

A size/fan-in trade-off for circuits

Pierre Meyer¹

¹CISPA Helmholtz Centre for Information Security, GERMANY. pierre.meyer@cispa.de

We consider (boolean or arithmetic) circuits in which every gate may compute an arbitrary function of its input gates. We show a novel tradeoff between a circuit’s size and its fan-in:

Any function which may be computed using s fan-in 2 gates can alternatively be computed using $s/3 (1 + O(1/\sqrt{k}))$ fan-in 2^k gates.

This asymptotically improves, at every fan-in 2^k , the previous bound of $2s/5 (1 + O(1/k))$ by Charbit, Couteau, Meyer, and Naserasr [TCC’24]. Among other applications, this improves the communication complexity of secure multiparty computation in the correlated randomness model. *Any* n -input m -output circuit with s internal gates (over arbitrary binary gates) can be securely computed in the correlated randomness model with per party communication $s/3 + n + m$ and computation $\tilde{O}(s)$.

Our paper stands at the intersection of cryptography, complexity theory, and graph theory, but our main technical contribution is one to extremal combinatorics: we establish that every order- n 2-degenerate graph admits a planarising set of size at most $\lfloor n/3 \rfloor$.

Our main conceptual contribution is to relate the existence of sublinear-size (directed) k -path transversals to well-studied graph parameters. Along the way, we uncover a recurrently overstated lemma throughout the literature on sublinear-size vertex-separators, graph divisions, and hyperfinite graphs. According to this lemma, any monotone graph class admitting sublinear-size balanced vertex separators should be weakly hyperfinite. However, this is contradicted by a graph class put forward by [Moshkovitz and Shapira, Random Structures & Algorithms’18]. Unfortunately, the lemma appears in highly influential works such as [Henzinger, Klein, Rao, and Subramanian, STOC’94 & JCSS’97], or the textbook of Nešetřil and Ossona de Mendez [Sparsity, Springer’12], and in turn it is used in a significant number of papers. Thankfully a slightly weaker version of this lemma is true, with a caveat on how sublinear the vertex separators need to be. We provide the correction as a service to the community, but emphasise that we do little more than point to existing results.

Contents

1	Introduction	2
1.1	Our results.	3
1.2	Related works.	4
2	Technical Overview	5
2.1	A mostly graph-theoretic template.	5
2.2	Characterising the easy classes for the k -path hitting set problem.	6
2.3	A planarising set for all 2-degenerate graphs.	8
3	Characterising the easy classes for the k-path hitting set problem.	8
3.1	Linking k -path transversals to balanced vertex separators	8
3.2	Linking k -path transversals in the “for all k ” regime to weak hyperfiniteness	14
4	A planarising set for all 2-degenerate graphs	16
5	A trade-off between the size and the fan-in of circuits	19

1 Introduction

Secure Multiparty Computation (MPC) [Yao86, GMW87] enables mutually distrusting parties to jointly compute a function of their private inputs, while revealing nothing beyond the function’s output. For our purposes, we may consider the *feasibility* of MPC against computationally unbounded parties (this is known as the setting of *information-theoretic security*) corrupting a subset of the parties to be an essentially closed question given enough communication between the parties [BGW88, CCD88, RB89]. Specifically, the total communication of all known protocols scales linearly with the computational complexity of the desired function (understood as the size of the smallest binary circuit computing the function). This stands in stark contrast with “insecure” computation, where the parties need only exchange their inputs, and the computation can thence be done locally, no matter how complex. One of the most fundamental questions in MPC is to understand whether this linear dependency is inherent. Settling this question one way or another seems out of reach of current techniques.

Meaningful progress is possible in the *correlated randomness model* [IKM⁺13], where the parties have access to secret random strings drawn from some joint distribution. This model is both of theoretical interest, as it bypasses the impossibility of dishonest-majority information-theoretic MPC [BGW88, CK89, Kus89], and of practical interest, as the correlated randomness can be distributed far in advance during an input-independent pre-processing phase. Ishai, Kushilevitz, Meldgaard, Orlandi, and Paskin-Cherniavsky [IKM⁺13] showed how to securely compute *any* function with constant communication, but their protocol uses an amount of computation (and correlated randomness) proportional to the size of the function’s look-up table. Couteau [Cou19] later noted this could be refined into a fine-grained doubly exponential trade-off between communication and computation for functions with a “nice structure”: any size- s *layered*¹ circuit can be securely computed using communication $O(s/k)$ and computation $O(s2^{2^k})$. The question remains whether such a trade-off is possible for *all functions*.

Charbit, Couteau, Meyer, and Naserasr [CCMN24] proposed to address this problem from a complexity-theoretic point of view. From here on out, the reader need know nothing about cryptographic protocols beyond the following statement: if a function is computable by a circuit of fan-in 2^ℓ with s_ℓ gates, then it can be securely computed with communication $O(s_\ell)$ and computation $O(s_\ell 2^{2^\ell})$ [IKM⁺13, Cou19, CCMN24]. The question now is whether a trade-off is possible between a circuit’s size and fan-in.

If a function can be computed by a size s binary circuit, what can be said about the size s_ℓ of the smallest fan-in 2^ℓ circuit which computes it? In particular, can s_ℓ be significantly smaller than s for a given value of ℓ ?

For the application to cryptography, the main regime of interest is $\ell = O(\log \log s)$, so computation remains polynomial. Arguably, setting ℓ to be an “arbitrarily small function in $\omega(1)$ ” is also interesting, so computation is quasilinear.

Charbit, Couteau, Meyer, and Naserasr [CCMN24] proposed a template for achieving such a trade-off. The idea is that if there is a way to select (or *pebble*) at most α gates of a fan-in 2 circuit C according to the two following rules, then this same circuit can alternatively be computed by an α -gate fan-in 2^k circuit C' : (1) every input gate and every output gate must be selected, and (2) there can be no chain of k consecutive unselected gates. The idea is that each gate of C' is going to compute the same subcircuit as one of the selected gates of C . Because C has fan-in at most 2 and because the longest chain of unselected gates has length less than k , composing the unselected gates of C in such a way that each selected gate is a function of values computed by selected gates only yields a fan-in 2^k gate. It is worth noting that we must allow C_k ’s gates to be arbitrary 2^k -ary functions, but that the template is agnostic to the computation’s underlying ring.

Choosing which gates to select boils down to solving a graph-theoretic problem. Consider the binary circuit’s underlying directed acyclic graph (DAG), then remove all the vertices corresponding to the circuit’s inputs and outputs (since those must always be selected), and then look for a *directed k -path transversal* (a.k.a. a *directed k -path hitting set*, i.e. a set of vertices whose removal yields a graph which does not contain a directed k -path as a subgraph). In all generality the DAG obtained from a general binary circuit could be any DAG with fan-in at most 2, which leads us to consider the following concrete technical question:

¹A circuit is layered if its gates can be partitioned into ordered layers such that wires only connect each layer to the next. Note that any function computed by a size- s depth- D circuit can be computed by size- $O(sD)$ layered circuit.

Can we improve the best known bound (as a function of k and n) on the directed k -path transversal number of order- n fan-in 2 DAGs?

The asymptotically best bound before this work was $2n/5 + O(n/k)$, obtained by [BDBS14, CCMN24].

1.1 Our results.

We successfully improve the bound from $2n/5 + O(n/k)$ to $n/3 + O(n/\sqrt{k})$, but the steps we took to achieve this result are arguably more interesting than this bottomline.

Firstly, we provide a near-complete characterisation of which graph classes admit sublinear k -path transversals, by relating this number to well-studied graph parameters such as treewidth, balanced vertex separators, graph divisions, fragmentability, and weak hyperfiniteness. The main message of Section 3 is this near-tight bidirectional link:

- If a hereditary graph class admits balanced vertex separators of size $O(n/[(\log n)(\log \log n)^{1+\varepsilon}])$, then it admits sublinear k -path hitting sets for any $k = \omega(1)$
- Conversely, any size- α k -path hitting set yields a balanced vertex separator of size $\alpha + k - 1$.

First of all this paves the way for further constant-factor improvements on our asymptotic $1/3$. Improving our result is now possible by completely disregarding the k -path hitting set problem and address a self-contained question in extremal combinatorics:

Does every order- n 2-degenerate graph admit an induced subgraph satisfying of (1) size at least εn , where $\varepsilon > 2/3$, and (2) treewidth $O(n/[(\log n)(\log \log n)^{1+\varepsilon}])$?

In fact, our own improvement from $2/5$ to $1/3$ was obtained by showing that every order- n 2-degenerate graph admits an induced planar subgraph of size at least $\lceil 2n/3 \rceil$, which may be of independent interest.

Additionally, our characterisation gives a very concrete avenue towards *actually breaking the circuit-size barrier*, beyond trying to optimise our constant factor. This would work by providing a functionality-preserving transform to get any circuit into one with sublinear treewidth. Admittedly, this remains a longshot as treewidth is an inherently global parameter, whereas known circuit transformations tend to be local.

Substituting [CCMN24]’s trade-off for our own in [CCMN24, Corollary 1] yields the following application to MPC.

