implication

Claim 5. ($\star$)*If there exists a clique of size k in the graph G, then there is a parity respecting embedding of σ into π.*

Reverse implication Suppose now that there exists a parity respecting embedding between σ and π and let f be the associated injective mapping. We want to show that we have enough structure in the permutations to infer that there must be a clique of size k in the graph G. In order to do so, we will prove the following sequence of claims that will restrict the map f.

Claim 6. *Any parity respecting embedding f from σ to π must map w_i to w'_i, for $i \in \{1,2,3,4\}$.*

Proof of claim. Since the pattern matching needs to respect parity, f must map the w_i's to even elements of π. Towards a contradiction, assume first that $f(w_i) \neq w'_j$, $i,j \in \{1,2,3,4\}$. That means that $f(w_i) = w'_{i',j}$ or $x'_{i',t}$, for some indices i',j or i',t. But then, the odd elements to the right of w_i in σ cannot map to elements to the right of $f(w_i)$ in π (as there would be at most 2 odd elements to the right of $f(w_i)$ and there are strictly more than 2 odd elements to the right of w_i), so f cannot be an embedding of σ into π. Finally, since both the w_i's and the w'_i's form 4231 subsequences, it is clear that there exists a unique way to embed the w_i's into the w'_i's, which is mapping each w_i to its corresponding w'_i, for every $i \in \{1,2,3,4\}$. Thus, if $f(w_i) \neq w'_i$, f cannot be an embedding. $\lhd$

Claim 7. *All the elements to the left (resp. to the right) of the w_i's in σ map to elements to the left (resp. to the right) of w'_i's in π. Similarly, the elements above (resp. below) w_2 in σ map to elements above (resp. below) w'_2 in π.*

Proof of claim. This is a direct corollary of Claim 6. $\lhd$

Claim 8. *Any parity respecting embedding f from σ to π must map vertex blocks of σ to vertex blocks of π.*

Proof of claim. By Claim 7, since elements to the left of w_2 in σ map to elements to the left of w'_2 in π, we have that $f(v_{i,j}) = v'_{i',j'}$, for $i \in \{1,...,k\}, i' \in \{1,...,n\}$ and $j,j' \in \{1,2\}$. Assume that $f(v_{i,1}) = v'_{i',j'}$ and $f(v_{i,2}) = v'_{i'',j''}$, with $i' \neq i''$. Since $v_{i,1}$ is to the left of $v_{i,2}$, it means that f must map $v_{i,2}$ to an element placed to the right of $f(v_{i,1}) = v'_{i',j'}$. But $v_{i,1} > v_{i,2}$ and every element which is to the right of $v'_{i',j'}$ and which does not belong to the vertex block of $v'_{i'}$, is greater than $v'_{i',j'}$. Thus, if $i'' \neq i'$, then f would not be an embedding. $\lhd$

Claim 9. *Any parity respecting embedding f from σ to π must map edge blocks of σ to edge blocks of π.*

Proof of claim. Again, we have that $f(e_{i,j}) = e'_{i',j'}$ for some pair i',j', and since the structure of the gadget has the same properties as the vertex gadget, we can use the same argument as in the proof of Claim 8. $\lhd$

Claim 10. *Any parity respecting embedding f from σ to π must map $h_{i,j}$ to $h'_{i',j}$, where $e'_{i'}$ is the edge associated to the edge block where $e_{i,1}$ maps to.*

Proof of claim. By Claim 7, we have that necessarily, $f(e_{i,j}) = e'_{i',j}$, for $i \in \{1,...,l\}, i' \in \{1,...,m\}$ and $j \in \{1,2,3,4\}$.

First, since $f(e_{i,2}) = e'_{i',2}$ and $f(e_{i,3}) = e'_{i',3}$, and f is an embedding, the fact that $h_{i,1}$ is in between $e_{i,2}$ and $e_{i,3}$ implies that it must map to an element between $e'_{i',2}$ and $e'_{i',3}$. Similarly, $h_{i,2}$ must map to an element in between $e'_{i',3}$ and $e'_{i',4}$. Since edge blocks map to edge blocks, there is at most one element that satisfies each of these conditions. And these elements are $h'_{i',1}$ and $h'_{i',2}$, respectively. $\triangleleft$

Claim 11. *All the even elements to the right of the edge gadgets in σ must map to even elements to the right of the edge gadgets in π.*

Proof of claim. This follows from Claim 6. Since $f(w_i) = w'_i$ for $i \in \{1,2,3,4\}$ and f has to respect parity, the rest of the even elements cannot map anywhere else. $\triangleleft$

Now, suppose that there is a parity respecting embedding f of σ into π and assume, towards a contradiction, that G does not contain a clique of size k. Since there is no clique of size k, it means that we cannot have $l = \binom{k}{2}$ edges between the k vertices of G which are in the image of f (that is, the vertices associated to the images of the k vertex boxes of σ).

