



Streaming deletion problems parameterized by vertex cover ^{☆, ☆☆}

Jelle J. Oostveen ^{*}, Erik Jan van Leeuwen

Dept. Information and Computing Sciences, Utrecht University, the Netherlands



ARTICLE INFO

Article history:

Received 4 February 2022

Received in revised form 20 January 2023

Accepted 4 September 2023

Available online 15 September 2023

Communicated by G.F. Italiano

Keywords:

Graph streaming

Parameterized complexity

Vertex cover

Vertex deletion

ABSTRACT

Streaming is a model where an input graph is provided one edge at a time, instead of being able to inspect it at will. In this work, we take a parameterized approach by assuming a vertex cover of the graph is given, building on work of Bishnu et al. [COCOON 2020]. We show the further potency of combining this parameter with the Adjacency List streaming model to obtain results for vertex deletion problems. This includes kernels, parameterized algorithms, and lower bounds for the problems of Π -FREE DELETION, H -FREE DELETION, and the more specific forms of CLUSTER VERTEX DELETION and ODD CYCLE TRANSVERSAL. We focus on the complexity in terms of the number of passes over the input stream, and the memory used. This leads to a pass/memory trade-off, where a different algorithm might be favourable depending on the context and instance. We also discuss implications for parameterized complexity in the non-streaming setting.

© 2023 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Streaming is an algorithmic paradigm to deal with data sets that are too large to fit into main memory [1]. Instead, elements of the data set are inspected in a fixed order¹ and aggregate data is maintained in a small amount of memory (much smaller than the total size of the data set). It is possible to make multiple passes over the data set. The goal is to design algorithms that analyze the data set while minimizing the combination of the number of passes and the required memory. We note that computation time is not measured in this paradigm. Streaming has proved very successful and is extensively studied in many diverse contexts [2,3]. In this work, we focus on the case where the data sets are graphs and the streamed elements are the edges of the graph.

A significant body of work on graph streaming works in the semi-streaming model, where $\tilde{O}(n)$ memory² is allowed, with the aim of limiting the number of necessary passes to one or two. This memory requirement might still be too much for the largest of networks. Unfortunately, many basic problems in graphs require $\Omega(n)$ or even worse space [4,5] to compute

[☆] This article belongs to Section A: Algorithms, automata, complexity and games, Edited by Paul Spirakis.

^{☆☆} This work is based on the master thesis "Parameterized Algorithms in a Streaming Setting" by the first author. This work is partially supported by the NWO grant OCENW.KLEIN.114 (PACAN). An extended abstract of this article appeared in the proceedings of the 23rd International Symposium on Fundamentals of Computation Theory (FCT 2021), LNCS vol. 12867, Springer, 2021, pp. 413–426; see https://doi.org/10.1007/978-3-030-86593-1_29.

^{*} Corresponding author.

E-mail addresses: j.j.oostveen@uu.nl (J.J. Oostveen), e.j.vanleeuwen@uu.nl (E.J. van Leeuwen).

¹ We consider insertion-only streams throughout this paper.

² Throughout this paper, memory is measured in bits. The \tilde{O} notation hides factors polylogarithmic in n . Note that $\mathcal{O}(\log n)$ bits is the space required to store (the identifier of) a single vertex or edge.

in a constant number of passes. Therefore, Fafianie and Kratsch [6] and Chitnis et al. [7] introduced concepts and analysis from parameterized complexity [8] to the streaming paradigm. For example, it can be decided whether a graph has a vertex cover of size at most K using one pass and $\tilde{O}(K^2)$ space, which is optimal. This led to various further works [9–11] and the first systematic study by Chitnis and Cormode [12].

Our work continues this line of research and follows up on recent work by Bishnu et al. [10,13].³ They made two important conceptual contributions. First, they analyzed the complexity of parameterized streaming algorithms in three models that prescribe the order in which the edges arrive in the stream and that are commonly studied in the literature [10, 14,2,15]. The *Edge Arrival* (EA) model prescribes some permutation of all the edges of the graph. The *Vertex Arrival* (VA) requires that the edges appear per vertex: if we have seen the vertices $V' \subseteq V$ already, and the next vertex is w , then the stream contains the edges between w and the vertices in V' . Finally, the *Adjacency List* (AL) gives the most information, as it requires the edges to arrive per vertex, but when vertex v appears in the stream, we also see all edges incident to v . This means we effectively see every edge twice in a single pass, once for both of its endpoints. Due to the nature of streams consisting of edges, we assume no degree-0 vertices are present in our input graphs.

The second and more important contribution of Bishnu et al. [13] was to study the size K of a vertex cover in the graph as a parameter. This has been broadly studied in parameterized complexity (see e.g. the PhD thesis of Jansen [16]). They showed that the very general \mathcal{F} -SUBGRAPH DELETION and \mathcal{F} -MINOR DELETION problems all admit one pass, $\tilde{O}(\Delta(\mathcal{F}) \cdot K^{\Delta(\mathcal{F})+1})$ space streaming algorithms in the AL model, by computing small kernels to which then a straightforward exhaustive algorithm is applied. On the other hand, such generic streaming algorithms are not possible in the EA and VA models, as then (super-) linear lower bounds exist even if the size of a smallest vertex cover is constant [13].

We focus on the induced subgraph version of the vertex deletion problem, parameterized by the size of a vertex cover. Here, Π is a collection of graphs.

Π -FREE DELETION [VC]

Input: A graph G with a vertex cover X , and an integer $\ell \geq 1$.

Parameter: The size $K := |X|$ of a vertex cover.

Question: Is there a set $S \subseteq V(G)$ of size at most ℓ such that $G[V(G) \setminus S]$ does not contain a graph in Π as an induced subgraph?

To avoid triviality,⁴ we assume $K \geq \ell$. We assume the vertex cover is given as input; if only the size is given, we can use one pass and $\tilde{O}(K^2)$ space or 2^K passes and $\tilde{O}(K)$ space to obtain it [7,12] (this does not meaningfully impact our results). The unparameterized version of this problem is well known to be NP-hard [17] for any nontrivial and hereditary property Π . It has also been well studied in the parameterized setting (see e.g. [18–20]). When parameterized by the vertex cover number, it has been studied from the perspective of kernelization: while a polynomial kernel cannot exist in general [21,22], polynomial kernels exist for broad classes of families Π [22,23]. Moreover, the algorithm of Cai [18] implies an FPT algorithm for Π -FREE DELETION [VC] for any fixed Π (but is XP in the maximum size of any graph in Π). As far as we are aware, parameterized algorithms for this parameterization have not been explicitly studied.

In the streaming setting, Chitnis et al. [9] showed for the unparameterized version of this problem in the EA model that any p -pass algorithm needs $\Omega(n/p)$ space if Π satisfies a mild technical constraint. For some Π -FREE DELETION [VC] problems, the results by Bishnu et al. [13] imply single-pass, $\text{poly}(K)$ space streaming algorithms (through their kernel for \mathcal{F} -SUBGRAPH DELETION [VC]) in the AL model and lower bounds in the EA/VA model. They also provide an explicit kernel for CLUSTER VERTEX DELETION [VC] in the AL/EA/VA models. However, this still leaves the streaming complexity of many cases of the Π -FREE DELETION [VC] problem open.

Our contributions. We determine the streaming complexity of the general Π -FREE DELETION [VC] problem. Our main positive result is a unified approach to a single-pass polynomial kernel for Π -FREE DELETION [VC] for a broad class of families Π . In particular, we show that the kernelization algorithms by Fomin et al. [22] and Jansen and Kroon [23] can be adapted to the streaming setting. The kernels of Fomin et al. [22] consider the case when Π can be characterized by few adjacencies, which intuitively means that for any vertex of any member of Π , adding or deleting edges between all but a few (say at most c_Π) distinguished other vertices does not change membership of Π . The exponent of the polynomial kernels depends on c_Π . Jansen and Kroon [23] considered even more general families Π . We show that these kernels can be computed in the AL model using a single pass and polynomial space (where the exponent depends on c_Π). This generalizes the previous results by Bishnu et al. [13] as well as their kernel for \mathcal{F} -SUBGRAPH DELETION [VC].

To complement the kernels, we take a direct approach to find more memory-efficient algorithms, at the cost of using many passes. We show novel parameterized streaming algorithms that require $\tilde{O}(K^2)$ space and $\mathcal{O}(K)^{\mathcal{O}(K)}$ passes. Here, all hidden constants depend on Π and c_Π . Crucially, however, the exponent of the space usage of these algorithms does not, which provides an advantage over computing the kernel. We also provide explicit streaming algorithms for CLUSTER VERTEX DELETION [VC] and ODD CYCLE TRANSVERSAL [VC] that require $\tilde{O}(K)$ space (both) and $2^K K^2$ and 3^K passes respectively, as

³ As the Arxiv version contains more results, we refer to this version from here on.

⁴ Otherwise, removing the entire vertex cover is a trivial solution.

well as streaming algorithms for Π -FREE DELETION $[VC, |V(H)|]$ when $\Pi = \{H\}$ and the problem is parameterized by K and $|V(H)|$. A crucial ingredient to these algorithms is a streaming algorithm that finds induced subgraphs isomorphic to a given graph H . For these algorithms, a Π of large size can lead to large memory use. To this end, we also give algorithms using an oracle to learn information about Π , similar to an algorithm by Cai [18]. Further details are provided in Section 3.

The above results provide a trade-off in the number of passes and memory complexity of the algorithm used. However, we should justify using both the AL model and the parameter vertex cover. To this end, in Section 4, we investigate lower bounds for streaming algorithms for Π -FREE DELETION $[VC]$. The (unparameterized) linear lower bound of Chitnis et al. [9] in the EA model requires that Π contains a graph H for which $|E(H)| \geq 2$ and no proper subgraph is a member of Π . We prove that the lower bound extends to both the VA and AL models, with only small adjustments. Hence, parameterization is necessary to obtain sublinear passes and memory for most Π . Since VERTEX COVER is one of the few natural graph parameters that has efficient parameterized streaming algorithms [6,7], this justifies the use of the vertex cover parameter. We also extend the reductions by Bishnu et al. [13] to general hardness results for Π -FREE DELETION in the VA and EA model when the size of the vertex cover is a constant (dependent on Π), justifying the use of the AL model for most Π .

We also consider the parameterized complexity of H -FREE DELETION $[VC]$ in the non-streaming setting. While polynomial kernels were known in the non-streaming setting [22], we are unaware of any investigation into explicit parameterized algorithms for these problems. We give a general $2^{O(K^2)} \text{poly}(n, |V(H)|)$ time algorithm. This contrasts the situation for H -FREE DELETION parameterized by the treewidth t of the graph, where a $2^{o(t^{|V(H)|-2})} \text{poly}(n, |V(H)|)$ time lower bound is known under the Exponential Time Hypothesis (ETH) [20]. We also construct a graph property Π for which we provide a lower bound of $2^{o(K \log K)} \text{poly}(n, |V(H)|)$ for Π -FREE DELETION $[VC]$ under ETH. Further details are provided in Section 3.

Preliminaries. We work on undirected graphs $G = (V, E)$, where $|V| = n, |E| = m$. We denote an edge $e \in E$ between $v \in V$ and $u \in V$ with $uv \in E$. For a set of vertices $V' \subseteq V$, denote the subgraph induced by V' as $G[V']$. Denote the neighbourhood of a vertex v with $N(v)$ and we write $N[v]$ for $N(v)$ including v , so $N[v] = N(v) \cup \{v\}$. The vertex cover of a graph H is denoted with $VC(H)$.

We denote the parameters of a problem in $[\cdot]$ brackets, a problem A parameterized by vertex cover number and solution size is denoted by $A[VC, \ell]$.

Further related work. Our main algorithm for Π -FREE DELETION $[VC]$ uses a procedure that finds an occurrence of an induced subgraph H in the input graph G . To this end, there is related work on the INDUCED SUBGRAPH ISOMORPHISM problem. In general, finding an induced copy of a graph H of size k can be brute-forced in $\mathcal{O}(n^k)$ time. Nešetřil and Poljak [24] show that finding an induced copy of a graph H can be reduced in time $\mathcal{O}(k^2 n^2)$ to the k -CLIQUE or k -INDEPENDENT SET problem, which makes a k -clique and a k -independent set essentially the hardest patterns to detect. With this in mind, it is logical that knowing a vertex cover of the input graph can help in our approach. Dalirrooyfard et al. [25] show that under the Hadwiger conjecture, the hardness of INDUCED SUBGRAPH ISOMORPHISM can be reduced to the hardness of k -CLIQUE. Also assuming ETH, this implies that for almost all H of size k , finding an H in a graph G cannot be done in time $n^{o(k/\log k)}$.

In terms of upper bounds, the running time of detecting a k -clique is related to fast matrix multiplication [26,24,27]. There is also work on INDUCED SUBGRAPH ISOMORPHISM on specific graph classes [28,29], and on specific small patterns H [30,31,27,32–37]. Eppstein [38] investigates the parameterized complexity of the INDUCED SUBGRAPH ISOMORPHISM problem with the parameter treewidth, and shows an algorithm listing all occurrences of H of size w that contain a vertex of a set $S \subseteq V(G)$ in time $2^{\mathcal{O}(w \log w)} n + \mathcal{O}(kw)$ for graphs of treewidth w , where k is the number of isomorphisms. This result is related to our algorithm in that a vertex cover of size k implies bounded treewidth, and our running time is similar. However, our result is more general because we allow freedom in the size of H as compared to the graph parameter.

1.1. Memory complexity of branching algorithms

In this work, we will see branching algorithms which branch on, for example, including a certain vertex in the solution set. The stated memory complexities can only be attained by making good use of memory. For example, when branching, we re-use the branch set already in memory, and when returning out of recursion, we select again the subset of memory which (still) corresponds to the memory of that branch. When returning from recursion to continue with another branching option, it might also pose a memory problem to have in memory what these branching options are (with b branching options along k recursion steps this requires $\tilde{\mathcal{O}}(bk)$ bits of memory). To overcome this, we can recompute the branching options when returning out of recursion, increasing the time complexity or number of passes by a factor b . Such memory complexity tricks are also used in [11, Appendix B.1].

2. Adapting existing kernels

We first show that very general kernels for vertex cover parameterization admit straightforward adaptations to the AL streaming model. The kernels considered are those by Fomin et al. [22] and by Jansen and Kroon [23]. Fomin et al. [22] provide general kernelization theorems that make extensive use of a single property, namely that some graph properties can be characterized by few adjacencies.

Definition 1. ([22, Definition 3]) A graph property Π is *characterized by $c_\Pi \in \mathbb{N}$ adjacencies* if for all graphs $G \in \Pi$, for every vertex $v \in V(G)$, there is a set $D \subseteq V(G) \setminus \{v\}$ of size at most c_Π such that all graphs G' that are obtained from G by adding or removing edges between v and vertices in $V(G) \setminus D$, are also contained in Π .

Fomin et al. show that graph problems such as Π -FREE DELETION [VC], can be solved efficiently through kernelization when Π is characterized by few adjacencies (and meets some other demands), by making heavy use of the REDUCE algorithm they provide. The idea behind the REDUCE algorithm is to save *enough* vertices with specific adjacencies in the vertex cover, and those vertices that we forget have equivalent vertices saved to replace them. The sets of adjacencies we have to consider can be reduced by making use of the characterization by few adjacencies, as more than c_Π adjacencies are not relevant. The number of vertices we retain is ultimately dependent on $\ell \leq K$. The kernel by Fomin et al. [22] is given as follows.

Algorithm 1 REDUCE(Graph G , Vertex Cover $X \subseteq V(G)$, $r \in \mathbb{N}$, $c \in \mathbb{N}$) [22, Algorithm 1].

for all $Y \in \binom{X}{\leq c}$ and a partition of Y into $Y^+ \cup Y^-$ **do**
 $Z := \{v \in V(G) \setminus X \mid v \text{ is adjacent to all of } Y^+ \text{ and none of } Y^-\}$
 Mark r arbitrary vertices of Z (if $|Z| < r$ then mark all of them)
 Delete from G all unmarked vertices that are not contained in X

Notice that if the input contains no degree-0 vertices, then neither does the kernel output by Algorithm 1, as any vertex in X has at least some marked neighbour. In the AL streaming model, we have enough information to compute this kernel, by careful memory management in counting adjacencies towards specific subsets of the vertex cover. We now go into detail on the streaming adaptation of this kernel, as given in Algorithm 2.

Theorem 1. For a fixed constant c_Π , REDUCESTR(G, X, r, c_Π), where G is provided as a stream in the AL model, is a streaming equivalent of REDUCE, using one pass, and $\mathcal{O}((|X| + |X|^{c_\Pi}) \log(n))$ bits of memory resulting in a kernel on $\mathcal{O}(|X| + r \cdot |X|^{c_\Pi})$ vertices, output as an EA stream.