Corollary 1 (MPC protocols through the depth-reduction lens, updated).

1. **Information-theoretic MPC in the correlated randomness model.** *Let C be an n -input, m -output, $\mathbb{F}$ -arithmetic circuit with s non-output computation gates (over any basis of binary gates). There exists an N -party protocol for perfectly securely computing C in the correlated randomness model in the presence of a semi-honest adversary corrupting up to $N - 1$ parties. The protocol uses the following resources (in bits):*

- Total communication (in bits):

$$\left(\left(\frac{1}{3} + o(1) \right) s + n + m \right) \cdot N \cdot \log |\mathbb{F}|$$

- Correlated randomness per party: $\tilde{O}(s + n) \cdot \log |\mathbb{F}|$ bits of function-dependent correlated randomness per party (or alternatively $B^{O(1)} \cdot \log |\mathbb{F}|$ bits of function-independent correlated randomness, where B is some a priori bound on $s + n + m$),

- Local computation per party (in bits): $\tilde{O}(s \cdot N \cdot \log^2 |\mathbb{F}|)$.

2. **The 1-out-of- M -OT complexity of 2PC.** Let $f: \{0,1\}^n \rightarrow \{0,1\}$ be a function which can be computed by a boolean circuit with s computation gates (over any basis of binary gates). For every $M \geq 2$, there exists a two-party protocol for computing C with perfect passive security in the 1-out-of- M -OT hybrid model; this protocol makes $s/3 + O(s/\sqrt{\log \log M})$ calls to the oblivious transfer functionality.

1.2 Related works.

Depth reduction, under its many names. The combinatorial problem at the heart of this work has been investigated in several communities and under different names. In its undirected form (removing as few vertices as possible so that no path on k vertices remains) it is the *k -path vertex cover* problem of Brešar, Kardoš, Katrenič, and Semanišin [BKKS11], and is also called the *P_k -hitting set*, *P_k -vertex deletion*, or *k -path transversal* problem. Some special cases have received particular attention: $k = 2$ is the *vertex cover* problem, and $k = 3$ is the complement of the *dissociation number*. We refer to the survey of Tu [Tu22] for a partial overview, and to the work of Bujtás, Jakovac, and Tuza [BJT22] for general upper bounds in terms of degrees and sparsity (but note that none of these results is directly relevant for the graph classes arising from our applications).

In its directed form it is called the *DAG vertex deletion* problem $DVD(k)$ introduced by Paik, Reddy, and Sahni [PRS94] under the name “deleting vertices to bound path length”. Their motivation was already a circuit one, albeit from *Very Large Scale Integration* (VLSI) design rather than cryptography: one wishes to *upgrade* as few gates of a circuit as possible so that no electrical signal traverses k consecutive un-upgraded gates. The problem was recently renamed the *DAG depth reduction* problem [CCMN24] in the cryptography literature. Note that since a directed k -path of a DAG is also a k -path of its underlying undirected graph, every undirected upper bound transfers to the directed setting.

A very close variant of the problem is studied in the literature on fully homomorphic encryption under the name of the *bootstrap(ping) problem*. In the lattice-based blueprint originating with Gentry [Gen09], each ciphertext carries a noise term which grows slightly at linear gates and much more significantly at non-linear ones. Once noise exceeds a threshold, correctness is lost unless the ciphertext is refreshed by a (computationally heavy) step known as *bootstrapping*. The *bootstrap(ping) problem*, formalised by Benhamouda, Lepoint, Mathieu, and Zhou [BLMZ17], asks for the minimum number of bootstraps needed to evaluate a given circuit. In graph-theoretic terms one is given a DAG whose non-source vertices are coloured according to whether they carry an addition or a multiplication, and one seeks a small set of vertices meeting every path that contains more than L multiplication vertices. This can be seen as a weighted version of our problem, capable of capturing *multiplicative-depth reduction* rather than simple depth reduction.

Optimisation versions. The question we ask is *extremal*: we bound the transversal number of the worst graph in a class (and it so happens our proofs happen to be constructive). The bulk of the literature instead studies the *optimisation* problem of finding a minimum transversal in a given graph. For every $k \geq 2$ the directed version is NP-complete on general DAGs [PV16], and Svensson [Sve12] showed that for every constant k it is UGC-hard to approximate $DVD(k)$ within a ratio better than k . Since a k -approximation is easy, the approximability of $DVD(k)$ is thereby completely settled for constant k . We note Svensson’s lower bounds do not apply if k is allowed to grow with n , which is the setting we investigate in this paper.

The edge-deletion analogue has seen more recent activity. Kenkre, Pandit, Purohit, and Saket [KPPS17] introduced the *DAG edge-deletion* problem $DED(k)$ (which asks to delete a minimum-weight set of edges so that no path of length k remains). They gave a k -approximation and studied its maximisation variant *Max k -Ordering*. Klein and Wexler [KW15] improved the ratio to $\frac{2}{3}(k+1)$, and Kalantarzadeh, Klein, and Reis [KKR25] brought it to $0.549(k+1)$ by randomised rounding, showing along the way that no independent rounding distribution can go below $0.542(k+1)$.

Separators, divisions, and hyperfiniteness. Our characterisation of the graph classes admitting sub-linear transversals draws on three strands that have so far been developed separately. The first is the theory

of *sublinear balanced vertex separators*, beginning with the planar separator theorem of Lipton and Tarjan [LT79] and including the equivalence, due to Dvořák and Norin [DN19], between admitting sublinear separators hereditarily and having small treewidth. The second is the theory of *graph divisions*: Frederickson [Fed87] introduced r -divisions to speed up shortest-path computation in planar graphs, and Klein, Rao, Rauch/Henzinger, and Subramanian [KRRS94, HKRS97] extended the construction to any class with sublinear separators. It is the latter extension whose statement we correct in Section 3.1.1. The third is the theory of *hyperfiniteness* and *fragmentability*, developed independently in property testing and in the structural theory of sparse classes: Czumaj, Shapira, and Sohler [CSS09] showed that hereditary classes with sufficiently sublinear edge separators are hyperfinite, Moshkovitz and Shapira [MS18] showed that this fails just beyond that threshold, and the notion appears under the name *fragmentability* in [EM94, EF01] and as *weak hyperfiniteness* in [NOdM10, NOdM12]. Our contribution here is limited to transcribing results from one terminology to another: we observe that the counterexample of [MS18] refutes the division lemma of [HKRS97, NOdM12] as stated, and that the positive results of [CSS09] deliver a near-tight repair.

Planarisation and degenerate graphs. Our extremal bound (Theorem 19) states that every 2-degenerate graph has a planarising set of size $\lfloor n/3 \rfloor$, equivalently a planar induced subgraph on $\lceil 2n/3 \rceil$ vertices. Its direct ancestor is the theorem of Borowiecki, Drgas-Burchardt, and Sidorowicz [BDBS14] that every 2-degenerate graph has a feedback vertex set of size at most $2n/5$, a bound they show to be tight; our argument is a refinement of theirs, exploiting that planarity is a much weaker requirement than acyclicity. Bounds of this shape — deleting a constant fraction of the vertices of a sparse graph to reach a prescribed minor-closed class — also appear in the parameterised-complexity literature on vertex deletion problems, and the fragmenting lemma we borrow from Gupta, Lee, Li, Manurangsi, and Włodarczyk [GLL⁺19] is of this kind.

More applications to cryptography. Although we decided to illustrate the applicability of our size/fan-in trade-off in the correlated randomness model, there are many other applications in computationally secure MPC, as noted in [CCMN24]. Broadly speaking, our result improves every result on sublinear-communication MPC which does not use FHE. For the most part, this means protocols based on *homomorphic secret sharing* (HSS) from DDH [BGI16], from DCR [OSY21, RS21], and from class groups of imaginary quadratic fields [ADOS22], from lattices [BKS19], and code-based assumptions [CM21, COS⁺22, BCM23, DIJL23]. This also includes the correlated-SPIR route of [BCM22]. Additionally, our results means the garbling scheme of Meyer, Orlandi, Roy, and Scholl [MORS25] achieves a factor-3 compression over the circuit’s size “in the clear” for *all circuits* (without having to assume a layered structure).

2 Technical Overview

How does one (slightly) decrease a circuit’s size by (slightly) increasing its fan-in?

2.1 A mostly graph-theoretic template.

Charbit, Couteau, Meyer, and Naserasr [CCMN24] proposed a three-step template for solving this problem.

Circuit preprocessing. The first step is to *preprocess the circuit* using functionality-preserving (and size-preserving) transformations. This changes the circuit’s underlying DAG without changing the function being computed. For instance, we may assume without loss of generality that circuits are triangle-free². However, both [CCMN24] and this paper skip the preprocessing step and simply consider the circuit’s underlying fan-in-2 DAG, after removing the circuit’s inputs and outputs, can be arbitrary.

²Indeed, if a circuit contains a triangle, it means that one of the gates is computed as $f(a, b)$ where b is itself computed as a gate of the form $g(a, c)$. Up to rewiring the circuit and changing the nature of a gate, but without altering the total gate count, $f(a, b)$ can instead be computed as $h(a, c) := f(a, g(a, c))$. By appropriately sweeping through the graph in reverse topological order, we can eliminate all triangles in the graph (the ordering is important to ensure that any triangle introduced by our rewiring rule must later be caught by the sweep).