We know that the k vertex blocks of σ map to k vertex blocks in π and the $\binom{k}{2}$ edge blocks of σ map to $\binom{k}{2}$ edge blocks of π. Since G does not contain a clique, one of the k vertices corresponding to the k vertex blocks in the image of f will have degree strictly smaller than $k-1$ when we restrict G to the k selected vertices. Let i' be the vertex with degree strictly smaller than $k-1$ and suppose it is the image of vertex block i in σ. Then, there are two possible cases. The first case is that in the image of f, between the values $f(v_{i,1})$ and $f(v_{i,2})$, there are less than k odd elements (these elements are necessarily of the form $h'_{i,j}$). Since in between $v_{i,1}$ and $v_{i,2}$ in σ there are k odd elements of the form $h_{i,j}$, this would imply that f cannot be a parity respecting embedding. The second possibility is that in between the values $f(v_{i,1})$ and $f(v_{i,2})$ there are k odd elements (which again are necessarily of the form $h'_{i,j}$) but one of them is not in between $f(e_{l,2})$ and $f(e_{l,3})$, or $f(e_{l,3})$ and $f(e_{l,4})$, for some $l \in \{1,\ldots,m\}$. This would also contradict the fact that f is a parity respecting isomorphism, as all the $h_{i,j}$ in σ are between some pair $e_{l,2}, e_{l,3}$, or $e_{l,3}, e_{l,4}$ (with respect to the x-axis). Therefore, if there is a parity respecting embedding of σ into π, it must map the k vertex boxes of σ into k vertex boxes of π associated to k vertices that form a clique in G. $\square$

Corollary 12. ($\star$)*Given a pattern $\sigma \in S_k$ and a text $\pi \in S_n$,* Parity PPM *cannot be solved in time $f(k) \cdot n^{o(\sqrt{k})}$ for any computable function f, under the Exponential Time Hypothesis (ETH).*

Note that reducing from Subgraph Isomorphism instead of k-Clique in the proof of Theorem 4 to get a better lower bound under the ETH is not trivial since there is a notion of order of the pattern in Parity PPM (i.e., two isomorphic subgraphs can result in different permutations depending on the ordering of their vertices).

3.2 Parameterized hardness for 4321-avoiding patterns

In this subsection, we complement the previous hardness result by showing that the problem remains hard for patterns belonging to the class of 4321-avoiding permutations. Our proof uses a colored version of PPM defined in [16], proven W[1]-hard parameterized by $k = |\sigma|$ in [16].

Definition 13. 2-colored 2IPP *(*2 Increasing Permutation Pattern*) consists on, given a 321-avoiding permutation $\sigma \in S_k$ and an arbitrary permutation π such that both σ and π are 2-colored permutations, finding a color-preserving embedding of σ into π.*

Theorem 14. ($\star$) Parity PPM *is* W[1]*-hard parameterized by the length k of the pattern, even if the pattern is 4321-avoiding.*

3.3 Parameterized algorithm for fixed pattern avoiding text

In the previous subsection, we showed that restricting the pattern does not necessarily reduce the complexity of the problem. However, we now see that restricting the text allows us to use the twin-width meta-theorem [5] to have a positive result. In fact, to see that Parity PPM is FPT if the text avoids a fixed pattern x, it suffices to show that we can describe the problem using first-order (FO) logic, i.e., that we can express it as a formula which uses quantified variables over non-logical objects, and sentences (formulas without free variables) that contain the variables. Indeed, adding unary relations to mark the odd and even values preserves bounded twin-width, and therefore FPT tractability. The result follows from [5]:

Lemma 15 ([5]). *FO model checking is* FPT *on every hereditary proper subclass of permutation graphs.*

This implies that FO model checking is FPT in the class of permutations avoiding a fixed pattern. Here, FO model checking refers to the problem of, given a first-order sentence ϕ of FO and a finite model $\mathcal{M}$ of FO (which specifies the domain of disclosure of the variables), deciding whether $\mathcal{M}$ satisfies ϕ, i.e., whether there exists an assignment of the variables which respects the domain imposed by $\mathcal{M}$ and that satisfies ϕ. Therefore, we can state the following theorem:

Theorem 16. ($\star$)Parity PPM *is in* FPT *if the text π avoids a fixed permutation.*

4 Classical complexity

Even though Parity PPM is harder than PPM from the parameterized point of view, we will show that this is not the case concerning its classical complexity.

4.1 Hardness

A nice quite recent result showed that PPM remains NP-hard, even if the pattern is 321-avoiding and the text is 4321-avoiding [17]. In the following, we show that it remains true for Parity PPM.