Proof. Let us first elaborate on the working of Algorithm 2. The set Z stores all the partitions considered in the original REDUCE, together with two counters per partition. The first counter, x , tracks the total amount of vertices already ‘marked’ with this partition, which does not exceed r . The second counter is reset at every vertex, and tracks whether v is adjacent to the entirety of Y^+ . This means these two counters mimic exactly the marking behaviour that REDUCE applies, except that the marking is not on arbitrary vertices, but dependent on the order of the stream (this does not impact the correctness). The rest of the algorithm merely interacts with Z correctly and uses some storage, S and V' , to make sure the output is constructed correctly without using too much memory. V' remembers which vertices in X we have already seen, to avoid outputting the same edge twice. S saves the set of edges adjacent to a vertex v , and if we mark v , we output S . As the output of REDUCE contains no degree-0 vertices, the stream output by REDUCESTR contains all marked vertices and those in X .

Let it be clear by the above motivation that the output of REDUCESTR can also be an output of REDUCE, and therefore, the algorithm works correctly. Let us analyse the space usage. The main concern is the space usage of the set Z , which contains partitions of sets in $\binom{X}{\leq c}$. There are $\mathcal{O}(|X|^c)$ such sets, each with at most 2^c partitions, and each set has at most c elements (using $\log n$ space). This means Z uses $\mathcal{O}(c \cdot 2^c \cdot |X|^c \cdot \log n) = \mathcal{O}(|X|^c \log n)$ bits of space. The set S is reset at every vertex v (which is not part of the vertex cover), and contains at most all edges incident on v . As v is not an element of the vertex cover, the degree of v is at most $|X|$. So the set S uses $\mathcal{O}(|X| \log n)$ bits of space. The set V' has size at most $|X|$, and so uses $\mathcal{O}(|X| \log n)$ bits of space. All in all, our memory never exceeds $\mathcal{O}((|X| + |X|^c) \log n)$ bits. \square

An interesting observation is that Algorithm 2 takes as input a stream and outputs the kernel as an EA stream, which we will now call a *streaming kernel*. This should not have a lot of impact, as a subsequent algorithm may want to store the entire kernel anyway. However, if we want the algorithm itself to store the entire kernel, this might increase the memory use, as we would have to store all edges contained in the kernel. The same goes for if we would want to output the kernel as an AL stream (which is how our input is provided). Let us shortly analyse how much memory would be needed for this. For a vertex cover X in a graph G , saving $G[X]$ entirely can take up to $\mathcal{O}(|X|^2 \log n)$ bits. Next to the vertex cover, the output kernel consists of $\mathcal{O}(r \cdot |X|^c)$ vertices, each of degree at most $|X|$, as these vertices are not part of the vertex cover. Therefore, the total memory use with this approach can be $\mathcal{O}((|X|^2 + r \cdot |X|^{c+1}) \log n)$ bits. To clarify, this is not the approach we take, our algorithms output a stream and do not save the kernel to memory.

The following theorem then shows how this algorithm leads to streaming kernels for Π -FREE DELETION [VC] as an adaptation of [22, Theorem 2]. We call a graph G *vertex-minimal* with respect to Π if $G \in \Pi$ and for all $S \subsetneq V(G)$, $G[S] \notin \Pi$.

Theorem 2. If Π is a graph property such that:

- (i) Π is characterized by c_Π adjacencies,
- (ii) every graph in Π contains at least one edge, and

Algorithm 2 REDUCESTR(Graph $G = (V, E)$ given as a stream in the AL model, Vertex Cover $X \subseteq V(G)$, $r \in \mathbb{N}$, $c \in \mathbb{N}$).

```

1: for each  $Y \in \binom{X}{\leq c}$  and a partition of  $Y$  into  $Y^+ \cup Y^-$  do ▷ Calculate and store partitions
2:   Store  $(Y^+, Y^-, 0, 0)$  in  $Z$ 
3:   Store the output vertices  $V' \leftarrow \emptyset$  ▷ Required for neatly outputting  $X$ 
4:   for each  $v \in V$  in the stream do ▷ This entire loop requires only one pass
5:     if  $v \in X$  then
6:       for each  $(v, w) \in E$  in the stream do
7:         if  $w \in V'$  then Output  $(v, w)$  as part of the kernel
8:        $V' \leftarrow V' \cup \{v\}$ 
9:     else ▷  $v \notin X$ 
10:      for each  $(Y^+, Y^-, x, y) \in Z$  do ▷ Reset local counters
11:         $y \leftarrow 0$ 
12:      Store an edge set  $S \leftarrow \emptyset$  ▷ Reset local edge memory
13:      for each  $(v, w) \in E$  in the stream do
14:        for each  $(Y^+, Y^-, x, y) \in Z$  where  $x < r$  and  $y \geq 0$  do ▷ Count 'correct' partitions
15:          if  $w \in Y^+$  then  $y \leftarrow y + 1$ 
16:          if  $w \in Y^-$  then  $y \leftarrow y - 1$ 
17:           $S \leftarrow S \cup \{(v, w)\}$ . ▷ If we mark  $v$ , then  $(v, w)$  needs to be added
18:        if  $\exists (Y^+, Y^-, x, y) \in Z$  where  $y = |Y^+|$  then ▷ Mark  $v$  and output what we can
19:           $x \leftarrow x + 1$  ▷ Increment  $x$  so that this partition marks at most  $r$  vertices
20:        Output  $S$  as part of the kernel

```

(iii) there is a non-decreasing polynomial $p : \mathbb{N} \rightarrow \mathbb{N}$ such that all graphs G that are vertex-minimal with respect to Π satisfy $|V(G)| \leq p(K)$,

then Π -FREE DELETION [VC] admits a kernel on $\mathcal{O}((K + p(K))K^{c_\Pi})$ vertices output as an EA stream in the AL streaming model using one pass and $\mathcal{O}((K + K^{c_\Pi}) \log(n))$ space, where K is the vertex cover number of the input graph.

Proof. Combine Theorem 1 with the proof of [22, Theorem 2], where instead of calling REDUCE($G, X, \ell + p(|X|), c_\Pi$) we call REDUCESTR($G, X, \ell + p(|X|), c_\Pi$). \square

We note that the theorem applies to \mathcal{F} -SUBGRAPH DELETION [VC] when $\Delta(\mathcal{F})$ (the maximum degree) is bounded as well as to CLUSTER VERTEX DELETION [VC]. As such, our streaming kernels generalize the kernels of Bishnu et al. [13] for these problems, while the memory requirements and kernel sizes are fairly comparable. Next we have a discussion and further implications for several general problems, following Fomin et al. [22].

Let us go over a few applications of Theorem 2. Consider CLUSTER VERTEX DELETION [VC]. This problem is exactly Π -FREE DELETION [VC] where $\Pi = \{P_3\}$, as any P_3 -free graph can only contain clusters. Notice that $c_\Pi = 2$ and $p(K) = 3$ suffice to meet the demands of Theorem 2. This implies that, given that we already have a vertex cover for our graph, we have a one-pass streaming algorithm for CVD in the AL model using $\mathcal{O}(K^3 \log n)$ space (or $\mathcal{O}(K^3 \log n)$ bits of space to save the kernel). Considering that this algorithm is a simple adaptation of known results, it is interesting to observe that it compares quite well to e.g. a more recent result of Bishnu et al. [13], who show that CVD admits a one-pass (randomized) streaming algorithm in the AL model using $\mathcal{O}(K^2 \log^4(n))$ space when parameterized by the size K of a vertex cover. Note that the kernel sizes differ, as the kernel given by Bishnu et al. [13] uses $\mathcal{O}(K^2 \log^2 n)$ bits of space, while the kernel from Algorithm 2 uses $\mathcal{O}(K^3 \log n)$ bits of space.

Next consider TRIANGLE DELETION [VC]. This problem is exactly Π -FREE DELETION [VC] where $\Pi = \{C_3\}$. Once again, it is easy to observe that $c_\Pi = 2$, and $p(K) = 3$ suffice to meet the demands of Theorem 2. This means that, given that we already have a vertex cover, we have a one-pass streaming algorithm for TD in the AL model using $\mathcal{O}(K^3 \log n)$ bits of space (if we save the entire kernel). Once again, this result compares very well to a result of Bishnu et al. [13], who show that TD admits a one-pass streaming algorithm in the AL model using $\tilde{\mathcal{O}}(K^3)$ bits of space when parameterized by the size K of a vertex cover.

Most interesting is that Theorem 2 also applies to \mathcal{F} -SUBGRAPH DELETION [VC] when $d = \Delta(\mathcal{F}) \leq K$ is also bounded. Bishnu et al. [13] showed that this problem admits a one-pass, $\mathcal{O}(d \cdot K^{d+1} \cdot \log n)$ bits memory, algorithm in the AL model. To apply Theorem 2 to \mathcal{F} -SUBGRAPH DELETION [VC], construct Π from \mathcal{F} by adding for each $F \in \mathcal{F}$ to Π both F and all graphs F' obtained by adding some set of edges to F (i.e. we end up with all graphs on the same vertex set that contain F as a subgraph). Notice that this Π is characterized by $c_\Pi = \Delta(\mathcal{F})$ adjacencies, and because $\Delta(\mathcal{F}) \leq K$, the size of every graph is also bounded by $p(K) = \mathcal{O}(K^2)$. Now Theorem 2 gives us a kernel of $\mathcal{O}(K^{\Delta(\mathcal{F})+2})$ vertices obtained in one pass and $\mathcal{O}(K^{\Delta(\mathcal{F})} \log n)$ space in the AL model. This implies a one-pass algorithm using $\mathcal{O}(K^{\Delta(\mathcal{F})+2} \log n)$ space for the \mathcal{F} -SUBGRAPH DELETION [VC] problem when $\Delta(\mathcal{F}) \leq K$ in the AL model. Notice that this result is comparable to that of Bishnu et al. [13], while our kernel is more generally applicable.

The second general result is for finding largest induced subgraphs.

LARGEST INDUCED Π -SUBGRAPH [VC]*Input:* A graph G with a vertex cover X , and an integer $\ell \geq 1$.*Parameter:* The size $K := |X|$ of the vertex cover.*Question:* Is there a set $P \subseteq V(G)$ of size at least ℓ such that $G[P] \in \Pi$?

With this definition, we give the following theorem as an adaptation of [22, Theorem 3].

Theorem 3. *If Π is a graph property such that:*

- (i) Π is characterized by c_Π adjacencies, and
- (ii) there is a non-decreasing polynomial $p : \mathbb{N} \rightarrow \mathbb{N}$ such that all graphs $G \in \Pi$ satisfy $|V(G)| \leq p(K)$,

then **LARGEST INDUCED Π -SUBGRAPH [VC]** admits a kernel on $\mathcal{O}(p(K) \cdot K^{c_\Pi})$ vertices output as an EA stream in the AL streaming model using one pass and $\mathcal{O}((K + K^{c_\Pi}) \log(n))$ bits of space, where K is the vertex cover number of the input graph.

Proof. Combine Theorem 1 with the proof of [22, Theorem 3], where instead of calling $\text{REDUCE}(G, X, p(|X|), c_\Pi)$ we call $\text{REDUCESTR}(G, X, p(|X|), c_\Pi)$. \square

Examples of problems fitting in the **LARGEST INDUCED Π -SUBGRAPH [VC]** category, that are also characterized by few adjacencies, are **LONG CYCLE**, **LONG PATH**, and **H-PACKING**.

The third general result is for graph partitioning problems.

PARTITION INTO q DISJOINT Π -FREE SUBGRAPHS [VC]*Input:* A graph G with a vertex cover X .*Parameter:* The size $K := |X|$ of the vertex cover.*Question:* Is there a partition of the vertex set into q sets $S_1 \cup S_2 \cup \dots \cup S_q$ such that for each $i \in [q]$ the graph $G[S_i]$ does not contain a graph in Π as an induced subgraph?

Note that q is regarded a constant.

Theorem 4. *If Π is a graph property such that:*

- (i) Π is characterized by c_Π adjacencies, and
- (ii) there is a non-decreasing polynomial $p : \mathbb{N} \rightarrow \mathbb{N}$ such that all graphs G that are vertex-minimal with respect to Π satisfy $|V(G)| \leq p(K)$,

then **PARTITION INTO q DISJOINT Π -FREE SUBGRAPHS [VC]** admits a kernel on $\mathcal{O}(p(K) \cdot K^{q \cdot c_\Pi})$ vertices output as an EA stream in the AL streaming model using one pass and $\mathcal{O}((K + K^{q \cdot c_\Pi}) \log(n))$ bits of space, where K is the vertex cover number of the input graph.

Proof. Combine Theorem 1 with the proof of [22, Theorem 4] and [22, Lemma 2], where instead of calling $\text{REDUCE}(G, X, q \cdot p(|X|), q \cdot c_\Pi)$ we call $\text{REDUCESTR}(G, X, q \cdot p(|X|), q \cdot c_\Pi)$. \square

Examples of problems fitting in the **PARTITION INTO q DISJOINT Π -FREE SUBGRAPHS [VC]** category, that are also characterized by few adjacencies, are **PARTITION INTO q INDEPENDENT SETS**, **PARTITION INTO q CLIQUES**, **PARTITION INTO q PLANAR GRAPHS**, and **PARTITION INTO q FORESTS**.

2.1. Kernel for characterization by low-rank adjacencies

We also give an adaptation of a more recent kernel by Jansen and Kroon [23], which has another broad range of implications for streaming kernels. This kernel uses a different characterization of the graph family, however, the adaptation to the AL streaming model is very similar. The adaptation of this kernel leads to streaming algorithms for problems like **PERFECT DELETION [VC]**, **AT-FREE DELETION [VC]**, **INTERVAL DELETION [VC]**, and **WHEEL-FREE DELETION [VC]**.

The characterization for Π in this case is a characterization by low-rank adjacencies. For this, however, we need the definition of a c -rank incidence vector.

Definition 2. ([23, Definition 3.1]⁵) Let G be a graph with vertex cover X and let $c \in \mathbb{N}$. Let $Q', R' \subseteq X$ such that $|Q'| + |R'| \leq c$ and $Q' \cap R' = \emptyset$. We define the c -incidence vector $\text{inc}_{(G,X)}^{c,(Q',R')}(u)$ for a vertex $u \in V(G) \setminus X$ as a vector over \mathbb{F}_2 that has an entry for each (Q, R) , where $Q, R \subseteq X$ with $Q \cap R = \emptyset$, such that $|Q| + |R| \leq c$, $Q' \subseteq Q$ and $R' \subseteq R$. It is defined as follows:

$$\text{inc}_{(G,X)}^{c,(Q',R')}(u)[Q, R] = \begin{cases} 1 & \text{if } N_G(u) \cap Q = \emptyset \text{ and } R \subseteq N_G(u), \\ 0 & \text{otherwise.} \end{cases}$$

The superscript (Q', R') is dropped when $Q' = R' = \emptyset$. Note that the order of the coordinates of the vector is fixed, but not explicit, as any order suffices. Therefore, we can also sum such incidence vectors coordinate-wise.

With this, we can give the definition of a graph property being characterized by rank- c adjacencies.

Definition 3. ([23, Definition 3.3]) Let $c \in \mathbb{N}$ be a natural number. A graph property Π is characterized by rank- c adjacencies if the following holds. For each graph H , for each vertex cover X of H , for each set $D \subseteq V(H) \setminus X$, for each $v \in V(H) \setminus (D \cup X)$, if

- $H - D \in \Pi$, and
- $\text{inc}_{(H,X)}^c(v) = \sum_{u \in D} \text{inc}_{(H,X)}^c(u)$ when evaluated over \mathbb{F}_2 ,

then there exists $D' \subseteq D$ such that $H - v - (D \setminus D') \in \Pi$. If there always exists such set D' of size 1, then we say Π is characterized by rank- c adjacencies with singleton replacements.

Jansen and Kroon note that the intuition here is that if we have a set D such that $H - D \in \Pi$, and the c -incidence vectors of D sum to the vector of some vertex v over \mathbb{F}_2 , then there exists $D' \subseteq D$ such that removing v from $H - D$ and adding back D' results in a graph that is still contained in Π . We can notice how there is some similarity in essential and non-essential adjacencies for occurrences of graphs in Π when comparing this characterization to that of few adjacencies (see Definition 1).

We will give the kernelization algorithm as given by Jansen and Kroon [23] next, but first, we recall a linear algebraic definition, that of a *basis*. As Jansen and Kroon state [23], “a *basis* of a set S of d -dimensional vectors over a field \mathbb{F} is a minimum-size subset $B \subseteq S$ such that all $\vec{v} \in S$ can be expressed as linear combinations of elements of B , i.e., $\vec{v} = \sum_{\vec{u} \in B} \alpha_{\vec{u}} \cdot \vec{u}$ for a suitable choice of coefficients $\alpha_{\vec{u}} \in \mathbb{F}$. When working over the field \mathbb{F}_2 , the only possible coefficients are 0 and 1, which gives a basis B of S the stronger property that any vector $\vec{v} \in S$ can be written as $\sum_{\vec{u} \in B'} \vec{u}$, where $B' \subseteq B$ consists of those vectors which get a coefficient of 1 in the linear combination”.

The essence of the kernel comes down to computing the basis of a set of incidence vectors of the remaining graph and adding vertices corresponding to the basis to the kernel. We give this kernel here as Algorithm 3.