Modularisation. The second step is to choose an intermediary graph class $\mathcal{E}$, and bound the size of a *modulator* to class $\mathcal{E}$. If $G = (V, E)$ is the graph obtained from the circuit pre-processing step, a modulator to $\mathcal{E}$ is a vertex set $X \subseteq V$ such that $G \setminus X \in \mathcal{E}$. In [CCMN24], $\mathcal{E}$ is the class of all directed forests.

Depth-reduction. The third step is to solve the directed k -path hitting set problem in $\mathcal{E}$. That is, we are looking to bound the size of $Y \subseteq V$ such that $(G \setminus X) \setminus Y$ has no directed k -path subgraph. For this reason, $\mathcal{E}$ should be chosen to be provably easy, which typically means we would want a guarantee that $|Y| = o(|V|)$.

After all these steps are completed, $X \cup Y$ constitutes a directed k -path hitting set for G . It should be noted that the template helps guide protocol design, but is not inherently restrictive (as we can always just choose $\mathcal{E}$ to be the class of all fan-in 2 DAGs). A simplified version of the template ignores all orientation and asks to hit every undirected k -path. This is a stronger requirement than the directed one: every directed path is also a path in the underlying undirected graph. Consequently, every solution to the undirected problem automatically yields a solution to the directed problem.

2.2 Characterising the easy classes for the k -path hitting set problem.

[CCMN24] uses *forests* as its modularisation target. Forests (whether directed or not) admit a very straightforward algorithm for finding a (directed) k -path hitting set: find the deepest³ sub-tree T containing a (directed) k -path, then remove its root u (or rather, the root should be added to the hitting set), and recurse. Observe that by removing u :

- $T \setminus u$ is *separated* from the rest of the graph⁴;
- $T \setminus u$ cannot contain a (directed) k -path, as T was chosen to be the deepest such sub-tree;
- $T \setminus u$ must contain at least $k - 1$ vertices; this is because T contained a (directed) k -path, and therefore at least k vertices.

For these reasons, whenever we select a vertex to be part of the (directed) k -path hitting set, we can essentially remove $k - 1$ others for free. Every order- n (directed) forest therefore contains a (directed) k -path hitting set of size at most $\lfloor n/k \rfloor$.

Bounded treewidth. When seeking to generalise this approach, it seems only natural to ask whether graph-theorists have studied the notion of “tree-like” graphs. For the reader well-versed in graph parameters, the notion of (bounded) *treewidth* [RS86] should immediately come to mind. For the others, let us start with the simpler notion of a *tree-partition*.

A graph $G = (V, E)$ admits a *tree-partition of width at most w* if V can be partitioned into sets (usually called *bags*) of size at most $w + 1$ such that contracting the nodes in each bag yields a forest H . Reformulating to avoid the *contraction* terminology, create a graph with one vertex for each bag, then connect two vertices if G had an edge connecting vertices in their corresponding bags. Observe that if we take a $\frac{k}{w+1}$ -path hitting set in the forest H (and there exists such a set of size $\lfloor n/(k/(w+1)) \rfloor$ [CCMN24, Lemma 7]) and remove all the vertices in the corresponding bags, the longest path can only intersect $\frac{k}{w+1}$ bags (each of size at most $w + 1$), and therefore through at most k vertices.

The notion of *treewidth at most w* generalises this by allowing the bags to overlap, forming a cover rather than a partition of the vertices. While it should not seem obvious to the reader apprehending this notion for the first time, the same algorithm works⁵ and shows that order- n graphs of treewidth at most w admit a k -path hitting set of size $O(\frac{w}{k}n)$. This can be sublinear, but w needs to be small (especially in the $k = O(\log \log n)$ regime, which is of particular interest in our cryptographic applications). The issue with this algorithm is that we remove only k vertices at a time but pay for the full removal cost of bag.

There is no universally correct way in the hierarchy of graph parameters to generalise algorithms working for bounded-treewidth graphs. In this case, the way forward is to consider graphs with small *separators*, *i.e.* sets of nodes whose removal splits the graph into at least two connected components, of bounded size.

³In the undirected case, first pick an arbitrary root for each tree so the notion of depth makes sense.

⁴Meaning there are no edges between it and the rest of the graph.

⁵The key is to use a *nice tree decomposition*, and to charge the cost of deletions to forget bags.

Sublinear balanced vertex separators. A *balanced vertex separator* is a set of vertices whose removal splits the graph into two roughly same-size components (usually formalised as each side needing to have at most two-thirds of the vertices), with no edge crossing from one to the other. This notion is useful for us because any remaining k -path must then live on one side of the cut, reducing the problem to finding a hitting set on each side independently.

However, the recursion is more subtle than it may seem at first. If $\mathcal{G}$ is a hereditary⁶, and every order- n graph in $\mathcal{G}$ admits a size- $\alpha(n)$ balanced vertex separator, then, if $\tau_k^{\mathcal{G}}(n)$ denotes the size of the smallest k -path hitting set in the worst order- n graph in $\mathcal{G}$:

$$\tau_k^{\mathcal{G}}(n) \leq \begin{cases} 2 \tau_k^{\mathcal{G}}(2n/3) + \alpha(n) & \text{if } n \geq k \\ 0 & \text{if } n < k \end{cases}$$

But this expression will not yield a sublinear upper bound on $\tau_k^{\mathcal{G}}(n)$ for any $\alpha > 0$. Fortunately, the problem of how to remove vertex separators recursively until each separated subgraph has at most r vertices (in our case, $r = k - 1$ guarantees k -path freeness) has been studied by Frederikson [Fed87, Lemma 1] in the context of fast algorithms for finding shortest paths in planar graphs. The main insight is that it is important to choose a separator which balances not only the size of the components, but also their boundary with the separating set. While this initial work focused on the special case of planar graphs, where $\alpha(n) = O(\sqrt{n})$, Henzinger, Klein, Rao, and Subramanian [HKRS97, Lemma 2.1] proposed an extension of the recursion’s analysis to any sublinear α , providing us with the following bound:

$$\tau_k^{\mathcal{G}}(n) = O\left(n \frac{\alpha(k)}{k}\right)$$

Unfortunately, the recursion is subtle and it turns out their lemma is wrong as stated, but nevertheless holds if α is sufficiently sublinear. Specifically, a separator bound of $n/[(\log n)(\log \log n)^{1+\varepsilon}]$ is sufficient for every $\varepsilon > 0$, whereas $(n/\log n)(\log \log n)^2$ is not. Perhaps surprisingly, these bounds were essentially known, but hiding in a different part of the literature. Specifically, Czumaj, Shapira, and Sohler [CSS09] showed that every hereditary graph class with size- $n/[(\log n)(\log \log n)^2]$ balanced *edge* separators is *hyperfinite*, while Moshkovitz and Shapira [MS18] put forward a hereditary graph class admitting a size- $(n/\log n)(\log \log n)^2$ balanced edge separator despite *not* being hyperfinite. Transposing their bounds into our setting is a straightforward technical exercise. Our humble contribution is strictly limited to pointing out the contradiction.

Returning to our problem, this gives us a sufficient condition for a *hereditary* graph class to admit a sublinear k -path hitting set. In particular, every planar graph of order n admits a k -path hitting set of size $O(n/\sqrt{k})$.

Interestingly, the existence of a k -path hitting set of size α implies the existence of a balanced vertex separator of size $k + \alpha - 1$. Indeed removing a k -path hitting set yields a graph with no k -path, and therefore of treewidth at most $k - 2$, and therefore admitting a balanced vertex separator of size $k - 1$.

We also obtain an interesting “there-exists-implies-for-all”-style collapse: if there exists even a single polynomially sublinear k^* for which a graph class admits a polynomially sublinear k^* -path hitting set, then this same class admits a sublinear k -path hitting set for *any* $k = \omega(1)$.

Weak hyperfiniteness. In the previous paragraph, we described a near-tight condition for a graph to admit a sublinear k -path hitting set *for a fixed function* $k = k(n)$. The near-tightness caveat lies in the range where k is “barely sublinear”. If we shift our attention from the “fixed- k ” regime to “for all $k = \omega(1)$ ”, we are instead able to provide an *exact characterisation* with essentially the same techniques.

A class $\mathcal{G}$ of graphs (resp. directed graphs) is called *fragmentable* or *weakly hyperfinite* (resp *hyperreducible*) if, for every $\varepsilon > 0$, there exists a constant c such that every $G = (V, E) \in \mathcal{G}$ admits a set of at most $\varepsilon |V|$ vertices whose removal causes every connected component to have size (resp. depth) at most c .

By heavily leaning on known results about the relationship between sublinear balanced vertex separators, treewidth, and hyperfiniteness we show that a graph class admits sublinear-size k -path transversals *for every* $k = \omega(1)$ if and only if it is weakly hyperfinite.

⁶A graph class is *hereditary* if it is stable under vertex deletion. In this paper, we use this property to guarantee that after removing a balanced vertex separator from a graph in $\mathcal{G}$, both “sides of the cut” remain in $\mathcal{G}$.