Theorem 17. ($\star$)Parity PPM *is* NP*-hard, even if* σ *is a* 321*-avoiding permutation and* π *is a* 4321*-avoiding permutation.*

4.2 Polynomial-time solvable cases

For some specific cases of Permutation Pattern Matching, polynomial time algorithms that solve the problem exactly have been proposed. Here, we show that some of these algorithms can be adapted to solve the problem Parity Permutation Pattern Matching while still running in polynomial time.

Theorem 18. ($\star$)*Let* σ *be a permutation in* S_k *and* π *be a permutation in* S_n*.* Parity PPM *can be solved in polynomial time in the following cases:*

1. *If both permutations are separable. In particular, if both permutations are (231, 213)-avoiding, it can be solved in linear time.*
2. *If both permutations are 321-avoiding.*

Acknowledgements Thanks to Édouard Bonnet and Eun Jung Kim for pointing out the link with the twin-width framework, and to the reviewers for their useful comments.

References

1. Ahal, S., Rabinovich, Y.: On complexity of the subpattern problem. SIAM J Discrete Math **22**(2), 629–649 (2008)
2. Albert, M.H., Lackner, M., Lackner, M., Vatter, V.: The Complexity of Pattern Matching for 321-Avoiding and Skew-Merged Permutations. Discret. Math. Theor. Comput. Sci. **18**(2) (2016)
3. Alexandersson, P., Fufa, S.A., Getachew, F., Qiu, D.: Pattern-avoidance and Fuss-Catalan numbers. arXiv preprint arXiv:2201.08168 (2022)
4. Berendsohn, B.A., Kozma, L., Marx, D.: Finding and Counting Permutations via CSPs. Algorithmica **83**(8), 2552–2577 (2021)
5. Bonnet, É., Kim, E.J., Thomassé, S., Watrigant, R.: Twin-width I: tractable FO model checking. J. ACM **69**(1), 3:1–3:46 (2022)
6. Bose, P., Buss, J.F., Lubiw, A.: Pattern Matching for Permutations. Inf. Process. Lett. **65**(5), 277–283 (1998)

7. Bruner, M.L., Lackner, M.: The computational landscape of permutation patterns. Pure Mathematics and Applications: Special Issue for the Permutation Patterns 2012 Conference **24**(2), 83–101 (2013)
8. Bruner, M., Lackner, M.: A Fast Algorithm for Permutation Pattern Matching Based on Alternating Runs. Algorithmica **75**(1), 84–117 (2016)
9. Bulteau, L., Fertin, G., Jugé, V., Vialette, S.: Permutation Pattern Matching for Doubly Partially Ordered Patterns. In: Proc. of CPM. LIPIcs, vol. 223, pp. 21:1–21:17 (2022)
10. Cormen, T.H., Leiserson, C.E., Rivest, R.L., Stein, C.: Introduction to algorithms. MIT press (2022)
11. Cygan, M., Fomin, F.V., Kowalik, L., Lokshtanov, D., Marx, D., Pilipczuk, M., Pilipczuk, M., Saurabh, S.: Parameterized Algorithms. Springer (2015)
12. Fox, J.: Stanley-wilf limits are typically exponential. arXiv preprint arXiv:1310.8378 (2013)
13. Gawrychowski, P., Rzepecki, M.: Faster Exponential Algorithm for Permutation Pattern Matching. In: 5th SOSA@SODA 2022. pp. 279–284. SIAM (2022)
14. Gil, J.B., Tomasko, J.A.: Restricted Grassmannian permutations. Enumerative Combinatorics and Applications **2**(4), Article #S4PP6 (2021)
15. Guillemot, S., Marx, D.: Finding small patterns in permutations in linear time. In: Proc. of SODA. pp. 82–101. SIAM (2014)
16. Guillemot, S., Vialette, S.: Pattern Matching for 321-Avoiding Permutations. In: Proc. of ISAAC. LNCS, vol. 5878, pp. 1064–1073. Springer (2009)
17. Jelínek, V., Kynčl, J.: Hardness of Permutation Pattern Matching. In: Proc. of SODA. pp. 378–396. SIAM (2017)
18. Jelínek, V., Opler, M., Pekárek, J.: Griddings of Permutations and Hardness of Pattern Matching. In: Proc. of MFCS. LIPIcs, vol. 202, pp. 65:1–65:22 (2021)
19. Kitaev, S.: Patterns in Permutations and Words. Monographs in Theoretical Computer Science. An EATCS Series, Springer (2011)
20. Tanimoto, S.: Combinatorics of the group of parity alternating permutations. Adv. Appl. Math. **44**(3), 225–230 (2010)