Algorithm 3 LOW-RANK REDUCE(Graph G , Vertex Cover $X \subseteq V(G)$, $\ell \in \mathbb{N}$, $c \in \mathbb{N}$) [23, Algorithm 1].

```

1: Let  $Y_1 := V(G) \setminus X$ 
2: for  $i \leftarrow 1$  to  $\ell$  do
3:   Let  $V_i = \{\text{inc}_{(G,X)}^c(y) \mid y \in Y_i\}$  and compute a basis  $B_i$  of  $V_i$  over  $\mathbb{F}_2$ .
4:   For each  $\vec{v} \in B_i$ , choose a unique vertex  $y_{\vec{v}} \in Y_i$  such that  $\vec{v} = \text{inc}_{(G,X)}^c(y_{\vec{v}})$ .
5:   Let  $A_i := \{y_{\vec{v}} \mid \vec{v} \in B_i\}$  and  $Y_{i+1} = Y_i \setminus A_i$ .
6: return  $G[X \cup \bigcup_{i=1}^{\ell} A_i]$ 

```

Jansen and Kroon show that Algorithm 3 runs in polynomial time in ℓ and the size of G (for a constant c), and returns a graph on $\mathcal{O}(|X| + \ell \cdot |X|^c)$ vertices. Also here, the output contains no degree-0 vertices, if the input does not contain any. We will now adapt Algorithm 3 to the streaming setting, and conclude it to be a streaming equivalent of [23, Theorem 3.5].

2.1.1. Adapting low-rank reduce

If we want to adapt Algorithm 3, LOW-RANK REDUCE, to the streaming setting, we are faced with a few challenges. For one, the set V_i consists of $\mathcal{O}(n)$ vectors, so saving this entire set is not desirable. We also have to consider how we can compute the basis of the set V_i if we do not want to save it. Luckily, the incidence vectors are computable from local information combined with the vertex cover, and computing a basis can be done incrementally by checking linear (in)dependence. With this small motivation, we give the adaptation of LOW-RANK REDUCE into the streaming setting, LOW-RANK REDUCESTR, in Algorithm 4.

Theorem 5. Algorithm 4 is a streaming equivalent of Algorithm 3, that is, given a graph G as an AL stream with a vertex cover X , and integer ℓ and a constant c , Algorithm 4 returns a graph on $\mathcal{O}(|X| + \ell \cdot |X|^c)$ vertices as an AL stream that could be the output of Algorithm 3 given the same input. Algorithm 4 uses $\ell + 1$ passes and $\mathcal{O}((|X| + \ell \cdot |X|^c) \log n + |X|^{2c})$ bits of memory.

⁵ We define each (Q, R) slightly different to [23], where they use $(Q, R) \in X \times X$. The notation $Q, R \subseteq X$ correctly reflects what the authors meant.

Algorithm 4 LOW-RANK REDUCESTR(Graph G as an AL stream, Vertex Cover $X \subseteq V(G)$, $\ell \in \mathbb{N}$, $c \in \mathbb{N}$).

```

1: Let  $A \leftarrow \emptyset$ 
2: for  $i \leftarrow 1$  to  $\ell$  do
3:    $A_i \leftarrow \emptyset$ 
4:    $B \leftarrow \emptyset$ 
5:   for each Vertex  $v \in V \setminus (X \cup A)$  in the stream do ▷ Entire loop in one pass
6:     Save the  $\leq |X|$  adjacencies of  $v$  in  $X$  (until the next vertex)
7:     Let  $\vec{v} \leftarrow \text{inc}_{(G,X)}^c(v)$  ▷ Can be computed from  $X$  and the adjacencies of  $v$ 
8:     If  $\vec{v}$  is linearly independent w.r.t.  $B$  over  $\mathbb{F}_2$ , do  $B \leftarrow B \cup \{\vec{v}\}$  and  $A_i \leftarrow A_i \cup \{v\}$ .
9:   Let  $A \leftarrow A \cup A_i$ .
10: return  $G[X \cup A]$ 

```

Proof. Let us first show the equivalence of Algorithm 4 to Algorithm 3. Let G be a graph with vertex cover X , and let ℓ and c be integers. Algorithm 4 and Algorithm 3 both do ℓ iterations over some process over all vertices, excluding the vertices in X and a set A (for both algorithms, $A = \bigcup_{i=1}^{\ell} A_i$). With these vertices, Algorithm 3 computes the incidence vector of each of them, and then computes a basis for this set. The new vertices added to A are then vertices with incidence vectors equivalent to those in the basis. Algorithm 4 computes the incidence vectors of these vertices one vertex at a time. It then checks whether the current incidence vector is linearly independent of a set B , and if so, the incidence vector is added to B . We can see that B must consist of a basis of all incidence vectors seen so far. Therefore, after Algorithm 4 has seen every vertex (excluding those in X and A), B is a basis that could also be found by Algorithm 3 for the same set of vertices. We conclude that in every iteration, Algorithm 4 adds to A a set of vertices which Algorithm 3 can also add to A in the corresponding iteration. As both algorithms return $G[X \cup A]$, the output of Algorithm 4 can also be an output of Algorithm 3.

As the incidence vector of a vertex v can be computed from the adjacencies of v together with X alone, and linear (in)dependence checking only requires the vectors to be in memory, we only require one pass for each $1 \leq i \leq \ell$. Another pass is used to compute the output, as we can produce an AL stream corresponding to $G[X \cup A]$ by simply using a pass and output only those edges between vertices in $X \cup A$. Therefore, Algorithm 4 uses $\ell + 1$ passes over the stream.

In terms of memory, Algorithm 4 stores A , B , X , adjacencies of a vertex v , and an incidence vector of v , \vec{v} . Clearly, the memory used by the adjacencies of v and \vec{v} is dominated by saving X and B in memory. The computation of \vec{v} also requires no more memory than X uses, as it iterates subsets of X to compute entries of the vector. Saving X requires $\mathcal{O}(|X|)$ bits of memory. B consists of at most $\mathcal{O}(|X|^c)$ vectors, and each vector consists of $\mathcal{O}(|X|^c)$ bits, as motivated by Jansen and Kroon [23, Proposition 3.4]. It follows that B requires $\mathcal{O}(|X|^{2c})$ bits of space. The set A consists of at most $\mathcal{O}(\ell \cdot |X|^c)$ vertices, as in every iteration B contains at most $\mathcal{O}(|X|^c)$ vectors. We conclude that the total memory usage is $\mathcal{O}((|X| + \ell \cdot |X|^c) \log n + |X|^{2c})$ bits. \square

With Theorem 5 we are ready to give the streaming equivalent of [23, Theorem 3.5].

Theorem 6. *If Π is a graph property such that:*

- (i) Π is characterized by rank- c adjacencies,
- (ii) every graph in Π contains at least one edge, and
- (iii) there is a non-decreasing polynomial $p : \mathbb{N} \rightarrow \mathbb{N}$ such that all graphs G that are vertex-minimal with respect to Π satisfy $|V(G)| \leq p(K)$,

then Π -FREE DELETION [VC] in the AL streaming model admits a kernel on $\mathcal{O}((K + p(K)) \cdot K^c)$ vertices output as an AL stream using $K + p(K) + 2$ passes and $\mathcal{O}((K + p(K)) \cdot K^c \log n + K^{2c})$ bits of memory, where K is the vertex cover number of the input graph.

Proof. See the proof of [23, Theorem 3.5], where instead of LOW-RANK REDUCE($G, X, \ell := k + 1 + p(|X|), c$) we call LOW-RANK REDUCESTR($G, X, \ell := k + 1 + p(|X|), c$). By Theorem 5 the theorem follows. \square

Let us shortly list some implications of Theorem 6, which consist of some problems admitting streaming kernels, following Jansen and Kroon [23].

The first result is for PERFECT DELETION [VC]. PERFECT DELETION [VC] is Π -FREE DELETION [VC] where Π is the set of all graphs that contain an odd hole or an odd anti-hole. An *odd hole* is a cycle consisting of an odd number of vertices, and an *odd anti-hole* is the complement graph of an odd hole.

Theorem 7. PERFECT DELETION [VC] in the AL streaming model admits a kernel on $\mathcal{O}(K^5)$ vertices output as an AL stream using $\mathcal{O}(K)$ passes and $\mathcal{O}(K^5 \log n + K^8)$ bits of memory, where K is the vertex cover number of the input graph.

Proof. See [23, Theorem 4.4], but instead of applying [23, Theorem 3.5] we apply Theorem 6. \square

The second result is for AT-FREE DELETION [VC]. This is Π -FREE DELETION [VC] where Π is the set of all graphs that contain an asteroidal triple. An *asteroidal triple* is a set of three vertices where every two vertices in the triple are connected by a path that avoids the neighbourhood of the third.

Theorem 8. AT-FREE DELETION [VC] in the AL streaming model admits a kernel on $\mathcal{O}(K^9)$ vertices output as an AL stream using $\mathcal{O}(K)$ passes and $\mathcal{O}(K^9 \log n + K^{16})$ bits of memory, where K is the vertex cover number of the input graph.

Proof. See [23, Theorem 4.12], but instead of applying [23, Theorem 3.5] we apply Theorem 6. \square

The third result is for INTERVAL DELETION [VC]. This is Π -FREE DELETION [VC] where Π is the set of all graphs that contain either an asteroidal triple or an induced cycle of length at least 4, or both.

Theorem 9. INTERVAL DELETION [VC] in the AL streaming model admits a kernel on $\mathcal{O}(K^9)$ vertices output as an AL stream using $\mathcal{O}(K)$ passes and $\mathcal{O}(K^9 \log n + K^{16})$ bits of memory, where K is the vertex cover number of the input graph.

Proof. See [23, Theorem 4.13], but instead of applying [23, Theorem 3.5] we apply Theorem 6. \square

The fourth result is for WHEEL-FREE DELETION [VC]. This is Π -FREE DELETION [VC] where Π is the set of all graphs that contain a wheel of size at least 3. A *wheel* of size $n \geq 3$ is a set of $n + 1$ vertices, where n vertices form a cycle, and one vertex is connected to all vertices on the cycle (the center of the wheel).

Theorem 10. WHEEL-FREE DELETION [VC] in the AL streaming model admits a kernel on $\mathcal{O}(K^5)$ vertices output as an AL stream using $\mathcal{O}(K)$ passes and $\mathcal{O}(K^5 \log n + K^8)$ bits of memory, where K is the vertex cover number of the input graph.

Proof. See [23, Theorem 4.25], but instead of applying [23, Theorem 3.5] we apply Theorem 6. \square

It is interesting to note that the above problems are not characterized by few adjacencies, but are characterized by rank- c adjacencies, therefore requiring Algorithm 4 to admit a streaming kernel.

3. A direct FPT approach

In this section, we give direct FPT streaming algorithms for Π -FREE DELETION [VC] for the same cases as Theorem 2. This is motivated by the fact that Chitnis and Cormode [12] found a direct FPT algorithm for VERTEX COVER using $\mathcal{O}(2^k)$ passes and only $\mathcal{O}(k)$ space in contrast to the kernel of Chitnis et al. [9] using one pass and $\mathcal{O}(k^2)$ space. Therefore, we aim to explore the pass/memory trade-off for Π -FREE DELETION [VC] as well.

For enumeration of sets in our algorithms in this section, we use a black box dictionary technique also used by Chitnis and Cormode [39,12]. One can implement such a dictionary using k pointers (for sets of size $\leq k$) to the universe, and increment the last pointer (or reset it and increment second to last, etc.) to implement the NEXT operation. This way the dictionaries are memory efficient.

Definition 4. ([39, Definition 9]) Let $U = \{u_1, u_2, \dots, u_n\}$ and $k \leq n$. Let $\mathcal{U}_{\leq k}$ denote the set of all $\sum_{i=0}^k \binom{|U|}{i}$ subsets of U which have at most k elements, and let $\text{DICT}_{\mathcal{U}_{\leq k}}$ be the dictionary ordering on $\mathcal{U}_{\leq k}$. Given a subset $X \in \mathcal{U}_{\leq k}$, let $\text{DICT}_{\mathcal{U}_{\leq k}}(\text{NEXT}(X))$ denote the subset that comes immediately after X in the ordering $\text{DICT}_{\mathcal{U}_{\leq k}}$. We denote the last subset in the dictionary order of $\mathcal{U}_{\leq k}$ by $\text{LAST}(\mathcal{U}_{\leq k})$, and similarly the first subset as $\text{FIRST}(\mathcal{U}_{\leq k})$, and use the notation that $\text{DICT}_{\mathcal{U}_{\leq k}}(\text{NEXT}(\text{LAST}(\mathcal{U}_{\leq k}))) = \spadesuit$. Similarly, we define \mathcal{U}_k as the set of all $\binom{|U|}{k}$ subsets of U with exactly k elements, and define \mathcal{U} as the set of all subsets of U , and analogously define the dictionary ordering on these sets.

3.1. P_3 -free deletion

We start with the scenario where $\Pi = \{P_3\}$, which means we consider the problem CLUSTER VERTEX DELETION [VC]. The general idea of the algorithm is to branch on what part of the given vertex cover should be in the solution. In a branch, we first check whether the ‘deletion-free’ part of the vertex cover (Y) contains a P_3 , which invalidates a branch. Otherwise, what remains is some case analysis where either one or two vertices of a P_3 lie outside the vertex cover, for which we deterministically know which vertices have to be removed to make the graph P_3 -free. We illustrate this step in Fig. 1. Case 1 and 2 have only one option for removal of a vertex. After Case 1 and 2 no longer occur, we can find Case 3 occurrences and show that we can delete all but one neighbour of y in such an occurrence. So, if this process can be executed in a limited number of passes, the algorithm works correctly.

We give the full algorithm in Algorithm 5.




Fig. 1. The different cases how a P_3 can exist with respect to Y , part of the vertex cover. Notice that the case where the entire P_3 is contained in Y is not included here. Case 3 assumes there are no Case 1 or Case 2 P_3 's in the graph anymore.

Algorithm 5 P_3 -FREE DELETION(Graph $G = (V, E)$ given as a stream in the AL model, integer ℓ , Vertex Cover $X \subseteq V(G)$).

```

1:  $S \leftarrow \text{FIRST}(\mathcal{X}_{\leq \ell})$ 
2: while  $S \in \mathcal{X}_{\leq \ell}, S \neq \blacklozenge$  do
3:    $Y \leftarrow X \setminus S$ 
4:    $S' \leftarrow S$ 
5:    $P \leftarrow \text{FIRST}(\mathcal{Y}_2)$ 
6:   while  $P = (y_1, y_2) \in \mathcal{Y}_2, P \neq \blacklozenge$  do
7:     for each Vertex  $v \in Y \setminus P$  do
8:       if  $v$  and  $P$  form a  $P_3$ ,  $Y$  is invalid, move to the next  $S$ 
9:        $P \leftarrow \text{DICT}_{\mathcal{Y}_2}(\text{NEXT}(P))$ 
10:   $P \leftarrow \text{FIRST}(\mathcal{Y}_2)$ 
11:  while  $P = (y_1, y_2) \in \mathcal{Y}_2, P \neq \blacklozenge$  do
12:    if  $y_1 y_2$  is an edge then
13:      for each Vertex  $v \in V \setminus (X \cup S')$  in the stream do
14:        if Either  $v y_1$  or  $v y_2$  is present and the other is not then  $S' \leftarrow S' \cup \{v\}$ 
15:      else
16:        for each Vertex  $v \in V \setminus (X \cup S')$  in the stream do
17:          if Both  $v y_1$  and  $v y_2$  are present then  $S' \leftarrow S' \cup \{v\}$ 
18:         $P \leftarrow \text{DICT}_{\mathcal{Y}_2}(\text{NEXT}(P))$ 
19:    for each  $y \in Y$  do
20:       $b \leftarrow \text{False}$ 
21:      for each Vertex  $v \in V \setminus (X \cup S')$  in the stream do
22:        if The edge  $v y$  is present and  $b = \text{False}$  then  $b \leftarrow \text{True}$ 
23:        else if The edge  $v y$  is present then  $S' \leftarrow S' \cup \{v\}$ 
24:    if  $|S'| \leq \ell$  then return  $S'$ 
25:   $S \leftarrow \text{DICT}_{\mathcal{X}_{\leq \ell}}(\text{NEXT}(S))$ 
26: return NO

```

To limit the number of passes, the use of the AL model is crucial. Notice that for every pair of vertices y_1, y_2 in the vertex cover, we can identify a Case 1 or 2 P_3 of Fig. 1, or these cases with v in the vertex cover as well, in a constant number of passes. This is because we can first use a pass to check the presence of an edge between y_1 and y_2 , and afterwards use a pass to check the edges of every other vertex towards y_1 and y_2 (which are given together because of the AL model). This means we can find P_3 's contained in the vertex cover or corresponding to Case 1 or 2 P_3 's in $\mathcal{O}(K^2)$ passes total. The remaining Case 3 can be handled in $\mathcal{O}(K)$ passes from the viewpoint of each $y \in Y$. So this algorithm takes $\mathcal{O}(2^K K^2)$ passes (including branching).

Theorem 11. We can solve CLUSTER VERTEX DELETION [VC] in the AL streaming model using $\mathcal{O}(2^K K^2)$ passes and $\mathcal{O}(K \log n)$ space, where K is the vertex cover number of the input graph.