2.3 A planarising set for all 2-degenerate graphs.

Having expanded the set of modularisation targets, that is graph classes for which we can show that there exists a sublinear-size k -path hitting set for some $k = O(\log \log n)$, it remains to show that 2-degenerate graphs do in fact admit a small modulator to one such class.

We establish that every n -vertex 2-degenerate graph admits a planarising set of size $\lfloor n/3 \rfloor$, and provide an algorithm for finding such a set in polynomial time. Our result is heavily inspired by the proof that every order- n 2-degenerate graph admits a feedback vertex set of size at most $\lfloor 2n/5 \rfloor$, by Borowiecki, Drgas-Burchardt, and Sidorowicz [BDBS14, Theorem 2].

Start by fixing an orientation of the graph such that every vertex has in-degree at most 2 (which is always possible since the graph is 2-degenerate). The main idea is that whenever we remove (*i.e.* add to the planarising set) the *last* node with out-degree at least 2, at least two nodes now have in-degree and out-degree at most 1. Without loss of generality, these will never need to be added to the planarising set. This can be formalised by showing that we can apply a series of vertex elimination rules to go from a graph G to H , such that:

- H contains none of G 's vertices with in-degree and out-degree at most 1.
- Any planarising set for H is also a planarising set for G .
- If G is 2-degenerate, then so is H .

When there no longer are any nodes with out-degree at least 2, the graph is a forest, and hence planar. For every node we put in the planarising set, we removed 2 more from the graph for free through the elimination rules, hence the rate of $1/3$.

3 Characterising the easy classes for the k -path hitting set problem.

The goal of this section is to determine some necessary and sufficient conditions for a graph class to admit a sublinear-size k -path hitting set. We investigate two parameter regimes:

- in Section 3.1 we relate the existence of a sublinear-size undirected k -path hitting set to the existence of sublinear-size balanced vertex separators,
- in Section 3.2 we provide characterisations of graph classes admitting sublinear-size k -path hitting sets for *every* $k \in \omega(1)$ (even arbitrarily small).

3.1 Linking k -path transversals to balanced vertex separators

Definition 2 (Balanced Vertex Separator). *A graph $G = (V, E)$ admits a balanced vertex separator if there exists a tripartition $V = A \sqcup S \sqcup B$ such that:*

- $E \cap (A \times B) = \emptyset$
- $|A|, |B| \leq \frac{2}{3}|V|$

We call such an S a balanced vertex separator for G . We then say that a graph class $\mathcal{G}$ is α -separable if every $G = (V, E) \in \mathcal{G}$ admits a balanced vertex separator of size at most $\alpha(|V|)$.

3.1.1 Correcting the record on when hereditary sublinear separators imply graph divisions

While investigating the literature on graph parameters, in search for anchors for k -path transversal numbers, we found a recurring mistake. At a high level, the question is whether the ability to recursively remove sublinear vertex separators is enough to claim that we can always shatter a graph into components of any constant prescribed size, while only paying for a controllable vertex- or edge-deletion cost. As it turns out, this was shown to be false (unless the vertex separators are sufficiently sublinear) by Moshkovitz and Shapira [MS18]. Possibly due to the slight variance in terminology, the contradiction between [MS18] and several relatively high profile works ([HKRS97, Lemma 2.1] and [NOdM12, Theorem 16.5]) was, to the best of our knowledge, never noted explicitly.

The issue with [NOdM12]. The textbook by Nešetřil and Ossona de Mendez [NOdM12, Section 16.3] presents a very detailed proof, making it easy to identify the flaw and to propose a near-tight repair. The unfortunate mistake was asserting that concave functions f such that $f(0) \geq 0$ are super-additive, whereas they are in fact sub-additive. This is a classic result, but omitting the full proof in this correction would be dangerously ironic, so let us recall that: for every $x, y \geq 0$, $f(x) \geq \frac{x}{x+y}f(x+y) + \frac{y}{x+y}f(0)$ (by concavity of f) and so $f(x) \geq \frac{x}{x+y}f(x+y)$ (because $f(0) \geq 0$), and symmetrically $f(y) \geq \frac{y}{x+y}f(x+y)$, and therefore by combining the two $f(x) + f(y) \geq \frac{x}{x+y}f(x+y) + \frac{y}{x+y}f(x+y) = f(x+y)$. Upon closer inspection however:

1. The statement of [NOdM12, Lemma 16.1] is wrong.
2. The statement of [NOdM12, Theorem 16.4] is false, and the proof that it follows from [NOdM12, Lemma 16.1] is entirely correct⁷ up until the last line of the proof (which incorrectly invokes super-additivity).
3. The statement of [NOdM12, Theorem 16.5] is false, even though the proof that it follows from [NOdM12, Lemma 16.4] is entirely correct.

However, instead of fixing this one directly, we use these insights to fix [HKRS97, Lemma 2.1] whose bottomline is strictly stronger.

When hereditary sublinear separators imply graph divisions.

Definition 3 (Division of a graph [Fed87, HKRS97]). *A division of a graph is a partition of its edges. Each of its sets is called a region, and a vertex of the graph belongs to any region containing an edge to which it is incident. A vertex is called an interior vertex if it belongs to a single region, and a boundary vertex otherwise.*

- r -division [Fed87]. *An r -division of an order- n graph is a division into $O(n/r)$ regions, each containing at most r vertices and at most $\sqrt{r}$ boundary vertices.*
- (strict) (r, s) -division [HKRS97]. *An (r, s) -division of an order- n graph is a division into $O(n/r)$ regions, each containing at most $r^{O(1)}$ vertices and at most s boundary vertices. It is a strict (r, s) -division if each region contains at most r vertices.*

Lemma 4 ([HKRS97, Lemma 2.1] is wrong as stated). *There exists $f(n) = o(n)$, a subgraph-closed f -separable class of graphs $\mathcal{G}$, such that $\forall c \in \mathbb{N}$, there exists $r \geq 2$ such that $\mathcal{G}$ does not admit a strict $(r, cf(r))$ -division.*

Proof. By [MS18, Corollary 2.7], there is a hereditary graph class $\mathcal{G}$ such that each order- n graph in $\mathcal{G}$ admits a $\frac{2}{3}$ -balanced edge separator of size at most $f(n) := (n/\log n)(\log \log n)^2 + 48$ (meaning there is a set of at most $f(n)$ edges whose removal disconnects the graph into not-necessarily-separated subgraphs with at most $2n/3$ vertices each), but is not $(1/2, n^{1/4608})$ -hyperfinite (meaning there is no way to remove half its edges in such a way that all resulting connected components have at most $n^{1/4608}$ nodes). As a crucial feature, $\mathcal{G}$ has $\Omega(1)$ average degree (specifically at least 8) and $O(1)$ maximum degree (specifically at most 48) [MS18, Theorem 1].

Since the vertex separation number is upper-bounded by the edge separation number (as, given a balanced edge separator, we can always select one endpoint of each of its edges and obtain a balanced vertex separator), $\mathcal{G}$ is f -separable in the vertex sense. The closure $\mathcal{G}'$ of $\mathcal{G}$ under edge deletions is subgraph-closed, f -separable (because $\mathcal{G}$ is, and removing edges cannot increase a graph's edge separation number).

If [HKRS97, Lemma 2.1] were true, there would exist a $c > 0$ such that every graph in $\mathcal{G}'$ (and in particular every graph in $\mathcal{G}$) would admit a strict $(r, cf(r))$ -division for all $r \in \mathbb{N}^*$. To be precise, there exists $C > 0$ such that every order- n graph in $\mathcal{G}'$ admits a division into at most $C \frac{n}{r}$ regions, each containing at most r nodes in total, and at most $cf(r)$ boundary nodes. Let us now show this would contradict the fact $\mathcal{G}$ is not $(1/2, n^{1/4608})$ -hyperfinite. Let G be an order- n graph in $\mathcal{G}$, and let r be large enough that

⁷Up to a minor typo, where f should be $\widehat{\sigma}$.

$C \cdot c \cdot \frac{f(r)}{r} < 1/24$ (which is possible because $f(r) = o(r)$). Note that r is a fixed constant, independent of n . Fix a strict $(r, cf(r))$ -division of G , and delete every edge incident to a boundary vertex. There are at most $C \frac{n}{r} \cdot cf(r) \leq n/24$ such nodes, and because each has degree at most 48, we have removed at most $2n$ edges. But because the graph has average degree at least 8, it has at least $4n$ edges, and we have therefore removed at most half the edges of G . Because G is not $(1/2, n^{1/4608})$ -hyperfinite, the largest remaining component has size at least $n^{1/4608}$, which is larger than r for large enough n (recall that r is independent of n), which is a contradiction. $\square$

Instead of proving the fixed version of [HKRS97, Lemma 2.1] as intended (that is, as a direct generalisation of the algorithm described in [Fed87]), we use a slight variant where the graph division is obtained by recursively removing *weighted* balanced vertex separator, as it provides an arguably more elegant proof. This is a standard technique in this context (see e.g. [LT79]), and is the proof strategy used in [NOdM12, Section 16.3]. We start by recalling in Lemma 5 the fact that balanced vertex separators imply weighted balanced vertex separators for every weight function (this can be seen as a corrected analogue of [NOdM12, Lemma 16.1], which had this exact same goal).