Proof. We claim that Algorithm 5 does exactly this.

Let us first reason that the number of passes and memory use are as stated. Let X be the provided vertex cover, and let $|X| = K$. The number of different sets S can take is bounded by 2^K . The first and second loop enumerate all pairs of vertices in the vertex cover, and use a two passes per pair to detect a P_3 , which gives us $\mathcal{O}(K^2)$ passes. The last loop enumerates all K vertices in the vertex cover and uses one pass per iteration. Therefore, the total number of passes is bounded by $\mathcal{O}(2^K K^2)$.

No set in memory exceeds $\mathcal{O}(K \log n)$ bits, so the stated memory complexity is correct.

Let us now show the correctness of the algorithm. The main idea of the algorithm is to branch on what part of the vertex cover is contained in the solution S' . This is modelled through the use of the sets Y and S , where in each branch, we cannot add vertices in Y to S' . Therefore, we first check whether Y fully contains a P_3 , and if one is found we stop, as we may not delete any vertex of this P_3 in this branch. For a fixed pair (v, w) in the vertex cover, checking for a P_3 that contains v and w only takes two passes because the only necessary information is whether v and w are neighbours (one pass), and the adjacencies of v and w towards another vertex, which is provided in the stream local to that vertex (see also Case 1 and Case 2 in the following analysis).

What remains is a careful analysis of the different cases of the structure of P_3 's with respect to Y . An illustration is given in Fig. 1. The loop of line 11 considers all pairs of vertices in Y . There are two cases we are interested in: Case 1 and Case 2 in Fig. 1. If we look at a single pair of vertices y_1 and y_2 either there is an edge between them (Case 1) or a non-edge (Case 2). These two vertices can then form a P_3 with any vertex outside Y in a very specific manner, which the algorithm looks for. It is then trivial that the one vertex outside Y has to be removed to make the graph P_3 free if a P_3 is found.

If there are no Case 1 or Case 2 P_3 's in the graph any more, we move on to Case 3. Note that this is the only remaining way a P_3 can be in the graph at all, because $Y \subseteq X$ is (part of) a vertex cover. In Case 3 at first it seems undecided which of the two vertices outside Y to remove, as one might lead to a solution and the other not. Let y_1, v, w form a Case 3 P_3 , where $y_1 \in Y$. Let us consider the scenario where v has another adjacency $y_2 \in Y$. Because there are no Case 2 P_3 's, y_1 and y_2 must be adjacent. Because there are no Case 1 P_3 's, w must now also be adjacent to y_2 . This means the structure extends as illustrated on the right in Case 3 in Fig. 1. We can observe that we need to delete all but one of the vertices attached to y , which is what the algorithm does. It does not matter which vertex we do not delete, as this vertex forms triangles if it has multiple adjacencies. Therefore, after these cases have all been handled, no induced P_3 's remain in the graph. If during the process S' never exceeded size ℓ , this means we have found a solution; otherwise, we move on to the next branch.

By the above reasoning, if there exists a solution of size at most ℓ for the CLUSTER VERTEX DELETION [VC] problem, then this solution contains some subset of the vertex cover X , which corresponds to some branch in the algorithm. As the removal of vertices is deterministic in each branch (as in, the solution must contain these vertices), and there exists a solution, the algorithm must find a solution too in that branch. If there exists no solution of size at most ℓ , then there exists no subset of vertices S' such that $G \setminus S'$ is induced P_3 free, and so in each branch of the algorithm S' will exceed size ℓ at some point, which results in the algorithm returning NO. \square

Let us stress some details. The use of the AL model is crucial, as it allows us to locally inspect the neighbourhood of a vertex when it appears in the stream. The same approach would require more memory or more passes in other models to accomplish this result. As a side-note, we could implement this algorithm in a normal setting (the graph is in memory, and not a stream) to get an algorithm for CLUSTER VERTEX DELETION [VC] with a running time of $\mathcal{O}(2^K \cdot K^2 \cdot (n + m))$ (although faster algorithms exist, see e.g. [40]).

3.2. H -free deletion

We now consider a more generalized form of Π -FREE DELETION [VC], where $\Pi = \{H\}$, a single graph. Unfortunately, the approach when $H = P_3$ does not seem to carry over to this case, because the structure of a P_3 is simple and local.

Theorem 12. *We can solve H -FREE DELETION [VC] in $2^{O(K^2)} \text{poly}(n, |V(H)|)$ time, where H contains at least one edge and K is the vertex cover number of the input graph.*

Proof. Let X be a vertex cover of $G = (V, E)$ of size K . Then $G[V \setminus X]$ has no edges and thus does not contain an occurrence of H . It follows that there is a solution of size at most K . Now call two vertices of $V \setminus X$ equivalent if their neighbourhood in X is the same. This yields 2^K equivalence classes. Observe that vertices in an equivalence class are interchangeable with respect to a solution for H -FREE DELETION [VC]: one can be exchanged for another without changing the validity of the solution. Hence, we may select the vertices of a solution from the first K vertices of an equivalence class. This means that there is a set of at most $(2^K + 1)K$ vertices in G that form a superset of some solution. Then we can enumerate all possible such solutions, which have size at most K , in $2^{O(K^2)}$ time. The validity of a solution can be checked in $2^{O(K \log K)} \text{poly}(n, |V(H)|)$ time through the algorithm of Abu-Khzam [41]. \square

In order to analyze the complexity with respect to H more precisely and to obtain a streaming algorithm, we present a different algorithm that works off a simple idea. We branch on the vertex cover, and then try to find occurrences of H of which we have to remove a vertex outside the vertex cover. We branch on these removals as well, and repeat this find-and-branch procedure. In an attempt to keep the second branching complexity low, we start by searching for occurrences of H such that only one vertex lies outside the vertex cover, and increase this number as we find no occurrences. For clarity, we present the occurrence detection part of the algorithm first, a procedure we call FINDH. Note that this is not (yet) a streaming algorithm.

Lemma 13. *Given a graph G with vertex cover X , graph H with at least one edge, sets S and $Y \subseteq X$ where $S \cap Y = \emptyset$, and integer i , Algorithm 6 finds an occurrence of H in G that contains no vertices in S and $X \setminus Y$ and contains $|V(H)| - i$ vertices in Y , if it exists. It runs in $\mathcal{O}\left(\binom{h}{i} [i^2 + \binom{K}{h-i} (h-i)! ((h-i)^2 + Kn + (h-i)in)]\right)$ time, where $|V(H)| = h$ and $|X| = K$.*

Proof. The correctness of the algorithm follows from the enumeration of all possibilities.

Algorithm 6 The procedure FINDH.

```

1: function FINDH(solution set  $S$ , forbidden set  $Y \subseteq X$ , integer  $i$ )
2:   for each Set  $O$  of  $i$  vertices of  $H$  that can be outside  $X$  do ▷ Check non-edges
3:     Denote  $H' = H \setminus O$ 
4:     for each Occurrence  $H_Y$  of  $H'$  in  $Y$  do ▷ Check  $\mathcal{O}(\binom{|X|}{h-i}(h-i)!)^2$  options
5:        $S' \leftarrow \emptyset$ ,  $O' \leftarrow O$ 
6:       for each Vertex  $v \in V \setminus (S \cup X)$  do
7:         Check the edges/non-edges towards  $H_Y$ 
8:         if  $v$  is equivalent to some  $w \in O'$  for  $H'$  then
9:            $S' \leftarrow S' \cup \{v\}$ ,  $O' \leftarrow O' \setminus \{w\}$ 
10:        if  $O' = \emptyset$  then return  $S' \cup H_Y$  ▷ We found an occurrence of  $H$ 
11:   return  $\emptyset$  ▷ No occurrence of  $H$  found

```

Let us analyse the running time. Checking all possible sets O takes $\mathcal{O}(\binom{h}{i}i^2)$ time resulting in at most $\mathcal{O}(\binom{h}{i})$ options for O . There are at most $\mathcal{O}(\binom{K}{h-i}(h-i)!)^2$ options for H' in Y : checking all of them costs $\mathcal{O}(\binom{K}{h-i}(h-i)!(h-i)^2)$ time. Then we take $\mathcal{O}((K + (h-i)in)$ time to, for each vertex, save adjacencies to H' , and check whether the adjacencies match a vertex in O . Here the factor K is from enumerating the adjacency list of a vertex not in X , and the $(h-i)$ factor is for checking adjacencies towards H' . Therefore, the running time of FINDH is $\mathcal{O}\left(\binom{h}{i}[i^2 + \binom{K}{h-i}(h-i)!((h-i)^2 + Kn + (h-i)in)]\right)$. \square

Now let us give the complete FPT algorithm for H -FREE DELETION [VC] (not in the streaming setting), see Algorithm 7.

Algorithm 7 H -FREE DELETION FPT(Graph $G = (V, E)$, integer ℓ , Vertex Cover $X \subseteq V(G)$).

```

1: for each Partition of  $X$  into  $S, Y$  where  $|S| \leq \ell$  do
2:   if  $H$  is not contained in  $Y$  then ▷ Check all  $\mathcal{O}(\binom{|X|}{h}h!)$  options
3:     if BRANCH( $S, Y, 1$ ) then
4:       return YES ▷ If any returns YES, we also return YES
5:   return NO

6: function BRANCH(solution set  $S$ , forbidden set  $Y \subseteq X$ , integer  $i$ )
7:    $B \leftarrow \text{FINDH}(S, Y, i)$  ▷ Try to find an  $H$  with  $i$  vertices outside  $Y$ 
8:   if  $B = \emptyset$  and  $i = |V(H)|$  then return YES
9:   else if  $B = \emptyset$  then BRANCH( $S, Y, i + 1$ ) ▷ No  $H$  found
10:  else if  $|S| = \ell$  then return NO ▷ Found an  $H$  but cannot remove it
11:  else
12:    for each  $v \in B \setminus Y$  do
13:      if BRANCH( $S \cup \{v\}, Y, i$ ) then return YES

```

Together with Lemma 13 we can analyse the performance of Algorithm 7 in its entirety.

Theorem 14. Algorithm 7 is an FPT algorithm for H -FREE DELETION [VC] using $\mathcal{O}(2^K h^K K^{h+1} h! h^2 n)$ time or alternatively $\mathcal{O}(2^K h^K K! K h! h^2 n)$ time, where $|V(H)| = h$, H contains at least one edge, and K is the vertex cover number of the input graph.

Proof. Let us first go into detail on the correctness of the algorithm. Assume the algorithm returns YES for some instance G, H, ℓ, X where $|X| = K$ and $|V(H)| = h$. The only way the algorithm returns YES, is if in some partition of X into S and Y the BRANCH function returns YES. The BRANCH function only returns YES if any recursive call returns YES, or when $B = \emptyset$ and $i = h$. As the latter is the only base case, this must have occurred for this instance. As i starts at 1 and is only ever incremented, we can conclude that for every i at some point $B = \emptyset$ while $|S| \leq \ell$. The algorithm calls on FINDH for every i to find if there is an occurrence of H with i vertices outside of $Y \cup S$ and $|V(H)| - i$ vertices in Y . By Lemma 13, FINDH correctly finds occurrences of H where i vertices are outside Y . As the algorithm returned YES, FINDH must have returned an empty set for each i at some point, and so no occurrences of H are present in the graph $G[V \setminus S]$ (otherwise, such an occurrence must have i vertices outside $Y \cup S$ for some i). This means that the algorithm is correct in returning YES.

For the other direction, assume that for an instance G, H, ℓ, X where $|X| = K$ there exists a smallest set S_{opt} such that $G[V \setminus S_{opt}]$ is H -free and $|S_{opt}| \leq \ell$. Then S_{opt} must contain some part of the vertex cover X , and as we enumerate all possibilities, the algorithm considers this option. As $G[V \setminus S_{opt}]$ is H -free, clearly, for every set B the function FINDH finds, at least one vertex in B is also in S_{opt} . As we branch on each possibility of the vertices in B , the algorithm also considers exactly the option where the set S in the algorithm is a subset of S_{opt} . This means there is a branch where the algorithm terminates with $S = S_{opt}$, which means it returns YES as $G[V \setminus S_{opt}]$ is H -free. We conclude that the algorithm solves H -FREE DELETION correctly.

Let us analyse the running time of the algorithm. There are $\mathcal{O}(2^K)$ possible partitions of X into S and Y . Checking whether H is contained in Y takes $\mathcal{O}(\binom{K}{h}h!h^2)$ time. Because H contains at least one edge, we can assume that $\ell \leq K$, as otherwise X is a trivial solution. The function BRANCH is called in worst case $\mathcal{O}(h^\ell) = \mathcal{O}(h^K)$ times (branching on at most h vertices each time). FINDH is called at most once for every i in every branch. By Lemma 13, FINDH runs in

$\mathcal{O}\left(\binom{h}{i}[i^2 + \binom{K}{h-i}(h-i)!((h-i)^2 + Kn + (h-i)in)]\right)$ time. Here is where there is some variance in how we round these complexities, namely when concerning e.g. $\binom{K}{h}$. This is because we can both say $\binom{K}{h} = \mathcal{O}(K^h)$, and $\binom{K}{h} = \mathcal{O}(K!)$. Which of these is a tighter bound comes down to the value of h in comparison to K . The total complexity of the algorithm comes down to

$$\mathcal{O}\left(2^K \left(\binom{K}{h} h! h^2 + h^K \sum_{i=1}^h \binom{h}{i} [i^2 + \binom{K}{h-i} (h-i)! ((h-i)^2 + Kn + (h-i)in)]\right)\right)$$

time, which we can shorten to either $\mathcal{O}(2^K h^K K^{h+1} h! h^2 n)$ or $\mathcal{O}(2^K h^K K! K h! h^2 n)$ time. \square

FINDH is adaptable to the streaming setting, as is the complete algorithm, see Algorithm 8. All the actions FINDH takes are local inspection of edges, and many enumeration actions, which lend itself well to usage of the AL streaming model. The number of passes of the streaming version is closely related to the running time of the non-streaming algorithm. This then leads to the full find-and-branch procedure.

It should be clear that the functionality of Algorithm 8 is the same as that of Algorithm 7, but translated to the streaming model using as little memory as possible. Once again we make use of dictionary orderings, see Definition 4 for the formal definition.

Algorithm 8 *H*-FREE DELETION STREAM(Graph $G = (V, E)$ in the AL model, integer ℓ , Vertex Cover $X \subseteq V(G)$).

```

1:  $S \leftarrow \text{FIRST}(\mathcal{X}_{\leq \ell})$ 
2: while  $S \in \mathcal{X}_{\leq \ell}$ ,  $S \neq \clubsuit$  do
3:    $Y \leftarrow X \setminus S$ 
4:    $S' \leftarrow S$ 
5:   if  $\neg \text{CHECK}(H, Y)$  then
6:     if  $\text{BRANCH}(S', Y, 1)$  then
7:       return YES
8:    $S \leftarrow \text{DICT}_{\mathcal{X}_{\leq \ell}}(\text{NEXT}(S))$ 
9: return NO

```

▷ If any returns YES, we also return YES

```

10: function CHECK(set to find  $H$ , search space  $Y$ )
11:    $P \leftarrow \text{FIRST}(\mathcal{Y}_h)$ 
12:   while  $P \in \mathcal{Y}_h$ ,  $P \neq \clubsuit$  do
13:     for each Permutation  $p$  of the vertices of  $H$  do
14:       Use a pass to check if  $p$  matches  $P$ 
15:       if  $p$  matches  $P$  then return YES
16:      $P \leftarrow \text{DICT}_{\mathcal{Y}_h}(\text{NEXT}(P))$ 
17:   return NO

```

▷ Go to the next p if some edges do not match

```

18: function BRANCH(solution set  $S$ , forbidden set  $Y \subseteq X$ , integer  $i$ )
19:    $B \leftarrow \text{FINDH}(S, Y, i)$ 
20:   if  $B = \emptyset$  and  $i = h$  then return YES
21:   else if  $B = \emptyset$  then  $\text{BRANCH}(S, Y, i + 1)$ 
22:   else if  $|S| = \ell$  then return NO
23:   else
24:     for each  $v \in B \setminus Y$  do
25:       if  $\text{BRANCH}(S \cup \{v\}, Y, i)$  then return YES

```

▷ Try to find an H with i vertices outside Y
 ▷ No H found
 ▷ Found an H but cannot remove it

```

26: function FINDH(solution set  $S$ , forbidden set  $Y \subseteq X$ , integer  $i$ )
27:   for each Set  $O$  of  $i$  vertices of  $H$  that can be outside  $Y$  do
28:     Denote  $H' = H \setminus O$ , and  $h' = |V(H')|$ 
29:      $P \leftarrow \text{FIRST}(\mathcal{Y}_{h'})$ 
30:     while  $P \in \mathcal{Y}_{h'}$ ,  $P \neq \clubsuit$  do
31:       for each Permutation  $p$  of the vertices of  $H'$  do
32:         Use a pass to check if  $p$  matches  $P$ 
33:         if  $p$  matches  $P$  then
34:            $S' \leftarrow \emptyset$ ,  $O' \leftarrow O$ 
35:           for each Vertex  $v \in V \setminus (S \cup Y)$  do
36:             Check the edges/non-edges towards  $H' \in Y$ 
37:             if  $v$  is equivalent to some  $w \in O'$  for  $H'$  then
38:                $S' \leftarrow S' \cup \{v\}$ ,  $O' \leftarrow O' \setminus \{w\}$ 
39:               if  $O' = \emptyset$  then return  $S' \cup P$ 
40:              $P \leftarrow \text{DICT}_{\mathcal{Y}_{h'}}(\text{NEXT}(P))$ 
41:   return  $\emptyset$ 

```

▷ Check non-edges in H
 ▷ We found an H occurrence
 ▷ No occurrence of H found

Theorem 15. We can solve *H*-FREE DELETION [VC] in the AL model, where H contains at least one edge, using $\mathcal{O}(2^K h^{K+2} K^h h!)$ or alternatively $\mathcal{O}(2^K h^{K+2} K! h!)$ passes and $\mathcal{O}((K + h^2) \log n)$ space, where $|V(H)| = h$ and K is the vertex cover number of the input graph.