Lemma 5. *Let f be a non-decreasing function, and let $\mathcal{G}$ be a subgraph-closed f -separable graph class. There exists a constant c such that, for all $n \in \mathbb{N} \setminus \{0, 1\}$, for every $G = (V, E) \in \mathcal{G}_n$, and for every weight function $w: V \mapsto [0, +\infty)$, there exists a partition $V = A \sqcup S \sqcup B$ such that (1) there is no edge between A and B , (2) $w(A), w(B) \leq \frac{2}{3}w(V)$, and (3) $|S| \leq cf(n) + 1$.*

Lemma 5 follows from the fact that graphs admitting a size- f balanced vertex separator have treewidth $O(f)$ [DN19], and that treewidth- t graphs admit a w -balanced separation of order at most $t + 1$ for every non-negative weight function w [RS86]. Armed with this lemma, we are ready to state and prove Theorem 6.

Theorem 6 (Fix of [HKRS97, Lemma 2.1]). *Let f be a non-decreasing function such that $\tilde{f}(n) := \sum_{j=0}^{+\infty} 2^{-j} f(2^j n) = o(n)$. For any subgraph-closed f -separable class of graphs, there is a constant c such that every graph in the family has a strict $(r, c\tilde{f}(r))$ -division for any $r \geq 2$.*

Proof. Let $\mathcal{G}$ be a subgraph-closed f -separable class of graphs and let $G = (V, E) \in \mathcal{G}$ be of order n .

Normalisation. If G has no edge, the empty division satisfies the statement; we therefore assume throughout that G has at least one edge. Since $\mathcal{G}$ is subgraph-closed, $K_2 \in \mathcal{G}$; and in any balanced vertex separation $\{u, v\} = A \sqcup S \sqcup B$ of K_2 we have $|A|, |B| \leq \frac{2}{3} \cdot 2 < 2$, so that $S = \emptyset$ would force the edge uv to run between A and B . Hence $S \neq \emptyset$, and $f(2) \geq 1$. Since f is non-decreasing and $f \leq \tilde{f}$ (the latter being the $j = 0$ term of the series defining $\tilde{f}$), we record that

$$\forall n \geq 2, \quad \tilde{f}(n) \geq f(n) \geq f(2) \geq 1, \quad (1)$$

and that $\tilde{f}$ is non-decreasing, being a termwise sum of non-decreasing functions.

Constants. Let a_0 be a constant such that Lemma 5, applied to a graph of order $m \geq 2$, always returns a separator with $|S| \leq a_0 f(m)$: the constant c of Lemma 5 plus one will do, as Eq. (1) allows its additive $+1$ to be absorbed into $a_0 f(m) \geq a_0 f(2) \geq a_0$. We set $a := 144 a_0$, and note that $a \geq 2$ and that

$$(C1) \quad \frac{2}{3} + \frac{a_0}{a} \leq \frac{3}{4}, \quad (C2) \quad \frac{18a_0}{a} \leq \frac{1}{2}, \quad (C3) \quad \frac{144a_0}{a} \leq 1, \quad (2)$$

since $\frac{2}{3} + \frac{1}{144} \leq \frac{3}{4}$, $\frac{18}{144} = \frac{1}{8} \leq \frac{1}{2}$, and $\frac{144}{144} = 1$. Finally, since $\tilde{f}(n) = o(n)$ by assumption, we may fix $r_0 \geq 2$ such that $a\tilde{f}(m) \leq m$ for every $m \geq r_0$; in the notation of Theorem 6 this amounts to setting $\gamma := 1/a$.

We establish the two following statements, the first of which is used in the proof of the second:

1. for every $r \geq r_0$, G admits a strict $(r, a\tilde{f}(r))$ -division into at most $14\lceil n/r \rceil$ regions;
2. for every $2 \leq r < r_0$, G admits a strict $(r, a\tilde{f}(r))$ -division into at most $C_2\lceil n/r \rceil$ regions, where $C_2 := 14\binom{r_0}{2}$ depends only on f .

Setting $c := a$ and $C := \max(14, C_2)$, the two statements together prove the theorem for every $r \geq 2$.

Case 1: $r \geq r_0$. Put $s := a\tilde{f}(r)$, and note that $s \geq a \geq 2$ by Eq. (1). If $n \leq r$ then the single region $E(G)$ has at most r vertices and no boundary vertex, so we may assume $n > r$.

The splitting process. We grow a rooted binary tree whose nodes carry *pieces*, i.e. subsets of $V(G)$ (distinct nodes may carry equal sets), the root carrying $V(G)$. The *current* pieces are those carried by the leaves. For a current piece P , we let $D(P)$ denote the set of vertices of P that also lie in some other current piece. We run two phases in order:

- *Phase 1.* While some current piece P satisfies $|P| > r$, apply Lemma 5 to $G[P]$ with the uniform weight $w \equiv 1$, obtaining $P = A \sqcup S \sqcup B$, and give the node carrying P the two children $P_1 := A \cup S$ and $P_2 := B \cup S$.
- *Phase 2.* Then, while some current piece P satisfies $|D(P)| > s$, do the same with the weight $w := \mathbf{1}_{D(P)}$, which is legitimate since $|P| \geq |D(P)| > s \geq 2$.

Splitting a piece P into P_1 and P_2 has the following three effects. First, the current pieces remain a family covering every edge of G : an edge with both endpoints in P is not an A – B edge, and hence has both its endpoints in P_1 or both in P_2 . Secondly, $D(Q)$ is unchanged for every *other* current piece Q , because $P_1 \cup P_2 = P$; consequently $D(P)$ is constant throughout the lifetime of P , which makes the notation unambiguous. Thirdly, $D(P_1) = (D(P) \cap A) \sqcup S$ and $D(P_2) = (D(P) \cap B) \sqcup S$, whence

$$|P_1| + |P_2| = |P| + |S| \quad \text{and} \quad |D(P_1)| + |D(P_2)| \leq |D(P)| + 2|S|. \quad (3)$$

Shrinkage and termination. In a Phase-1 split we have $|P| > r \geq r_0$, so $a\tilde{f}(|P|) \leq |P|$ and therefore $|S| \leq a_0 f(|P|) \leq a_0 \tilde{f}(|P|) \leq \frac{a_0}{a} |P|$; both children thus satisfy $|P'| \leq (\frac{2}{3} + \frac{a_0}{a}) |P| \leq \frac{3}{4} |P|$ by (C1). In a Phase-2 split we have $|P| \leq r$, Phase 1 having ended, and $|D(P)| > s = a\tilde{f}(r) \geq af(r) \geq af(|P|)$, so that $|S| \leq a_0 f(|P|) \leq \frac{a_0}{a} |D(P)|$ and, by Eq. (3) and (C1), $|D(P')| \leq \frac{2}{3} |D(P)| + |S| \leq \frac{3}{4} |D(P)|$; moreover $P' \subseteq P$ gives $|P'| \leq r$, so that Phase 2 never recreates an oversized piece. Along any root-to-leaf path the successive Phase-1 sizes lie in $(r, n]$ and the successive Phase-2 values $|D(\cdot)|$ lie in $(s, n]$, and each drops by a factor at least $\frac{4}{3}$; the tree therefore has depth $O(\log n)$ and the process terminates. The final pieces form a family covering every edge of G , each with at most r vertices and at most s shared vertices. Assigning every edge to one final piece containing both of its endpoints induces a division whose regions have at most r vertices and at most $s = a\tilde{f}(r)$ boundary vertices — a boundary vertex of the division lies in two distinct pieces, and is therefore a shared vertex of each — and which has no more regions than there are final pieces.

An antichain bound. Let $\Gamma := \sum_P a_0 f(|P|)$, the sum ranging over the nodes that were split; it is a finite sum, and it dominates $\sum_P |S_P|$. Call a set of nodes a *cut* if it is an antichain meeting every root-to-leaf path. If P is a deepest split node lying strictly above a cut $\mathcal{A}$, then both children of P belong to $\mathcal{A}$: the child on the side of a node of $\mathcal{A}$ below P is itself in $\mathcal{A}$, since otherwise it would be a split node strictly above $\mathcal{A}$ and deeper than P ; and the path through the other child must meet $\mathcal{A}$, which by the same argument can only happen at that child itself. Replacing those two children by P therefore yields a cut with one fewer split node above it, decreasing $\sum |P|$ by exactly $|S_P|$ and $\sum |D(P)|$ by at most $2|S_P|$, by Eq. (3). Inducting down to the cut $\{V(G)\}$, whose two sums are n and 0, gives for every cut $\mathcal{A}$

$$\sum_{P \in \mathcal{A}} |P| \leq n + \Gamma \quad \text{and} \quad \sum_{P \in \mathcal{A}} |D(P)| \leq 2\Gamma; \quad (4)$$

and the same holds for an arbitrary antichain, which may be enlarged to a cut by adjoining every leaf incomparable with it, an operation that only increases both sums.