Proof. The algorithm in question is Algorithm 8. When we say ‘check if p matches P ’, this means the permutation p of vertices in H or H' specifies an order of those vertices, and we check if associating the vertices of P with those in H or H' in that order contains exactly those edges/non-edges between vertices in P that are in H or H' .

Let it be clear from the algorithm that the approach to solving H -FREE DELETION has not changed from Algorithm 7. Therefore, if the graph stream is handled correctly, we can conclude that H -FREE DELETION is solved correctly. As we have H in memory, we only require to use passes of the stream to determine (parts of) X and G . The dictionary orderings require no passes because we have the vertices of X in memory. The only places where we require passes of the stream thus are when concerning edges of the vertex cover, and edges/vertices in the rest of the graph ($V \setminus X$). Notice that at such points in Algorithm 8 we correctly mention the use of a pass. The loop over all vertices in $V \setminus (S \cup Y)$ only requires one pass because of the use of the AL model. Therefore, the algorithm is a correct adaptation of Algorithm 7 to the streaming model.

What remains is to analyse the number of passes and memory use. Let us analyse the memory use per function. The entire algorithm keeps track of the vertex cover X , and the forbidden graph H . Denote $|X| = K$ and $|V(H)| = h$. The vertex cover uses $\tilde{O}(K)$ bits, and because we save the entirety of H we use $\tilde{O}(h^2)$ bits. The main function uses the set S of at most ℓ elements of X , likewise S' , and the set Y of size $O(K)$. The function CHECK uses a set P of h elements, a permutation p of h elements, and $O(1)$ bits to check if p matches P , as it can stop when it does not match. The function BRANCH uses a set B of at most h elements and increases the size of the set S , but only when it does not exceed ℓ elements. The function FINDH uses sets O, H', P, O', p , all of at most h vertices. It also uses $O(1)$ bits to check if p matches P , and $O(h)$ bits to save the adjacencies to H' to check if some vertex matches one in O' . S' can only contain an element for each element in O' . Therefore, S' contains at most h elements.

We can conclude the memory use of a single branch is bounded by $O((K + h^2) \log n)$ bits. However, in this algorithm we branch on h options, which is not a constant. Therefore, to be able to return out of recursion when branching and continue where we left off, we need to save the set B , or recompute it when we return. Saving the sets B takes $\tilde{O}(hK)$ bits because we have at most K active instances. Alternatively, recomputing B adds a factor h to the number of passes. Then we only need to save a counter which branch to do next, and call FindH again to find the same set B . Seeing as we aim to be memory efficient, we opt for this second option here.

The number of passes used by the algorithm is closely aligned with the running time of Algorithm 7. There are only three places in which we use a pass of the stream, namely, line 14 in the function CHECK, and line 32 and line 35 in the function FINDH. The loop of line 35 requires only one pass because the stream is given in the Adjacency List model. The number of passes is clearly dominated by the number of times the passes in line 32/35 are used. Now we can use the same analysis as for Algorithm 7, but make some nuances in the running time, as we have to distinguish running time which leads to more passes and running time which will be ‘hidden’ in the allowed unbounded computation. Consider the running time of FINDH, as given by Lemma 13, $O\left(\binom{h}{i}[i^2 + \binom{K}{h-i}(h-i)!((h-i)^2 + (h-i)in)]\right)$. The running time factors i^2 and $(h-i)^2$ are checks over some amount of edges, which will be hidden by the unbounded computation. The factor $(h-i)in$ comes from the finding of vertices that fit the current form of H , and can be done in one pass, which means this factor falls away as well. We now have that the FINDH function costs us $O\left(\binom{h}{i}\binom{K}{h-i}(h-i)!\right)$ passes. This can be shortened to $O(K^h h!)$ by expanding the $\binom{h}{i}$ factor, and bounding $\binom{K}{h-i} = O(K^h)$. We can also bound this by $O(K!h!)$ as discussed before. As FINDH is called for every $1 \leq i \leq h$ in every branch, we get an extra h factor in the total number of passes. Another factor h is added for recomputing B when returning out of recursion at every step. This means the total number of passes comes down to either $O(2^K h^{K+2} K^h h!)$ or $O(2^K h^{K+2} K! h!)$. \square

3.3. Towards Π -free deletion

An issue with extending the previous approach to the general Π -FREE DELETION problem is the dependence on the maximum size h of the graphs $H \in \Pi$. Without further analysis, we have no bound on h . However, we can look to the preconditions used by Fomin et al. [22] on Π in e.g. Theorem 2 to remove this dependence.

The first precondition is that the set $\Pi' \subseteq \Pi$ of graphs that are vertex-minimal with respect to Π have size bounded by a function in K , the size of the vertex cover. That is, for these graphs $H \in \Pi'$ we have that $|V(H)| \leq f(K)$, where $f(K)$ is some function. We can prove that it suffices to only remove vertex-minimal elements of Π to solve Π -FREE DELETION, see Lemma 16. Note that Fomin et al. [22] require that this is a polynomial, we have no need to demand this.

Lemma 16. *Let Π be some graph property, and denote the set of vertex-minimal graphs in Π with Π' . Let G be some graph and $S \subseteq V(G)$ some vertex set. Then $G[V(G) \setminus S]$ is Π -free if and only if $G[V(G) \setminus S]$ is Π' -free.*

Proof. Assume the preconditions in the lemma, and assume that $G[V(G) \setminus S]$ is not Π' -free. As $\Pi' \subseteq \Pi$, clearly, $G[V(G) \setminus S]$ is not Π -free.

Now assume that $G[V(G) \setminus S]$ is Π' -free. Assume there is some $H \in \Pi$ such that $H \in G[V(G) \setminus S]$ (that is, H is isomorphic to an induced subgraph of $G[V(G) \setminus S]$). As the graph is Π' -free, H is a non-vertex-minimal graph with respect to Π . By definition of vertex-minimal, removal of a specific set of one or more vertices of H results in a vertex-minimal graph $I \in \Pi'$. But if we ignore the same set of vertices in $H \in G[V(G) \setminus S]$, clearly, $I \in G[V(G) \setminus S]$. This contradicts our assumption that

$G[V(G) \setminus S]$ is Π' -free. Therefore, there cannot be a $H \in \Pi$ such that $H \in G[V(G) \setminus S]$, which means that $G[V(G) \setminus S]$ is Π -free. \square

If we also assume that we know the set Π' , we obtain the following result.

Theorem 17. *If Π is a graph property such that:*

- (i) *we have explicit knowledge of $\Pi' \subseteq \Pi$, which is the subset of q graphs that are vertex-minimal with respect to Π , and*
- (ii) *there is a non-decreasing function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that all graphs $G \in \Pi'$ satisfy $|V(G)| \leq f(K)$, and*
- (iii) *every graph in Π contains at least one edge,*

then Π -FREE DELETION [VC] can be solved using $\mathcal{O}(q \cdot 2^K \cdot f(K)^K \cdot K! \cdot K \cdot f(K)! \cdot f(K)^2 \cdot n)$ time, where K is the vertex cover number of the input graph.

Proof. This result can be achieved by adjusting Algorithm 7 to search for each of the q graphs in Π' instead of only H . This increases the complexity by a factor q as we search for more graphs than just one. Then, using Theorem 14 and Lemma 16, the theorem follows. \square

Note that we require explicit knowledge of Π' and $f(K)$ to achieve Theorem 17. Also note that we chose to write the $K!$ alternative here, as likely $K^{f(K)} > K!$ (i.e. for $f(K) = K$, a linear function, we have that $K^{f(K)} = K^K > K!$).

We argue this algorithm is essentially tight, under the Exponential Time Hypothesis (ETH) [42], by augmenting a reduction by Abu-Khzam et al. [43]. Recall that the INDUCED SUBGRAPH ISOMORPHISM problem is defined as follows: given two graphs G_1 (the pattern) and G_2 (the host), does G_2 contain an induced subgraph that is isomorphic to G_1 . Abu-Khzam et al. [43] proved that INDUCED SUBGRAPH ISOMORPHISM[VC] has no $2^{o(K \log K)} \text{poly}(|G_1|, |G_2|)$ time algorithm unless ETH fails, where K is the sum of the vertex cover numbers of G_1 and G_2 . Crucially, in Π -FREE DELETION [VC] the graph property Π is part of the problem specification, not of the input. Hence, the original reduction of Abu-Khzam et al. [43] does not directly apply. We strengthen their reduction to yield the following result.

Theorem 18. *There is a graph property Π for which we cannot solve Π -FREE DELETION [VC] in $2^{o(K \log K)} \text{poly}(n)$ time, unless ETH fails, where K is the vertex cover number of G , even if each graph that has property Π has size quadratic in its vertex cover number.*

Proof. We augment the reduction by Abu-Khzam et al. [43]. Their reduction is from the $k \times k$ PERMUTATION CLIQUE problem. In this problem, the input is a graph $G = (V, E)$ on the vertex set $V = [k] \times [k]$ and the question is whether it has a clique that contains exactly one vertex from each row and column. That is, whether there exists a set $C \subseteq [k] \times [k]$ such that for each distinct $(i, j), (i', j') \in C$, it holds that $i \neq i', j \neq j'$, and $(i, j), (i', j') \in E$. Then we say that G has a permutation clique of size k . Lokshtanov et al. [44] proved that this problem has no $2^{o(k \log k)}$ time algorithm, unless ETH fails.

Abu-Khzam et al. show that, given an instance of (G, k) of $k \times k$ PERMUTATION CLIQUE, in polynomial time an equivalent instance of INDUCED SUBGRAPH ISOMORPHISM[VC] can be constructed with pattern graph G_1 and host graph G_2 that both have a vertex cover of size $K = O(k)$. By inspection of their reduction, we observe that G_1 has size quadratic in k . Moreover, both G_1 and G_2 have a unique clique (called \tilde{D}_r and D_r respectively) of size 6 with designated vertex \tilde{r} and r respectively that is the only vertex that has edges to vertices outside the clique. The isomorphism in the proof of equivalence maps \tilde{D}_r to D_r and \tilde{r} to r .

We augment the construction of G_1 and G_2 as follows. Add a clique of size 7 to G_1 and G_2 (called \tilde{A} and A respectively), with a designated vertex (called \tilde{a} and a) respectively. Then add a path \tilde{P} (P) of length k from \tilde{a} (a) to a designated vertex $\tilde{d} \neq \tilde{r}$ ($d \neq r$). Since \tilde{A} (A) is the unique clique of size 7 in G_1 (G_2), any isomorphism from the augmented versions of G_1 to G_2 will map \tilde{A} to A . Since a is the only vertex of A of degree 7, \tilde{a} will be mapped to a . Repeating such arguments, \tilde{P} will be mapped to P , \tilde{d} to d , \tilde{D}_r to D_r , and \tilde{r} to r . Then the remainder of the proof of Abu-Khzam et al. carries over without further modification, and we again obtain an instance of INDUCED SUBGRAPH ISOMORPHISM[VC] that is equivalent to the original instance of $k \times k$ PERMUTATION CLIQUE.

By inspecting the reduction of Abu-Khzam et al. and the above augmentation, we note that the construction of the pattern graph G_1 is actually independent of G , and only depends on k . So denote by G_1^ℓ the pattern graph constructed when $k = \ell$. Moreover, we note that the above isomorphism between \tilde{P} and P can only occur when their length is the same. Hence, an induced subgraph of G_2 is isomorphic to G_1^ℓ only if $\ell = k$. Now let Π be the set of graphs isomorphic to G_1^ℓ for any $\ell \geq 1$. Observe that this Π is independent of the original instance of INDUCED SUBGRAPH ISOMORPHISM and can thus be part of the problem specification of Π -FREE DELETION [VC].

To complete the argument, consider any algorithm for Π -FREE DELETION [VC] running in time $2^{o(K \log K)} \text{poly}(n)$. Let (G, k) be any instance of $k \times k$ PERMUTATION CLIQUE. Apply the reduction of Abu-Khzam et al. to G with the above augmentation to obtain the host graph G_2 . As argued above, G_2 contains an induced subgraph isomorphic to G_1^ℓ for any ℓ if and only if G has a permutation clique of size k . Hence, the instance of Π -FREE DELETION [VC] for the constructed property Π with input graph G_2 has solution size at least 1 if and only if G_2 contains an induced subgraph isomorphic to G_1^ℓ for any ℓ , if and only

if G has a permutation clique of size k . Hence, using this transformation on an instance of $k \times k$ PERMUTATION CLIQUE and then applying the assumed algorithm for Π -FREE DELETION [VC], we obtain an algorithm for $k \times k$ PERMUTATION CLIQUE with running time $2^{o(k \log k)} \text{poly}(n)$. Such an algorithm cannot exist unless ETH fails. \square

Next, we look to further improve the bound of Theorem 17. Note that so far, we have made no use at all of the characterization by few adjacencies of Π , as in Theorem 2. We now argue that there may be graphs in Π that cannot occur in G simply because it would not fit with the vertex cover.

Lemma 19. *If Π is a graph property such that*

- (i) *every graph in Π is connected and contains at least one edge, and*
- (ii) *Π is characterized by c_Π adjacencies,*

and G is some graph with vertex cover X , $|X| = K$, and $S \subseteq V(G)$ some vertex set. Then $G[V(G) \setminus S]$ is Π -free if and only if $G[V(G) \setminus S]$ is Π' -free, where $\Pi' \subseteq \Pi$ contains only those graphs in Π with $\leq (c_\Pi + 1)K$ vertices.

Proof. Assume the preconditions in the lemma. When $G[V(G) \setminus S]$ is Π -free, it must also be Π' -free, as $\Pi' \subseteq \Pi$ by definition.

Now assume that $G[V(G) \setminus S]$ is not Π -free, let us say some $H \in \Pi$ occurs in the graph. Consider a vertex v of the graph H . Because Π is characterized by c_Π adjacencies, there exists a set $D \subseteq V(H)$ with $|D| \leq c_\Pi$ such that changing the adjacencies between v and $V(H) \setminus D$ does not change the presence of $H \in \Pi$. Remove all adjacencies existing between v and $V(H) \setminus D$. Then $\deg(v) \leq c_\Pi$. Our new version of H is still contained in Π , so we can repeat this process for every vertex in H . But then every vertex v in H has $\deg(v) \leq c_\Pi$. Given that H can contain at most K vertices of the vertex cover in G , each with degree at most c_Π , we know that this edited version of H can have at most $(c_\Pi + 1)K$ vertices. Although this version of H is still in Π , it might not exactly be in G , as we might have deleted essential adjacencies. However, changing all the adjacencies did not increase or decrease the number of vertices, as every graph in Π is connected and H remains in Π at every step. Therefore, every H occurring in $G[V(G) \setminus S]$ is in Π' , and so the graph is not Π' -free either. \square

The precondition that every graph in Π is connected is necessary to obtain this result. If not every graph in Π is connected, the removal of adjacencies might leave a vertex without edges, but then this vertex might still be required for the presence in Π , which is problematic. Seeing that we want to bound the size of possible graphs in Π , to be able to bound the size of graphs we need that the vertex cover together with c_Π gives us information on the size of the graph, which is not the case for a disjoint union of graphs.

We can use Lemma 19 in combination with Theorem 17 to obtain a new result. Alternatively, using a streaming version of the algorithm instead of the non-streaming one, immediately also provides a streaming result.

Theorem 20. *Given a graph G with vertex cover X , $|X| = K$, if Π is a graph property such that*

- (i) *every graph in Π is connected and contains at least one edge, and*
- (ii) *Π is characterized by c_Π adjacencies, and*
- (iii) *we have explicit knowledge of $\Pi' \subseteq \Pi$, which is the subset of q graphs of at most size $(c_\Pi + 1)K$ that are vertex-minimal with respect to Π ,*

then Π -FREE DELETION [VC] can be solved using $\mathcal{O}(q \cdot 2^K \cdot ((c_\Pi + 1)K)^K \cdot K! \cdot K \cdot ((c_\Pi + 1)K)! \cdot ((c_\Pi + 1)K)^2 \cdot n)$ time. Assuming $c_\Pi \geq 1$ this can be simplified to $\mathcal{O}(q \cdot 2^K \cdot c_\Pi^K \cdot K^{K+3} \cdot K! \cdot (c_\Pi K)! \cdot n)$ time. In the streaming setting, Π -FREE DELETION [VC] can be solved using $\mathcal{O}(q \cdot 2^K \cdot c_\Pi^K \cdot K^{K+2} \cdot K! \cdot (c_\Pi K)!)$ passes in the AL streaming model, using $\tilde{\mathcal{O}}(q \cdot (c_\Pi \cdot K)^2)$ space.

Proof. Given some Π characterized by c_Π adjacencies, where every graph in Π is connected, we can see that through Lemma 19 we only need to consider those graphs with size $\leq (c_\Pi + 1)K$ in Π , and with this subset, using Theorem 17 where $f(K) = (c_\Pi + 1)K$, the fpt part of the theorem follows.

For the streaming part, we use Theorem 15 with the above analysis. The main factor $q \cdot (c_\Pi \cdot K)^2$ of the memory usage comes from the fact that we need to explicitly store Π' . \square

The required explicit knowledge of Π' might give memory problems. That is, we have to store Π' somewhere to make this algorithm work, which takes $\tilde{\mathcal{O}}(q \cdot (c_\Pi + K)^2)$ space. Note that q can range up to $K^{\mathcal{O}(K)}$. We adapt the streaming algorithm to the case when we have oracle access to Π in Section 3.4.

3.4. Π -free deletion without explicit Π

One of the issues in Section 3 is that having the graph property Π explicitly saved might cost us a lot of memory. To circumvent this, we can assume to be working with some *oracle*, which we can call to learn something about Π , and

generalizes saving Π explicitly. For some Π it may be easier to check membership than using the explicit graphs that are forbidden, hence an oracle would be applicable. In general, if we have an oracle algorithm \mathcal{A} , let us assume that it takes a graph G as input as a stream. We then denote $P_{\mathcal{A}}(n)$, $M_{\mathcal{A}}(n)$ as respectively the number of passes and the memory use of the oracle algorithm \mathcal{A} when called on a streamed (sub)graph G with n vertices. In this section, we will discuss two different oracles and their use for Π -FREE DELETION [VC]. We also assume that the graphs in Π have a maximum size ν , and that ν is known.

As the oracle algorithms take a stream as an input, but we usually might want to pass a subgraph of G , we require a full pass over the stream for each call to an oracle algorithm. This way, we can select exactly the sub-stream that corresponds to the graph we wish to pass to the oracle algorithm while being memory-less (for each edge we only need to decide whether or not it should appear in the oracle input, and pass it to the oracle if it should). Actually, if the oracle algorithm uses multiple passes over its input, we need to generate this stream every time it does so.

The first oracle model is where our oracle algorithm \mathcal{A}_1 , when called on a graph G , returns whether or not $G \in \Pi$. Our approach then has to be different from Algorithm 8, as in Algorithm 8 we rely on knowing whether a part of the vertex cover is contained in some $H \in \Pi$. This is not information the oracle can give us. The general idea is the following: we still branch on what part of the vertex cover is in the solution, and consider every subset of (the remaining part of) the vertex cover in each branch. To avoid having to test every combination of vertices outside the vertex cover together with this subset, we consider the notion of *twin vertices*. Two vertices $u, v \notin X$, where X is a vertex cover, are called *twins* when $N(u) = N(v) \subseteq X$, their neighbourhoods are equal. If EC is a set of vertices where each pair u, v are twins, and EC is maximal under this property, we call EC an *equivalence class*. When trying to test which graphs might be in Π , we can ignore twin vertices if we have tested one of them before. Also, if we delete a vertex from an equivalence class, we may need to delete the entire equivalence class. We can identify twins easily in the AL streaming model. The issue here is that saving the equivalence classes is not memory efficient. Nonetheless, we use this idea here in Algorithm 9.

In Algorithm 9, the set EC saves the sizes of each equivalence class, and can be seen as a set of key-value pairs (*key*, *val*), where *key* is a K -bit string representing the adjacencies towards the vertex cover, and *val* is the number of vertices in this equivalence class. Notice that we can find the set EC using a single pass over the stream. This can be done by, for each vertex, locally saving its adjacencies as a K -bit string, and then finding and incrementing the correct counter in EC . We only need one pass to find EC in its entirety.

With slight abuse of notation, we write $\text{Dict}_{EC \leq k}$ to denote a dictionary ordering on choosing $\leq k$ ‘vertices’ out of the 2^K equivalence classes such that if an equivalence class *key* is chosen twice, then *val* is at least 2. This is essentially enumerating picking $\leq k$ vertices out of $V \setminus X$ except that we do not pick vertices explicitly, but we pick the equivalence classes they are from. An entry of this ordering is a set of key-value pairs (*key*, *count*), such that *key* corresponds to some equivalence class, and *count* is the number of vertices we pick out of this equivalence class.

Algorithm 9 Π -FREE DELETION WITH \mathcal{A}_1 (Graph $G = (V, E)$ in the AL model, integer ℓ , integer ν , Vertex Cover $X \subseteq V(G)$).

```

1:  $S \leftarrow \text{FIRST}(\mathcal{X}_{\leq \ell})$ 
2: while  $S \in \mathcal{X}_{\leq \ell}$ ,  $S \neq \clubsuit$  do
3:    $Y \leftarrow X \setminus S$ 
4:    $S' \leftarrow S$ 
5:   if  $\forall Y' \subseteq Y: \mathcal{A}_1(Y') = \text{false}$  then                                     ▷ Test if Y is  $\Pi$ -free
6:      $EC \leftarrow$  Count the sizes of the equivalence classes of vertices in  $V \setminus X$  with their adjacencies towards Y   ▷ Use a pass for this
7:     if  $\text{SEARCH}(S', Y, EC)$  then return YES
8:      $S \leftarrow \text{Dict}_{\mathcal{X}_{\leq \ell}}(\text{NEXT}(S))$ 
9: return NO

10: function  $\text{SEARCH}$ (solution set  $S$ , forbidden set  $Y \subseteq X$ , equivalence class sizes  $EC$ )
11:    $J \leftarrow \text{FIRST}(\mathcal{Y}_{\leq \nu})$ 
12:    $I \leftarrow \text{FIRST}(EC_{\leq (\nu - |J|)})$ 
13:   while  $J \in \mathcal{Y}_{\leq \nu}$ ,  $J \neq \clubsuit$  do
14:     while  $I \in EC_{\leq (\nu - |J|)}$ ,  $I \neq \clubsuit$  do
15:       if  $\mathcal{A}_1(I, J)$  then                                                                                               ▷  $I \cup J \in \Pi$ 
16:         if  $|S| \geq \ell$  then return NO                                                                                   ▷ No budget to branch
17:         else
18:           for each  $(k, \text{count}) \in I$  do
19:             Update  $EC$  such that  $(k, \text{val}) \leftarrow (k, \text{count} - 1)$                                              ▷ Remove all but  $\text{count} - 1$  vertices from class  $k$ 
20:             if  $|S| + \text{val} - (\text{count} - 1) > \ell$  then return NO
21:             else
22:                $VS \leftarrow$  Find  $\text{val} - (\text{count} - 1)$  vertices that belong to class  $k$                                    ▷ Uses a pass
23:               if  $\text{SEARCH}(S \cup VS, Y, EC)$  then return YES                                                           ▷ Branch
24:              $I \leftarrow \text{Dict}_{EC_{\leq (\nu - |J|)}}(\text{NEXT}(I))$ 
25:            $J \leftarrow \text{Dict}_{\mathcal{Y}_{\leq \nu}}(\text{NEXT}(J))$ 
26:            $I \leftarrow \text{FIRST}(EC_{\leq (\nu - |J|)})$ 
27: return YES

```

Theorem 21. *If Π is a graph property such that the maximum size of graphs in Π is ν , and \mathcal{A}_1 is an oracle algorithm that, when given subgraph H on h vertices, decides whether or not H is in Π using $P_{\mathcal{A}_1}(h)$ passes and $M_{\mathcal{A}_1}(h)$ bits of memory, then Algorithm 9 solves Π -FREE DELETION [VC] on a graph G on n vertices given as an AL stream with a vertex cover of size K using $\mathcal{O}(3^K \nu^K 2^{\nu K} [1 + P_{\mathcal{A}_1}(\nu)])$ passes and $\mathcal{O}(2^K(K + \log n) + \nu \log n + M_{\mathcal{A}_1}(\nu))$ bits of memory.*

Proof. Let us argue on the correctness of Algorithm 9. In essence, the algorithm tries every option of how an occurrence of a graph in Π can occur in G , by enumerating all subsets of the vertex cover (J) combined with vertices from outside the vertex cover (I). This process is optimized by use of the equivalence classes, removing multiple equivalent vertices at once to eliminate such an occurrence. If the algorithm returns YES, then for every combination of I and J , no occurrence of a graph in Π was found or branching happened on any such occurrence. Each branch eliminates the found occurrence of a graph in Π , as enough vertices are removed from an equivalence class to make the same occurrence impossible. As removing vertices from the graph cannot create new occurrences of graphs in Π , this means that every occurrence found was removed, and no new occurrences were created. But then there is no occurrence of a graph in Π , as every combination of such a graph occurring was tried. So the graph is Π -free. If the algorithm returns NO, then in every branch at some point we needed to delete vertices to remove an occurrence of a graph in Π but had not enough budget left. Removing any less than $\text{val} - (\text{count} - 1)$ vertices in line 22 results in at least count vertices of that equivalence class remaining, which means the same occurrence of a graph in Π persists. Therefore, the conclusion that we do not have enough budget is correct, and so there is no solution to the instance. We can conclude that the algorithm works correctly.

Let us analyse the space usage. Almost all sets used are bounded by K entries. The exceptions are the sets EC , I , and J . J contains at most ν vertices, and so uses at most $\mathcal{O}(\nu \log n)$ bits of space. EC contains at most 2^K key-value pairs, each using $K + \log n$ bits, so EC uses $\mathcal{O}(2^K(K + \log n))$ bits. The set I contains only subsets of EC . The oracle algorithm \mathcal{A}_1 uses at most $M_{\mathcal{A}_1}(\nu)$ bits. Therefore, the memory usage of this algorithm is $\mathcal{O}(2^K(K + \log n) + \nu \log n + M_{\mathcal{A}_1}(\nu))$ bits of space. Note that we branch on at most ν options together spanning at most ν vertices, which is a constant factor in the memory use (for remembering what we branch on when returning out of recursion). The algorithm also repeats the process for each I and J in each branch to avoid having to keep track of these sets in memory when returning from recursion.

It remains to show the number of passes used by the algorithm. Let it be clear that the number of passes made by the algorithm is dominated by the number of calls made to \mathcal{A}_1 , which requires at least one pass. The number of calls made is heavily dependent on the total number of branches in the algorithm, together with how many options the sets I and J can span. This total can be concluded to be $\mathcal{O}(3^K \nu^K 2^{\nu K})$. Let us elaborate on this. The factor 3^K comes from the fact that any vertex in the vertex cover is either in J , in S , or in Y (and not in J). The factor ν^K comes from the worst case branching process, as we branch on at most ν different deletions, and we branch at most $\ell \leq K$ times. The remaining factor is the number of options we have for the set I , which picks, for each size i between 1 and ν , i times between at most 2^K options (the equivalence classes). This number of options is bounded by $\mathcal{O}(2^{\nu K})$. This means the total number of passes is bounded by $\mathcal{O}(3^K \nu^K 2^{\nu K} [1 + P_{\mathcal{A}_1}(\nu)])$. \square

Next, we look at a different oracle model, as the memory complexity of $\mathcal{O}(2^K K)$ is not ideal.

The second oracle model is where our oracle algorithm \mathcal{A}_2 , when called on a graph G , returns whether or not G is induced Π -free (that is, it returns YES when it is induced Π -free, and NO if there is an occurrence of a graph in Π in G).

We observe that both oracle models \mathcal{A}_1 and \mathcal{A}_2 admit a streaming algorithm parameterized by solution size by applying the approach of Cai [18]. By enumerating all subsets of vertices of size $\leq \nu$ we can find an occurrence of a graph in Π in $\mathcal{O}(n^\nu)$ calls to either oracle. Then we can branch on the $\leq \nu$ options for deleting a vertex and repeat. We get an extra ν factor for recomputing the set to branch on each time we return out of recursion, to be memory-efficient.

Observation 22. *If Π is a graph property such that the maximum size of graphs in Π is ν , for oracle model $\mathcal{A} = \mathcal{A}_1, \mathcal{A}_2$ using $P_{\mathcal{A}}(h)$ passes and $M_{\mathcal{A}}(h)$ bits of space when called on a graph on h vertices, the algorithmic approach by Cai [18] can be implemented as a streaming algorithm to solve Π -FREE DELETION [ℓ] using $\mathcal{O}(\nu^{\ell+1} n^\nu \cdot [1 + P_{\mathcal{A}}(\nu)])$ passes and $\mathcal{O}((\nu + \ell) \log n + M_{\mathcal{A}}(\nu))$ bits of space.*

Another approach in which we do not demand ν is known comes from the algorithm that Theorem 12 proposes. We work with oracle model \mathcal{A}_2 , which can check whether a given stream is Π -free. We can enumerate the $(2^K + 1) \cdot K$ possible solutions by saving $\mathcal{O}(K^2)$ bits corresponding to the deletion of $\mathcal{O}(K)$ vertices of equivalence classes, which each require $\mathcal{O}(K)$ bits to represent. We can enumerate all these $(2^K + 1) \cdot K$ options by using Dictionary Orderings. This immediately provides an algorithm, as for each possible solution, for each pass the oracle requires, we make a pass over the AL stream and pass on to \mathcal{A}_2 only the vertices from an equivalence class when the first i have been skipped (where i corresponds to the number of deletions in this equivalence class according to the current proposed solution). The correctness of this process follows from Theorem 12, and the number of passes comes down to $\mathcal{O}(2^{\mathcal{O}(K^2)} \cdot P_{\mathcal{A}_2}(n))$, with a memory usage of $\mathcal{O}(K^2 + K \log n + M_{\mathcal{A}_2}(n))$ bits.

We have now seen some algorithms for using an oracle to obtain information about Π instead of explicitly saving it. This way, we have gotten closer to a memory-optimal algorithm, but the number of passes over the stream has increased significantly in comparison to algorithms from Section 3.

3.5. Odd cycle transversal

Specific forms of Π -FREE DELETION [VC] allow for improvement over Theorem 20, which we illustrate for the problem of ODD CYCLE TRANSVERSAL [VC]. Note that odd cycle-free and induced odd cycle-free are equivalent.

ODD CYCLE TRANSVERSAL [VC]

Input: A graph G with a vertex cover X , and an integer ℓ .

Parameter: The size $K := |X|$ of the vertex cover.

Question: Is there a set $S \subseteq V(G)$ of size at most ℓ such that $G[V(G) \setminus S]$ contains no induced odd cycles?

The interest in this problem comes from the FPT algorithm using iterative compression provided in [45, Section 4.4], based on work by Reed et al. [46]. Although Chitnis and Cormode [12] have shown how iterative compression can be used in the streaming setting, adapting the algorithm out of Reed et al. seems difficult. The main cause for this is the use of a maximum-flow algorithm, which does not seem to lend itself well to the streaming setting because of its memory requirements. Instead, we present the following approach.

It is well known that a graph without odd cycles is a bipartite graph (and thus 2-colourable) and vice versa. In the algorithm, we guess what part of the vertex cover is in the solution, and then we guess the colouring of the remaining part. Then vertices outside the vertex cover for which not all neighbours have the same colour must be deleted. This step can be done in one pass if we use the AL streaming model. In the same pass, we can also check if the colouring is valid within the vertex cover. If the number of deletions does not exceed the solution size and the colouring is valid within the vertex cover, then the resulting graph is bipartite and thus odd cycle free.

The total number of guesses comes down to $\mathcal{O}(3^K)$ options, as any vertex in the vertex cover is either in the solution, coloured with colour 1 or coloured with colour 2. This directly corresponds to the number of passes, as only one pass is needed per guessed colouring.

The full algorithm is given in Algorithm 10.