Counting the splits. Say that a Phase-1 split node has *scale* $j \geq 0$ if $|P| \in (2^j r, 2^{j+1} r]$. Since Phase-1 sizes drop by a factor at least $\frac{4}{3}$ per split and $(\frac{4}{3})^3 > 2$, no root-to-leaf path carries more than three scale- j nodes; grouping the scale- j nodes according to their number of scale- j strict ancestors, which is 0, 1 or 2, therefore partitions them into at most three antichains. As each has $|P| > 2^j r$, Eq. (4) bounds the number of scale- j nodes by $3(n + \Gamma)2^{-j}/r$; using $f(|P|) \leq f(2^{j+1} r)$ and $\sum_{j \geq 0} 2^{-j} f(2^{j+1} r) \leq 2\tilde{f}(r)$, we obtain

$$\Gamma_1 := \sum_{\text{Phase-1}} a_0 f(|P|) \leq \frac{3a_0(n + \Gamma)}{r} \sum_{j \geq 0} 2^{-j} f(2^{j+1} r) \leq \frac{6a_0(n + \Gamma)}{r} \tilde{f}(r).$$

The same argument applies to Phase 2, once one observes that on any root-to-leaf path every Phase-1 split node is an ancestor of every Phase-2 split node — Phase 1 runs to completion before Phase 2 begins, and Phase 2 never recreates an oversized piece — so that only Phase-2 splits, each shrinking $|D(\cdot)|$ by a factor $\frac{3}{4}$, can separate two Phase-2 nodes on a common path. Scaling Phase-2 nodes by $|D(P)| \in (2^i s, 2^{i+1} s]$ and using the second half of Eq. (4) gives at most $6\Gamma 2^{-i}/s$ nodes of scale i , hence at most $12\Gamma/s$ Phase-2 splits in all; each has $|P| \leq r$ and so $a_0 f(|P|) \leq a_0 \tilde{f}(r) = a_0 s/a$, whence $\Gamma_2 \leq \frac{12a_0}{a}\Gamma$. Since $r \geq r_0$ gives $\tilde{f}(r)/r \leq 1/a$, combining the two yields

$$\Gamma = \Gamma_1 + \Gamma_2 \leq \frac{6a_0 n \tilde{f}(r)}{r} + \frac{18a_0}{a}\Gamma,$$

so that, by (C2) and the finiteness of Γ , we get $\Gamma \leq 12a_0 n \tilde{f}(r)/r \leq \frac{12a_0}{a}n \leq n$. The number of regions is at most one more than the number of splits, so, using $n > r$, $\Gamma \leq n$, $s = a\tilde{f}(r)$ and (C3),

$$\#\text{regions} \leq 1 + \frac{6(n + \Gamma)}{r} + \frac{12\Gamma}{s} \leq \frac{n}{r} + \frac{12n}{r} + \frac{144a_0}{a} \cdot \frac{n}{r} \leq \frac{14n}{r} \leq 14 \left\lceil \frac{n}{r} \right\rceil,$$

which establishes the first statement.

Case 2: $2 \leq r < r_0$. Take the division of G in which every edge forms its own region. Each region has exactly $2 \leq r$ vertices, so the division is strict, and has at most 2 boundary vertices, which is at most $a\tilde{f}(r)$ since $a \geq 2$ and $\tilde{f}(r) \geq 1$ by Eq. (1). It remains to bound the number of regions, which is $|E(G)|$. Applying Case 1 with r_0 in place of r partitions $E(G)$ into at most $14\lceil n/r_0 \rceil$ regions, each of which has at most r_0 vertices and therefore spans at most $\binom{r_0}{2}$ edges. Hence

$$|E(G)| \leq 14 \binom{r_0}{2} \left\lceil \frac{n}{r_0} \right\rceil \leq C_2 \left\lceil \frac{n}{r} \right\rceil,$$

the last step because $r \leq r_0$ implies $\lceil n/r_0 \rceil \leq \lceil n/r \rceil$. Since r_0 depends only on f and on the absolute constant a , so does C_2 . $\square$

Corollary 7 (Simple criterion for fixing of [HKRS97, Lemma 2.1]). *Suppose $f(n) \leq n^\varepsilon$ for some $\varepsilon < 1$. For any subgraph-closed f -separable class of graphs, there is a constant c such that every graph in the family has a strict (r, cr^ε) -division for any r .*

Proof. Let $f(n) = n^\varepsilon$, and define $\tilde{f}$ as in Theorem 6. Observe that:

$$\tilde{f}(r) = \sum_{j=0}^{+\infty} 2^{-j} (2^j r)^\varepsilon = r^\varepsilon \sum_{j=0}^{+\infty} (2^{1-\varepsilon})^{-j} = O(r^\varepsilon)$$

$\square$

If we try and relax the separator bound as much as possible, we recover a vertex-analogue of [CSS09, Corollary 3.2] [MS18, Theorem 7].

Lemma 8 (poly log log-tight criterion fix of [HKRS97, Lemma 2.1]). *Let $\varepsilon > 0$, and let $f: \mathbb{N}^* \rightarrow [0, +\infty)$ be a non-decreasing function for which there exists a constant $K > 0$ such that*

$$\forall n \geq 4, \quad f(n) \leq K \frac{n}{(\log n)(\log \log n)^{1+\varepsilon}}.$$

For any subgraph-closed f -separable class of graphs, there is a constant $c > 0$ such that every graph in the class admits, for every $r \geq 4$, a strict $(r, cr(\log \log r)^{-\varepsilon})$ -division.

Proof. Let $g(t) := \frac{1}{t(\log t)^{1+\varepsilon}}$ for $t > 1$. Note that $t \mapsto t(\log t)^{1+\varepsilon}$ is increasing on $(1, +\infty)$, so g is positive and decreasing on this part of its domain.

We start by deriving an explicit upper bound on $\tilde{f}$ as defined in Theorem 6. Let $n \geq 4$, for convenience of notation we denote $X := \log n \geq 2$. Observe that for all $j \in \mathbb{N}$, $2^j n \geq n \geq 4$, and so $f(2^j n) \leq$

$K2^j n / [(\log(2^j n))(\log \log(2^j n))^{1+\varepsilon}]$ by hypothesis. Using the identity $\log(2^j n) = j + \log n = j + X$, we obtain:

$$\tilde{f}(n) \leq K \sum_{j=0}^{+\infty} 2^{-j} \frac{2^j n}{(\log(2^j n))(\log \log(2^j n))^{1+\varepsilon}} = K n \sum_{j=0}^{+\infty} \frac{1}{(j+X)(\log(j+X))^{1+\varepsilon}} = K n S(X),$$

where $S(X) := \sum_{j=0}^{+\infty} g(j+X)$ is well defined for every *real* $X \geq 2$ (convergence follows from Eq. (5) below, and is an instance of Bertrand's criterion [Ber42]).

All that remains is to bound S appropriately. Since g is integrable and decreasing, $g(X+j) \leq \int_{X+j-1}^{X+j} g(t) dt$ for every $j \geq 1$. Thence, summing over $j \geq 1$ and adding the $j = 0$ term separately yields the bound:

$$\begin{aligned} S(X) &\leq g(X) + \int_{t=X}^{+\infty} \frac{dt}{t(\log t)^{1+\varepsilon}} && // \text{series-integral comparison, } g \text{ decreasing} \\ &= g(X) + (\ln 2)^{1+\varepsilon} \int_{u=\ln X}^{+\infty} u^{-1-\varepsilon} du && // \log t = \ln t / \ln 2, \text{ then } u \leftarrow \ln t \text{ and } du = \frac{dt}{t} \\ &= g(X) + \frac{(\ln 2)^{1+\varepsilon}}{\varepsilon (\ln X)^\varepsilon} = g(X) + \frac{\ln 2}{\varepsilon} (\log X)^{-\varepsilon} && // (\ln 2 / \ln X)^\varepsilon = (\log X)^{-\varepsilon} \end{aligned} \tag{5}$$

Finally, $X \geq 2$ gives $\log X \geq 1$, hence $g(X) = \frac{1}{X(\log X)^{1+\varepsilon}} \leq \frac{1}{2}(\log X)^{-\varepsilon}$, and therefore

$$\forall X \geq 2, \quad S(X) \leq \left(\frac{1}{2} + \frac{\ln 2}{\varepsilon} \right) (\log X)^{-\varepsilon}.$$

In turn, substituting $X = \log n$ into the two bounds above yields

$$\forall n \geq 4, \quad \tilde{f}(n) \leq K \left(\frac{1}{2} + \frac{\ln 2}{\varepsilon} \right) \frac{n}{(\log \log n)^\varepsilon},$$

and in particular $\tilde{f}(n) = o(n)$. Therefore, by Theorem 6, there exists a constant $c' > 0$ such that every graph in the class admits, for every $r \geq 2$, a division into regions of size at most r with at most $c' \tilde{f}(r)$ boundary vertices. Restricting to $r \geq 4$ and setting $c := c' K \left(\frac{1}{2} + \frac{\ln 2}{\varepsilon} \right)$ yields the desired bound. $\square$

3.1.2 Some sublinear separators are sufficient

Having corrected the relationship between balanced vertex separators and graph divisions in Section 3.1.1, we can derive sufficient conditions for a graph class to admit sublinear k -path hitting sets for any and all $k = \omega(1)$.