Algorithm 10 OCT(Graph $G = (V, E)$ in the AL model, integer ℓ , Vertex Cover $X \subseteq V(G)$).

```

1:  $S \leftarrow \text{FIRST}(\mathcal{X}_{\leq \ell})$ 
2: while  $S \in \mathcal{X}_{\leq \ell}$ ,  $S \neq \clubsuit$  do
3:    $Y \leftarrow X \setminus S$ 
4:    $Y_1 \leftarrow \text{FIRST}(\mathcal{Y})$ 
5:    $Y_2 \leftarrow Y \setminus Y_1$ 
6:   while  $Y_1 \subseteq Y$ ,  $Y_1 \neq \clubsuit$  do
7:      $\text{success} \leftarrow \text{true}$ ,  $S' \leftarrow S$ 
8:     for each  $v \in V \setminus S'$  do
9:       if  $v \in Y$  and  $v \in Y_i$  then
10:        Check that all neighbours of  $v$  in  $Y$  are in  $Y_{3-i}$ 
11:        If one is not,  $\text{success} \leftarrow \text{false}$ 
12:       else
13:        Check if all neighbours of  $v$  in  $Y$  are in the same  $Y_i$ 
14:        If not, and  $|S'| < \ell$ ,  $S' \leftarrow S' \cup \{v\}$ 
15:        Else,  $\text{success} \leftarrow \text{false}$ 
16:     if  $|S'| \leq \ell$  and  $\text{success}$  then return YES
17:      $Y_1 \leftarrow \text{DICT}_{\mathcal{Y}}(\text{NEXT}(Y_1))$ 
18:      $Y_2 \leftarrow Y \setminus Y_1$ 
19:      $S \leftarrow \text{DICT}_{\mathcal{X}_{\leq \ell}}(\text{NEXT}(S))$ 
20: return NO

```

\triangleright Reset local values
 \triangleright Use one pass
 $\triangleright v \in X$
 $\triangleright v \notin X$
 \triangleright Try the next colouring

Theorem 23. Given a graph G given as an AL stream with vertex cover X , $|X| = K$, we can solve ODD CYCLE TRANSVERSAL [VC] using $\mathcal{O}(3^K)$ passes and $\mathcal{O}(K \log n)$ space.

Proof. We claim that Algorithm 10 does exactly this.

Let us first prove the correctness of Algorithm 10. Let $G = (V, E)$ be a graph with vertex cover X , $|X| = K$. Let O , $|O| \leq \ell$ be a solution for ODD CYCLE TRANSVERSAL [VC]. Denote $Y' = X \setminus O$ as the part of X that is not contained in the solution. Because Algorithm 10 enumerates all possibilities for the set Y , at some point it must consider $Y = Y'$. As O is a solution, $G[X \setminus Y']$ must be bipartite, and so admits a proper 2-colouring Y'_1, Y'_2 . For $Y = Y'$, Algorithm 10 considers at some point $Y_1 = Y'_1$ and $Y_2 = Y'_2$ ($Y'_1 \cup Y'_2 = Y' = Y$ because Y'_1 and Y'_2 form a proper 2 colouring), because it considers all possibilities for the set $Y_1 \subseteq Y$. As Y_1 and Y_2 now form a proper 2-colouring of Y , the check in line 10 never fails. The check in line 13 only fails if a vertex outside of X is adjacent to two different coloured vertices in Y . But then this vertex must also be in O , as Y_1 and Y_2 mimic exactly the 2-colouring in $G[X \setminus Y']$. Therefore, Algorithm 10 can at least find O as well, which means it returns YES. For the reverse implication, when Algorithm 10 returns YES, it has found a set S such that $|S| \leq \ell$ and $G[V \setminus S]$ admits a proper 2-colouring given by deterministically adding vertices outside the vertex cover to Y_1 and Y_2 . But then S is a solution to ODD CYCLE TRANSVERSAL [VC]. We can conclude that Algorithm 10 works correctly.

Let us analyse the memory usage of Algorithm 10. All sets used in the algorithm have size at most K ($\ell \leq K$). The only worry is whether or not the checks in lines 10 and 13 require more memory. The first check only requires us to remember in what set the vertex v is contained in, and whether or not we have seen a ‘wrong’ colour yet, which should only take a constant number of bits. The second check merely needs to remember in what set all neighbours up until this point were, which should also only take a constant number of bits. Therefore, the algorithm uses $\mathcal{O}(K \log n)$ bits of memory.

Let us analyse the number of passes of Algorithm 10 does over the stream. Firstly, the for-loop over all vertices only requires a single pass because the checks only need to know what all the neighbours of the current vertex are, which is what the AL stream gives us. The total number of times this for-loop can be executed is bounded by $\mathcal{O}(3^K)$, as any vertex in the vertex cover can either be in S , in Y_1 or in Y_2 . We can conclude that Algorithm 10 uses $\mathcal{O}(3^K)$ passes over the graph stream. \square

If we think about this algorithm, we can notice that often the colouring we guess on the vertex cover is invalid. An alternative approach follows by noting that a connected component within the vertex cover can only have two possible valid colourings. We can exploit this to decrease the number of passes when the number of connected components in the vertex cover is low. This comes at a price: to easily find components of the vertex cover, we store it in memory, which increases the memory complexity. Alternatively, we can use $\mathcal{O}(K)$ passes to find the connected components of the vertex cover in every branch. We formalize this in Algorithm 11 and Theorem 24.

Algorithm 11 OCT-CC(Graph $G = (V, E)$ in the AL model, integer ℓ , Vertex Cover $X \subseteq V(G)$).

```

1:  $S \leftarrow \text{FIRST}(\mathcal{X}_{\leq \ell})$ 
2: while  $S \in \mathcal{X}_{\leq \ell}, S \neq \clubsuit$  do
3:    $Y \leftarrow X \setminus S$ 
4:   Use a pass to find and save the edges in  $Y$ 
5:   Find the connected components and their two 2-colourings
6:   If this fails, move to the next option for  $S$ 
7:   for each Combination of colourings of the CCs do
8:      $\text{success} \leftarrow \text{true}, S' \leftarrow S$ 
9:     for each  $v \in V \setminus (X \cup S')$  do
10:      Check if all neighbours of  $v$  in  $Y$  have the same colour
11:      If not, and  $|S'| < \ell, S' \leftarrow S' \cup \{v\}$ 
12:      Else,  $\text{success} \leftarrow \text{false}$ 
13:     if  $|S'| \leq \ell$  and  $\text{success}$  then return YES
14:    $S \leftarrow \text{DICT}_{\mathcal{X}_{\leq \ell}}(\text{NEXT}(S))$ 
15: return NO

```

\triangleright Does $2^{\#\text{CCs}}$ iterations
 \triangleright Reset local values
 \triangleright Use one pass

Theorem 24. Given a graph G given as an AL stream with vertex cover X , $|X| = K$, Algorithm 11 solves ODD CYCLE TRANSVERSAL [VC] using $\mathcal{O}(3^K)$ passes and $\mathcal{O}(K^2 \log n)$ bits of memory.

Proof. The correctness of Algorithm 11 quickly follows from the correctness of Algorithm 10. Where Algorithm 10 enumerates all 2-colourings of Y , Algorithm 11 only enumerates those which are feasible 2-colourings for Y (combinations of 2-coloured components). This means Algorithm 11 only leaves out colourings which are not feasible anyway, which means that it still works correctly.

The number of passes is heavily dependent on the amount of connected components in $G[Y]$ in an iteration. But, in worst case, this is $|Y|$, which would mean it considers all 2-colourings of Y . But this is a worst case where the behaviour exactly mimics that of Algorithm 10, and so the worst case number of passes is the same.

Let it be clear that the memory use is $\mathcal{O}(K^2 \log n)$ bits, as we save (a part of) the vertex cover with edges to enable easy colouring and component finding. The rest of the sets in the algorithm use $\mathcal{O}(K \log n)$ bits of memory. \square

Let us denote again that Algorithm 11 on paper is strictly worse than Algorithm 10 in worst case complexity. However, we believe there to be many cases where Algorithm 11 can outperform Algorithm 10 because of its behaviour of clever enumeration instead of trying all possibilities. We also want to mention that, instead of saving the vertex cover using $\mathcal{O}(K^2 \log n)$ bits of memory, we could find the connected components of the vertex cover in $\mathcal{O}(K)$ passes in each branch. This increases the number of passes by a factor K , but remains memory-optimal, using $\mathcal{O}(K \log n)$ bits of memory.

4. Lower bounds

We show lower bounds for Π -FREE DELETION. To prove lower bounds for streaming, we can show reductions from problems in communication complexity, as first shown by Henzinger et al. [1]. An example of such a problem is DISJOINTNESS.

DISJOINTNESS

Input: Alice has a string $x \in \{0, 1\}^n$ given by $x_1 x_2 \dots x_n$. Bob has a string $y \in \{0, 1\}^n$ given by $y_1 y_2 \dots y_n$.

Question: Bob wants to check if $\exists 1 \leq i \leq n$ such that $x_i = y_i = 1$. (Formally, the answer is NO if this is the case.)

The following proposition is given and used by Bishnu et al. [13], and gives us one important consequence of reductions from a problem in communication complexity to a problem for streaming algorithms.

Proposition 25. (Rephrasing of item (ii) of [13, Proposition 5.6]) *If we can show a reduction from DISJOINTNESS to problem P in streaming model \mathcal{M} such that in the reduction, Alice and Bob construct one model- \mathcal{M} pass for a streaming algorithm for P by communicating the memory state of the algorithm only a constant number of times to each other, then any streaming algorithm working in the model \mathcal{M} for P that uses p passes requires $\Omega(n/p)$ bits of memory, for any $p \in \mathbb{N}$ [47–49].*

The structure of these reductions is relatively simple: have Alice and Bob construct the input for a streaming algorithm depending on their input to DISJOINTNESS. If we do this in such a manner that the solution the streaming algorithm outputs gives us exactly the answer to DISJOINTNESS, we can conclude that the streaming algorithm must abide the lower bound of DISJOINTNESS.

Chitnis et al. [9, Theorem 6.3] prove hardness for many Π , those that abide to a small precondition. However, Chitnis et al. do not describe in their reduction how Alice and Bob give their ‘input’ as a stream to the algorithm for Π -FREE DELETION, and thus it would apply only to the EA streaming model. However, if we observe the proof closely, we can see it extends to the VA model.

We would also like it to extend to the AL model. However, this requires a slightly stronger precondition on the graph class Π .

Theorem 26. *If Π is a set of graphs such that each graph in Π is connected, and there is a graph $H \in \Pi$ such that*

- *H is a minimal element of Π under the operation of taking subgraphs, i.e., no proper subgraph of H is in Π , and*
- *H has at least two **disjoint** edges,*

then any p -pass (randomized) streaming algorithm working on the AL streaming model for Π -FREE DELETION [ℓ] needs $\Omega(n/p)$ bits of space.

Proof. We add onto the proof of [9, Theorem 6.3], by specifying how Alice and Bob provide the input to the p -pass streaming algorithm.

Let H be a minimal graph in Π which has at least two disjoint edges, say e_1 and e_2 . Let $H' := H \setminus \{e_1, e_2\}$. Create as an input for the streaming algorithm n copies of H' , where in copy i we add the edges e_1 and e_2 if and only if the input of DISJOINTNESS has a 1 for index i for Alice and Bob respectively.

As e_1 and e_2 are disjoint, e_2 is incident on two vertices v, w which are not incident to e_1 . For every pass the algorithm requires, we do the following. We provide all the copies of H as input to the streaming algorithm by letting Alice input all vertices $V(H) \setminus \{v, w\}$ as an AL stream. Note that Alice has enough information to do this, as the vertices incident on the edge e_2 in each copy of H is never included in this part of the stream. Then Alice passes the memory of the streaming algorithm to Bob, who inputs the edges incident to the vertices v, w for each copy of H (which includes e_2 if and only if the respective bit in the input of DISJOINTNESS is 1). This ends a pass of the stream.

Note that Alice and Bob have input the exact specification of a graph as described by Chitnis et al. [9, Theorem 6.3], but now as an AL stream. Hence, the correctness follows. \square

Theorem 26 provides a lower bound for, for example, EVEN CYCLE TRANSVERSAL [ℓ] (where Π is the set of all graphs that contain a C_4, C_6, \dots), and similarly ODD CYCLE TRANSVERSAL [ℓ] and FEEDBACK VERTEX SET [ℓ]. Theorem 26 does not hold for the scenario where Π contains only stars.

Notice that the lower bound proof makes a construction with a vertex cover size linear in n . Therefore, these bounds do not hold when the vertex cover size is bounded. We can prove lower bounds with constant vertex cover size for H -FREE DELETION with specific requirements on H , for the EA and VA models. These results follow by adapting the known lower bound construction by Bishnu et al. [13]. The idea is as follows. We take a copy of H , and extend some path of three vertices to the construction in Fig. 2, which we will call a *double fan*. The n vertices in the overlap of the two fans we will call the *center vertices*. The idea of this construction is that both Alice and Bob input the edges of one of the fans in the double fan, as those edges will be determined by the input to DISJOINTNESS. If there is some $1 \leq i \leq n$ such that the i -th bit in both Alice’s and Bobs input is 1, then the double fan will have a completed path from A to B , creating an induced copy of H in the graph. If we make sure the budget is $\ell = 0$, then the answer to H -FREE DELETION must be NO if and only if the answer to DISJOINTNESS is NO. If we manage this, Proposition 25 gives us a lower bound result.

It is important to note that this construction does not always make a correct reduction. Most importantly, it innately has trouble if H is some star. This is because the double fan construction can form exactly H even if Bobs input exists entirely of zeroes. Therefore, we must be careful to form conditions that exclude stars from the lower bounds proofs. Also, we cannot work on the AL model in this construction. That would mean that for each center vertex, we require some information from both Alice and Bob to construct the stream, which is something that neither Alice nor Bob can do. Avoiding this by making the center vertices edges increases the vertex cover size linearly.




Fig. 2. A reduction technique which we call a *double fan*. The right construction can imitate the input to the DISJOINTNESS instance, forming the left construction if and only if the answer to DISJOINTNESS is NO.

Theorem 27. *If H is a connected graph with at least 3 edges and a vertex of degree 2, then any algorithm for solving H -FREE DELETION [VC] on a graph G with $|VC(G)| \geq |VC(H)| + 1$ requires $\Omega(n/p)$ bits when using p passes in the VA/EA models, even when the solution size $\ell = 0$.*

Proof. Similar to the reductions given by Bishnu et al. [13], we give a reduction from DISJOINTNESS to H -FREE DELETION [VC] in the VA model when the solution size parameter $\ell = 0$. The idea is to build a graph G with bounded vertex cover size, and construct edges according to the input of DISJOINTNESS, such that G is H -free if and only if the output of DISJOINTNESS is YES.

Let \mathcal{A} be a streaming algorithm that solves H -FREE DELETION [VC] in the VA model, such that $|VC(G)| \leq |VC(H)| + 1$, and the space used is $o(n)$. Let G be a graph with $n + |V(H)| - 1$ vertices consisting of H where a degree-2 vertex in H is expanded to a double fan (i.e., the two adjacencies of this degree-2 vertex correspond to A and B , and the degree-2 vertex is replaced by the n center vertices of the double fan). Let x, y be the input strings, consisting of n bits each, of Alice and Bob for DISJOINTNESS, respectively.

Alice exposes to \mathcal{A} all the vertices of G , except for the vertex B . Here, Alice exposes an edge between A and the i -th center vertex of the double fan if and only if the i -th bit of x is 1. Notice that Alice can expose all these vertices according to the VA model, as only the addition of the vertex B will require information of the input of Bob, y . If Alice has exposed all vertices of G except for B , then she passes the memory of \mathcal{A} to Bob. Bob then exposes the vertex B , including an edge between B and the i -th center vertex of the double fan if and only if the i -th bit of y is 1. This completes one pass of the input for \mathcal{A} .

From the construction, observe that $|VC(G)| \leq |VC(H)| + 1$, as we may need to include both A and B in the vertex cover in G while it was optimal to include the degree-2 vertex in the vertex cover in H .

If the answer to DISJOINTNESS is NO, that is, there exists an index i such that $x_i = y_i = 1$, then in G the edges from A and B connect in the i -th center vertex, creating an induced copy of H in G , and so the graph is not H -free. Because $\ell = 0$, H -FREE DELETION [VC] must also be answered with NO.

If the answer to DISJOINTNESS is YES, that is, there is no index i such that $x_i = y_i = 1$, then there is no path from A to B through a center vertex in G . We will show that there is no induced occurrence of H in G . If the degree-2 vertex that we split into the double fan is contained in a cycle in H , then now this cycle is no longer present in the graph. As the rest of the graph simply consists of a (partial) copy of H , this means there cannot be enough cycles in the graph to get exactly H , and so H does not appear in G . Otherwise, the degree-2 vertex is not contained in a cycle, and so there is no path between A and B . Hence, for H to occur in G , it must occur in either the component containing A or the component containing B . If A or B was a degree-1 vertex in H , then the component containing it is now a star, and H is not a star by assumption. Otherwise, in a component, the number of vertices with degree 2 or higher is lower than the number of vertices in H with degree 2 or higher, as the double fan only adds degree-1 vertices in this case, and destroyed a degree-2 vertex, while A or B was already of degree 2 or higher. Hence, H cannot occur in either component of G . So the answer to H -FREE DELETION [VC] is YES.

Now, from Proposition 25 it follows that any algorithm for solving H -FREE DELETION [VC] on a graph G with $|VC(G)| \geq |VC(H)| + 1$ requires at least n/p bits when using p passes in the VA/EA models, even when $\ell = 0$. This can be generalized for every ℓ by adding ℓ disjoint copies of H to G , which also increases the vertex cover of G by a constant amount for each copy. \square

There are other conditions for which this reduction works as well.

Theorem 28. *If H is a graph with a vertex of degree at least 2 for which every neighbour has an equal or larger degree, then any algorithm for solving H -FREE DELETION [VC] on a graph G with $|VC(G)| \geq |V(H)|$ requires $\Omega(n/p)$ bits when using p passes in the VA/EA models, even when the solution size $\ell = 0$.*

Proof. Similar to the reductions given by Bishnu et al. [13], we give a reduction from DISJOINTNESS to H -FREE DELETION [VC] in the VA model when the solution size parameter $\ell = 0$. The idea is to build a graph G with bounded vertex cover size,

and construct edges according to the input of DISJOINTNESS, such that G is H -free if and only if the output of DISJOINTNESS is YES.

Let \mathcal{A} be a streaming algorithm that solves H -FREE DELETION [VC] in the VA model, such that $|VC(G)| \leq |V(H)|$, and the space used is $o(n)$. Let G be a graph with $n + |V(H)| - 1$ vertices consisting of H where a vertex degree at least 2 for which every neighbour has an equal or larger degree in H is expanded to a double fan, i.e., two adjacencies of this vertex correspond to A and B , and the vertex is replaced by the n center vertices of the double fan. All other adjacencies of this vertex are connected to all of the center vertices. Note that this is possible, as the vertex has degree at least 2. Let x, y be the input strings, consisting of n bits each, of Alice and Bob for DISJOINTNESS, respectively.

Alice exposes to \mathcal{A} all the vertices of G , except for the vertex B . Here, Alice exposes an edge between A and the i -th center vertex of the double fan if and only if the i -th bit of x is 1. Notice that Alice can expose all these vertices according to the VA model, as only the addition of the vertex B will require information of the input of Bob, y . If Alice has exposed all vertices of G except for B , then she passes the memory of \mathcal{A} to Bob. Bob then exposes the vertex B , including an edge between B and the i -th center vertex of the double fan if and only if the i -th bit of y is 1. This completes one pass to the input for \mathcal{A} .

From the construction, observe that $|VC(G)| \leq |V(H)|$, as the vertex cover of G can always be bounded by taking all vertices originally in H , which covers the edges towards the n center vertices.

If the answer to DISJOINTNESS is NO, that is, there exists an index i such that $x_i = y_i = 1$, then in G the edges from A and B connect in the i -th center vertex, creating an induced copy of H in G , and so the graph is not H -free. Because $\ell = 0$, H -FREE DELETION [VC] must also be answered with NO.

If the answer to DISJOINTNESS is YES, that is, there is no index i such that $x_i = y_i = 1$, then there is no path from A to B directly through a center vertex in G . We will show that there is no induced occurrence of H in G . Let us call the vertex that was expanded into the center vertices v , and say it has degree $d \geq 2$. Then all the neighbours of this vertex must also have degree at least d in H . However, the center vertices in G have degree at most $d - 1$, as no center vertex can be adjacent to both A and B . Hence, no center vertex can be used for v or any of its neighbours in an induced copy of H in G . Consider in H all vertices of degree at least d where the neighbours also have degree at least d , and say there are c many of these vertices. In any induced copy of H in G , these vertices must still have this relation of degrees. However, as none of the center vertices has degree at least d , G contains at most $c - 1$ such vertices, which means an induced copy of H cannot occur in G . Hence, the answer to H -FREE DELETION [VC] is YES.

Now, from Proposition 25 it follows that any algorithm for solving H -FREE DELETION [VC] on a graph G with $|VC(G)| \geq |V(H)|$ requires at least n/p bits when using p passes in the VA/EA models, even when $\ell = 0$. This can be generalized for every ℓ by adding ℓ disjoint copies of H to G , which also increases the vertex cover of G by a constant amount for each copy. \square

In Theorem 28 we only demand that the vertex cover size is at least $|V(H)|$, the number of vertices in H . One can wonder if this bound can be tightened, as in Theorem 27, where we only demand that the vertex cover size is at least $|VC(H)| + 1$. The problem in Theorem 28, is that we might split a vertex of high degree. To get a valid vertex cover without it having linear size in n , the only option is to include at least all adjacencies of the center vertices. This makes it that the vertex cover can get a size up to $|V(H)|$, and so this is the only safe demand.

For clarity, we summarize these lower bounds for bounded vertex cover size in a corollary.

Corollary 29. *If H is such that either:*

1. H is a connected graph with at least 3 edges and a vertex of degree 2, or,
2. H is a graph with a vertex of degree at least 2 for which every neighbour has an equal or larger degree,

then any algorithm for solving H -FREE DELETION [VC] on a graph G with $|VC(G)| \geq |V(H)|$ requires $\Omega(n/p)$ bits when using p passes in the VA/EA models, even when the solution size $\ell = 0$.

Corollary 29 gives a lower bound for COGRAPH DELETION [VC], as this is H -FREE DELETION [VC] with $H = \{P_4\}$. Examples for which Corollary 29 does not give a lower bound include CLUSTER VERTEX DELETION [VC] (indeed, then a kernel is known [13]), or more generally, H -FREE DELETION [VC] when H is a star.

We observe that these constructions can be straightforwardly modified to apply to the problems of ODD CYCLE TRANSVERSAL [VC], EVEN CYCLE TRANSVERSAL [VC], FEEDBACK VERTEX SET [VC], so these problems admit a similar lower bound to Corollary 29.

5. Conclusion

We have seen different streaming algorithms and lower bounds for Π -FREE DELETION and its more specific forms, making use of the minimum vertex cover as a parameter. We have seen the potency of the AL streaming model in combination with the vertex cover, where in other streaming models lower bounds arise. It is interesting that for very local structures like a P_3 ,

this combination works effortlessly, giving a very efficient memory-optimal algorithm. For more general structures troubles arise, but nonetheless, we can solve the more general problems with a many-pass, low-memory approach. Alternatively, the adaptations of kernels give rise to a few-pass, high-memory algorithm, which provides a possible trade-off when choosing an algorithm.

We also propose the following open problems. Can lower bounds be found expressing a pass/memory trade-off in the vertex cover size for the Π -FREE DELETION [VC] problem? Or alternatively, can we find an upper bound for Π -FREE DELETION [VC] using $\mathcal{O}(K \log n)$ bits of memory but only a polynomial in K number of passes? Essentially, here we ask whether or not our algorithm is reasonably tight, or can be improved to only use a polynomial number of passes in K . A lower bound expressing a trade-off in terms of the vertex cover size is a standalone interesting question, as most lower bound statements about streaming algorithms express a trade-off in terms of n .

We also ask about the unparameterized streaming complexity of CLUSTER VERTEX DELETION in the AL model. While lower bounds for most other Π -FREE DELETION problems in the AL model follow from our work (Theorem 26) and earlier work of Bishnu et al. [13], this appears an intriguing open case.

Finally, we ask if there is a $2^{\mathcal{O}(K \log K)}$ lower bound for Π -FREE DELETION [VC] when Π is characterized by few adjacencies?

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Jelle Oostveen reports financial support was provided by NWO grant OCENW.KLEIN.114 (PACAN).

Data availability

No data was used for the research described in the article.

Acknowledgements

The authors thank the anonymous reviewer of this paper for their helpful suggestions.

References

- [1] M.R. Henzinger, P. Raghavan, S. Rajagopalan, Computing on data streams, in: J.M. Abello, J.S. Vitter (Eds.), Proc. DIMACS 1998, in: DIMACS, vol. 50, DIMACS/AMS, 1998, pp. 107–118.
- [2] A. McGregor, Graph stream algorithms: a survey, SIGMOD Rec. 43 (1) (2014) 9–20.
- [3] S. Muthukrishnan, Data streams: algorithms and applications, Found. Trends Theor. Comput. Sci. 1 (2) (2005).
- [4] J. Feigenbaum, S. Kannan, A. McGregor, S. Suri, J. Zhang, On graph problems in a semi-streaming model, Theor. Comput. Sci. 348 (2–3) (2005) 207–216.
- [5] J. Feigenbaum, S. Kannan, A. McGregor, S. Suri, J. Zhang, Graph distances in the streaming model: the value of space, in: Proc. SODA 2005, SIAM, 2005, pp. 745–754.
- [6] S. Fafanie, S. Kratsch, Streaming kernelization, in: Proc. MFCS 2014, in: LNCS, vol. 8635, Springer, 2014, pp. 275–286.
- [7] R.H. Chitnis, G. Cormode, M.T. Hajiaghayi, M. Monemizadeh, Parameterized streaming: maximal matching and vertex cover, in: Proc. SODA 2015, SIAM, 2015, pp. 1234–1251.
- [8] R.G. Downey, M.R. Fellows, Parameterized Complexity, Monographs in Computer Science, Springer, 1999.
- [9] R. Chitnis, G. Cormode, H. Esfandiari, M.T. Hajiaghayi, A. McGregor, M. Monemizadeh, S. Vorotnikova, Kernelization via sampling with applications to finding matchings and related problems in dynamic graph streams, in: Proc SODA 2016, SIAM, 2016, pp. 1326–1344.
- [10] A. Bishnu, A. Ghosh, S. Kolay, G. Mishra, S. Saurabh, Fixed parameter tractability of graph deletion problems over data streams, in: Proc. COCOON 2020, in: LNCS, vol. 12273, Springer, 2020, pp. 652–663.
- [11] A. Agrawal, A. Biswas, É. Bonnet, N. Brettell, R. Curticapean, D. Marx, T. Miltzow, V. Raman, S. Saurabh, Parameterized streaming algorithms for min-ones d -sat, in: Proc. FSTTCS 2019, in: LIPIcs, vol. 150, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019, 8.
- [12] R. Chitnis, G. Cormode, Towards a theory of parameterized streaming algorithms, in: Proc. IPEC 2019, in: LIPIcs, vol. 148, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019, 7.
- [13] A. Bishnu, A. Ghosh, S. Kolay, G. Mishra, S. Saurabh, Fixed-parameter tractability of graph deletion problems over data streams, CoRR, arXiv:1906.05458 [abs], 2019, arXiv:1906.05458.
- [14] G. Cormode, J. Dark, C. Konrad, Independent sets in vertex-arrival streams, in: Proc. ICALP 2019, in: LIPIcs, vol. 132, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019, 45.
- [15] A. McGregor, S. Vorotnikova, H.T. Vu, Better algorithms for counting triangles in data streams, in: Proc. PODS 2016, ACM, 2016, pp. 401–411.
- [16] B.M. Jansen, The power of data reduction: Kernels for fundamental graph problems, Ph.D. thesis, Utrecht University, 2013.
- [17] J.M. Lewis, M. Yannakakis, The node-deletion problem for hereditary properties is NP-complete, J. Comput. Syst. Sci. 20 (2) (1980) 219–230.
- [18] L. Cai, Fixed-parameter tractability of graph modification problems for hereditary properties, Inf. Process. Lett. 58 (4) (1996) 171–176.
- [19] F.V. Fomin, P.A. Golovach, D.M. Thilikos, On the parameterized complexity of graph modification to first-order logic properties, Theory Comput. Syst. 64 (2) (2020) 251–271.
- [20] I. Sau, U. dos Santos Souza, Hitting forbidden induced subgraphs on bounded treewidth graphs, Inf. Comput. 281 (2021) 104812, <https://doi.org/10.1016/j.ic.2021.104812>.
- [21] H.L. Bodlaender, B.M.P. Jansen, S. Kratsch, Kernelization lower bounds by cross-composition, SIAM J. Discrete Math. 28 (1) (2014) 277–305.
- [22] F.V. Fomin, B.M.P. Jansen, M. Pilipczuk, Preprocessing subgraph and minor problems: when does a small vertex cover help?, J. Comput. Syst. Sci. 80 (2) (2014) 468–495.
- [23] B.M.P. Jansen, J.J.H. de Kroon, Preprocessing vertex-deletion problems: characterizing graph properties by low-rank adjacencies, J. Comput. Syst. Sci. 126 (2022) 59–79, <https://doi.org/10.1016/j.jcss.2021.12.003>.
- [24] J. Nešetřil, S. Poljak, On the complexity of the subgraph problem, Comment. Math. Univ. Carol. 26 (2) (1985) 415–419.

- [25] M. Dalirrooyfard, T.D. Vuong, V.V. Williams, Graph pattern detection: hardness for all induced patterns and faster noninduced cycles, *SIAM J. Comput.* 50 (5) (2021) 1627–1662, <https://doi.org/10.1137/20M1335054>.
- [26] A. Itai, M. Rodeh, Finding a minimum circuit in a graph, *SIAM J. Comput.* 7 (4) (1978) 413–423, <https://doi.org/10.1137/0207033>.
- [27] F. Eisenbrand, F. Grandoni, On the complexity of fixed parameter clique and dominating set, *Theor. Comput. Sci.* 326 (1–3) (2004) 57–67, <https://doi.org/10.1016/j.tcs.2004.05.009>.
- [28] M. Konagaya, Y. Otachi, R. Uehara, Polynomial-time algorithms for subgraph isomorphism in small graph classes of perfect graphs, *Discrete Appl. Math.* 199 (2016) 37–45.
- [29] S. Kijima, Y. Otachi, T. Saitoh, T. Uno, Subgraph isomorphism in graph classes, *Discrete Math.* 312 (21) (2012) 3164–3173.
- [30] D.G. Corneil, Y. Perl, L.K. Stewart, A linear recognition algorithm for cographs, *SIAM J. Comput.* 14 (4) (1985) 926–934, <https://doi.org/10.1137/0214065>.
- [31] S. Olariu, A simple linear-time algorithm for computing the center of an interval graph, *Int. J. Comput. Math.* 34 (3–4) (1990) 121–128, <https://doi.org/10.1080/00207169008803870>.
- [32] T. Kloks, D. Kratsch, H. Müller, Finding and counting small induced subgraphs efficiently, *Inf. Process. Lett.* 74 (3–4) (2000) 115–121, [https://doi.org/10.1016/S0020-0190\(00\)00047-8](https://doi.org/10.1016/S0020-0190(00)00047-8).
- [33] M. Kowaluk, A. Lingas, E. Lundell, Counting and detecting small subgraphs via equations, *SIAM J. Discrete Math.* 27 (2) (2013) 892–909, <https://doi.org/10.1137/110859798>.
- [34] H. Moser, D.M. Thilikos, Parameterized complexity of finding regular induced subgraphs, *J. Discret. Algorithms* 7 (2) (2009) 181–190.
- [35] T. Kloks, D. Kratsch, H. Müller, Finding and counting small induced subgraphs efficiently, *Inf. Process. Lett.* 74 (3–4) (2000) 115–121.
- [36] V.V. Williams, J.R. Wang, R.R. Williams, H. Yu, Finding four-node subgraphs in triangle time, in: P. Indyk (Ed.), *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4–6, 2015*, SIAM, 2015, pp. 1671–1680.
- [37] M. Bläser, B. Komarath, K. Sreenivasaiiah, Graph pattern polynomials, in: S. Ganguly, P.K. Pandya (Eds.), *38th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS2018, December 11–13, 2018, Ahmedabad, India*, in: *LIPICs*, vol. 122, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018, 18.
- [38] D. Eppstein, Subgraph isomorphism in planar graphs and related problems, in: *Graph Algorithms and Applications I*, World Scientific, 2002, pp. 283–309.
- [39] R. Chitnis, G. Cormode, Towards a theory of parameterized streaming algorithms, *CoRR*, arXiv:1911.09650 [abs], 2019, arXiv:1911.09650, <http://arxiv.org/abs/1911.09650>.
- [40] D. Tsur, Faster parameterized algorithm for cluster vertex deletion, *Theory Comput. Syst.* 65 (2) (2021) 323–343, <https://doi.org/10.1007/s00224-020-10005-w>.
- [41] F.N. Abu-Khzam, Maximum common induced subgraph parameterized by vertex cover, *Inf. Process. Lett.* 114 (3) (2014) 99–103.
- [42] R. Impagliazzo, R. Paturi, On the complexity of k -sat, *J. Comput. Syst. Sci.* 62 (2) (2001) 367–375.
- [43] F.N. Abu-Khzam, É. Bonnet, F. Sikora, On the complexity of various parameterizations of common induced subgraph isomorphism, *Theor. Comput. Sci.* 697 (2017) 69–78.
- [44] D. Lokshtanov, D. Marx, S. Saurabh, Slightly superexponential parameterized problems, *SIAM J. Comput.* 47 (3) (2018) 675–702.
- [45] M. Cygan, F.V. Fomin, L. Kowalik, D. Lokshtanov, D. Marx, M. Pilipczuk, M. Pilipczuk, S. Saurabh, *Parameterized Algorithms*, Springer, 2015.
- [46] B.A. Reed, K. Smith, A. Vetta, Finding odd cycle transversals, *Oper. Res. Lett.* 32 (4) (2004) 299–301.
- [47] R.H. Chitnis, G. Cormode, H. Esfandiari, M. Hajiaghayi, M. Monemizadeh, Brief announcement: new streaming algorithms for parameterized maximal matching & beyond, in: *Proc. SPAA 2015, ACM*, 2015, pp. 56–58.
- [48] A. Bishnu, A. Ghosh, G. Mishra, S. Sen, On the streaming complexity of fundamental geometric problems, *CoRR*, arXiv:1803.06875 [abs], 2018, arXiv:1803.06875.
- [49] D. Agarwal, A. McGregor, J.M. Phillips, S. Venkatasubramanian, Z. Zhu, Spatial scan statistics: approximations and performance study, in: *Proc. SIGKDD 2006, ACM*, 2006, pp. 24–33.