Lemma 9 (From divisions to path transversals). *If a graph $G = (V, E)$ admits a strict (r, s) -division for some $r, s \in \mathbb{N}$, then it admits an $(r+1)$ -path hitting set of size $O(|V| s/r)$.*

Proof. Let $r, s \in \mathbb{N}$, and let $G = (V, E)$ be a graph admitting a strict (r, s) -division. Fix such a division and let S be the set of all its boundary nodes.

- S is an $(r+1)$ -path hitting set. Any path in $G \setminus S$ may only use edges from a single region. Indeed, any vertex incident to edges from two different regions is by definition a boundary vertex, and hence lies in S . Every vertex of such a path therefore belongs to that one region, which contains at most r vertices (recall that the division is *strict*); so no path in $G \setminus S$ has more than r vertices, and S is an $(r+1)$ -path hitting set for G .
- $|S| = O(|V| s/r)$. There are by definition $O(|V|/r)$ regions, and each of them has at most s boundary vertices, so $|S| = O(|V| s/r)$.

$\square$

By combining Corollary 7 and Lemma 9, we obtain Corollary 10.

Corollary 10 (Polynomially sublinear separators are sufficient). *Let α be a polynomially sublinear function, and let $k = \omega(1)$. If $\mathcal{G}$ is an α -separable (Definition 2) hereditary graph class, then every graph $G = (V, E) \in \mathcal{G}$ admits a k -path hitting set of size $O(|V| \alpha(k-1)/k)$.*

In particular, the criterion of Corollary 10 is satisfied by the class of planar graphs, which admit balanced vertex separators of size $O(\sqrt{n})$ [LT79].

Corollary 11 (Planarity is sufficient). *Every planar graph $G = (V, E)$ admits a k -path hitting set of size $O(|V|/\sqrt{k})$.*

If we return to the results of Section 3.1.1 and focus on getting a criterion which is very close to tight (even if it means getting no significantly better bound than $o(n)$ for the size of the hitting set), then we get Corollary 12 (which is obtained by combining Lemmata 8 and 9).

Corollary 12 (Slightly sublinear separators are sufficient). *Let $\varepsilon > 0$, let $\alpha = O\left(\frac{n}{(\log n)(\log \log n)^{1+\varepsilon}}\right)$, and let $k = \omega(1)$. If $\mathcal{G}$ is an α -separable (Definition 2) hereditary graph class, then every graph $G = (V, E) \in \mathcal{G}$ admits a k -path hitting set of size $o(|V|)$.*

3.1.3 Sublinear separators are necessary

If a graph class admits a sublinear k -path hitting set for some sublinear k , then it has a sublinear balanced vertex separator.

Lemma 13. *Every P_k -free graph has tree-width at most $k - 2$.*

Proof. By [NOdM12, Proposition 6.1], any P_k -free graph has treedepth at most $k - 1$, and therefore treewidth at most $k - 2$. $\square$

Theorem 14 (Sublinear separators are necessary). *Let $k, \alpha: \mathbb{N}^* \rightarrow \mathbb{N}^*$. If a graph class admits a size- α k -path hitting set, then it is $(\alpha + k - 1)$ -separable.*

Proof. Let $G = (V, E)$ be a graph. If X is a k -path hitting set for G , then $G \setminus X$ is P_k -free by definition. But any P_k -free graph class has treewidth at most $k - 2$ (Lemma 13), and therefore admits a balanced vertex separator of size $k - 1$ [RS86]. That is, there exists a partition $A \sqcup B \sqcup S$ of $V \setminus X$ such that (1) $|A|, |B| \leq 2|V \setminus X|/3 \leq 2|V|/3$, (2) $|S| \leq k - 1$, and (3) $(A \times B) \cap E = \emptyset$. Therefore $S \cup X$ constitutes a size- $(\alpha + k - 1)$ balanced vertex separator for G . $\square$

As a corollary (of Corollary 12 and Theorem 14) of particular interest, note that if a graph class admits a size- α k -path hitting set for *some* $k, \alpha \in O\left(\frac{n}{(\log n)(\log \log n)^{1+\varepsilon}}\right)$ (where $\varepsilon \in (0, 1)$), then this graph class admits sublinear k -path hitting sets for *every* $k = \omega(1)$.

3.2 Linking k -path transversals in the “for all k ” regime to weak hyperfiniteness

In this section, we leave the “fixed- k ” regime and instead study the graph classes admitting sublinear-size k -path hitting sets for *every* $k = \omega(1)$ (even arbitrarily small). Interestingly these graph classes can be fully characterised by a single parameter. This characterisation uncovers a partial there-exists- k -to-for-all- k collapse in the undirected case.

3.2.1 The undirected case: weak hyperfiniteness

Definition 15 (Fragmentability [EM94, EF01] / Weak Hyperfiniteness [NOdM10] [NOdM12, Chapter 16]). *A graph class $\mathcal{G}$ is weakly hyperfinite (or fragmentable) if for every $\varepsilon > 0$, there exists $c_\varepsilon \geq 0$ such that for any $G = (V, E) \in \mathcal{G}$, there exists a set $S \subseteq V$ of size $|S| \leq \varepsilon|V|$ such that every connected component of $G \setminus S$ has at most c_ε vertices.*

Theorem 16. *Let $\mathcal{G}$ be a class of simple undirected graphs. The following statements are equivalent:*

- For every function $k = \omega(1)$, there exists a sublinear function α such that every n -vertex graph in $\mathcal{G}$ admits a $k(n)$ -path hitting set of size $\alpha(n)$.
- $\mathcal{G}$ is fragmentable. ▷ Terminology of [EM94].
- $\mathcal{G}$ is weakly hyperfinite. ▷ Terminology of [NOdM10].

Proving that weak hyperfiniteness is sufficient for a graph class to admit sublinear k -path transversals for all diverging k is a straightforward exercise in manipulating limits and quantifiers. The core observation is that if a graph is sufficiently fragmented, it cannot contain long paths. Conversely, weak hyperfiniteness is necessary because P_k -free graphs have bounded treedepth.

We start by introducing some convenient notations. If $\mathcal{G}$ is a graph class, then for every $n \in \mathbb{N}$, we denote $\mathcal{G}_n$ as the set of order- n graphs in $\mathcal{G}$. Furthermore, for every function $k: \mathbb{N} \rightarrow \mathbb{N}$, we define

$$\alpha_k(n) := \begin{cases} \max_{G \in \mathcal{G}_n} \{\tau_{k(n)}(G)\} & \text{if } \mathcal{G}_n \neq \emptyset \\ 0 & \text{if } \mathcal{G}_n = \emptyset \end{cases}$$

and note that α_k is well-defined since $|\mathcal{G}_n|$ is finite for any fixed choice of n .

Weak hyperfiniteness is sufficient. Let $\mathcal{G}$ be a weakly hyperfinite graph class, and let $\mathcal{G}$ and $\alpha_k(\cdot)$ be defined as above. Let us now show that $\forall k \in \omega(1), \alpha_k(n) = o(n)$, or equivalently that $\forall k \in \omega(1), \forall \varepsilon > 0, \exists N > 0, \forall n \geq N, \alpha_k(n) \leq \varepsilon n$:

$\mathcal{G}$ is weakly hyperfinite

$$\begin{aligned} \Leftrightarrow & \quad \forall \varepsilon > 0, \exists c_\varepsilon > 0, \forall G = (V, E) \in \mathcal{G}, \exists S \in \binom{V}{\leq \varepsilon |V|}, \\ & \quad \text{every connected component of } G \setminus S \text{ has } \leq c_\varepsilon \text{ vertices} \\ & \quad // \text{ By definition of weakly hyperfinite} \\ \Rightarrow & \quad \forall \varepsilon > 0, \exists c_\varepsilon > 0, \forall G = (V, E) \in \mathcal{G}, \exists S \in \binom{V}{\leq \varepsilon |V|}, G \setminus S \text{ is } (c_\varepsilon + 1)\text{-path free} \\ & \quad // \text{ Each path is contained in a single connected component.} \\ \Rightarrow & \quad \forall \varepsilon > 0, \exists c_\varepsilon > 0, \forall n \in \mathbb{N}, \forall G \in \mathcal{G}_n, \tau_{c_\varepsilon + 1}(G) \leq \varepsilon n \\ & \quad // \text{ By definition of the } (c_\varepsilon + 1)\text{-path transversal number } \tau_{c_\varepsilon + 1} \\ \Rightarrow & \quad \forall \varepsilon > 0, \forall k \in \omega(1), \exists n_\varepsilon \in \mathbb{N}, \forall n \geq n_\varepsilon, \forall G \in \mathcal{G}_n, \tau_{k(n)}(G) \leq \varepsilon n \\ & \quad // \text{ If } k \text{ diverges to infinity, } \exists n_\varepsilon \in \mathbb{N}, \forall n \geq n_\varepsilon, k(n) \geq c_\varepsilon + 1 \\ & \quad // \text{ and for all } n \in \mathbb{N}, \kappa \mapsto \tau_\kappa(n) \text{ is non-increasing} \\ \Rightarrow & \quad \forall \varepsilon > 0, \forall k \in \omega(1), \exists n_\varepsilon \in \mathbb{N}, \forall n \geq n_\varepsilon, \alpha_k(n) \leq \varepsilon n \\ \Rightarrow & \quad \forall k \in \omega(1), \forall \varepsilon > 0, \exists n_\varepsilon \in \mathbb{N}, \forall n \geq n_\varepsilon, \alpha_k(n) \leq \varepsilon n \\ & \quad // \text{ Commuting quantifiers} \end{aligned}$$

Weak hyperfiniteness is necessary. Let $\mathcal{G}$ be a graph class admitting sublinear k -path transversals for every diverging k , which implies that for all $k \in \omega(1)$, $\alpha_k(n) = o(n)$. Because the property we wish to prove (namely, weak hyperfiniteness or fragmentability) has a flavour of “uniform convergence” (in the sense that c_ε does not depend on n), we start by showing that our assumption on $\mathcal{G}$ allows us to asymptotically bound $\tau_k(G)/|V(G)|$ without requiring k to be a function of n .

Lemma 17.

$$\forall \delta > 0, \exists K_\delta > 0, \exists N_\delta \in \mathbb{N} \setminus \{0, 1\}, \forall n \geq N_\delta, \forall G \in \mathcal{G}_n, \tau_{K_\delta}(G) \leq \delta n \quad (6)$$

Proof of Lemma 17. Towards a contradiction, assume that Eq. (6) is false. Therefore, there exists $\delta > 0$ such that

$$[\text{Stated for contradiction}] \quad \forall K_\delta > 0, \forall N_\delta \in \mathbb{N} \setminus \{0, 1\}, \exists n \geq N_\delta, \exists G \in \mathcal{G}_n, \tau_{K_\delta}(G) > \delta n \quad (7)$$

Applying Eq. (7) with $K_\delta^{(0)} := 1$ and $N_\delta^{(0)} := 2$ guarantees the existence of an $n_0 \geq N_\delta^{(0)}$ and a $G_0 \in \mathcal{G}_{n_0}$ such that $\tau_{K_\delta^{(0)}}(G_0) > \delta n_0$. For every $i \in \mathbb{N}^*$, recursively applying Eq. (7) with $K_\delta^{(i)} := \max\{i, K_\delta^{(i-1)} + 1\}$ and $N_\delta^{(i)} := n_{i-1} + 1$ guarantees the existence of an $n_i \geq N_\delta^{(i)}$ and a $G_i \in \mathcal{G}_{n_i}$ such that $\tau_{K_\delta^{(i)}}(G_i) > \delta n_i$. Lastly consider the function k defined as follows:

$$\begin{aligned} k: \mathbb{N}^* &\rightarrow \mathbb{N}^* \\ n &\mapsto \begin{cases} \max\{i \in \mathbb{N}^* \mid n_i \leq n\} & \text{if } n \geq n_0 \\ 1 & \text{otherwise} \end{cases} \end{aligned}$$

Note that the sequence $(n_i)_{i \in \mathbb{N}^*}$ is increasing (which in particular means k is well-defined), and therefore (being integer-valued) diverges to $+\infty$. Note that k is non-decreasing, integer-valued, and admits a diverging subsequence (*e.g.* because $\forall i \in \mathbb{N}, k(n_i) = i$), and so $k = \omega(1)$. And so (by assumption on $\mathcal{G}$) $\alpha_k(n) = o(n)$. And yet α_k diverges on the subsequence $(n_i)_{i \in \mathbb{N}^*}$, which is a contradiction: $\forall i \in \mathbb{N}^*, \alpha_k(n_i) \geq \tau_{k(n_i)}(G_i) = \tau_i(G_i) \geq \tau_{K_\delta^{(i)}}(G_i) > \delta n_i$ (where the inequality $\tau_i(G_i) \geq \tau_{K_\delta^{(i)}}(G_i)$ follows from $i \leq K_\delta^{(i)}$ by definition of $K_\delta^{(i)}$, and from the monotony of $\tau(n)$ for any fixed value of n). $\square$

The second lemma we need states that bounded-treewidth graphs (such as those resulting from removing a k -path hitting set) are weakly hyperfinite. We borrow a lemma from [GLL⁺19, Lemma 10], which is more conveniently expressed in terms of individual graphs rather than graph classes.

Lemma 18 (Bounded treewidth implies small fragmenting sets, immediate corollary of [GLL⁺19, Lemma 10]). *If an order- n graph G has treewidth at most d , then for every $\eta \in (0, 1)$ there exists a subset X of at most ηn nodes such that every component of $G \setminus X$ has at most $\lceil d/\eta \rceil$ vertices. Furthermore, such an X can be found in polynomial time.*

We now resume the proof that $\mathcal{G}$ is weakly hyperfinite. Let $\varepsilon > 0$. Applying Lemma 17 to $\delta := \varepsilon/2$ guarantees the existence of $K > 0$ and $N \geq 2$ such that $\forall n \geq N, \alpha_K(n) \leq \frac{\varepsilon}{2}n$. We are now ready to define $c_\varepsilon := 2(K-2)/\varepsilon$, and to show that it satisfies the requirements of Definition 15. Let $G = (V, E) \in \mathcal{G}$. By Lemma 17 (with the aforementioned parameterisation), there exists $Y \subseteq V$ such that $G \setminus Y$ is K -path free and $|Y| \leq \frac{\varepsilon}{2}|V|$. Being K -path free, $G \setminus Y$ has treewidth at most $K-2$ Lemma 13, and therefore, by applying Lemma 18 to $d := K-2$ and $\eta := \varepsilon/2$, there exists $X \subseteq V$ such that every component of $(G \setminus Y) \setminus X$ has at most $\lceil d/\eta \rceil = 2(K-2)/\varepsilon = c_\varepsilon$ vertices. The set $S := X \cup Y$ is therefore an appropriately fragmenting set of G , and $|S| \leq |X| + |Y| \leq \varepsilon|V|$.

4 A planarising set for all 2-degenerate graphs

The goal of this section is to show that every n -vertex 2-degenerate graph contains a planarising set (*i.e.* a set of vertices whose removal yields a planar graph) of size at most $\lfloor n/3 \rfloor$, and that furthermore such a set can be found in polynomial time.

Theorem 19. *Every 2-degenerate graph of order n admits a planar induced subgraph of order at least $\lfloor 2n/3 \rfloor$, and one such subgraph can be found in polynomial time.*

In order to prove Theorem 19, let us start by establishing some helpful lemmata.

Lemma 20. *The following (addition) operations preserve a directed multigraph's planarity⁸:*

⁸In this paper, a directed multigraph is said to be *planar* if and only if its underlying undirected multigraph is, which in turn is equivalent to its underlying simple underlying graph being planar. The orientation and multiplicity of arcs are completely irrelevant as far as planarity is concerned. We find the directed multigraph terminology convenient, as it allows us to add such instructions as “add an arc” without discussing the cases where such an arc already exists.

- Add isolated vertex. *Add a fresh vertex (with no incident arcs).*
- Add pendant. *Add a fresh vertex u , then choose any vertex $v \neq u$, and add either the arc (u, v) or the arc (v, u) .*
- Subdivide arc. *Add a fresh vertex u to the graph then replace any arc (x, y) with arcs (x, u) and (u, y) .*

As a corollary, if S is a planarising vertex set for a directed multigraph G , and H can be obtained by successively applying these rules to G , then S is also a planarising set for H .

Proof. It should be clear from inspection that these operations preserve planarity, so we focus our efforts on the corollary. It suffices to consider the case where H can be obtained from G by applying a single step (and induction handles the rest). In short, no matter which operation is applied, $H \setminus S$ can be obtained from $G \setminus S$ by applying one of the lemma's operations, which have already been shown to preserve planarity. Therefore, if S is a planarising set for G , then it is also one for H . In more detail, we proceed on a case analysis on how vertex u was added to G :

- *As an isolated vertex.* Then $H \setminus S$ is exactly $G \setminus S$ plus an isolated vertex.
- *As a pendant anchored with v .* If $v \in S$, then $H \setminus S$ is exactly $G \setminus S$ plus an isolated vertex. If $v \notin S$, then $H \setminus S$ is exactly $G \setminus S$ plus the pendant (u, v) or (v, u) .
- *Subdividing an arc (x, y) .* If $x \in S \wedge y \in S$, then $H \setminus S$ is exactly $G \setminus S$ plus an isolated vertex. If $x \in S \wedge y \notin S$, then $H \setminus S$ is exactly $G \setminus S$ plus the pendant (u, y) anchored in y . If $x \notin S \wedge y \in S$, then $H \setminus S$ is exactly $G \setminus S$ plus the pendant (x, u) anchored in x . If $x \notin S \wedge y \notin S$, then $H \setminus S$ can be obtained by subdividing arc (x, y) in $G \setminus S$.

□

Lemma 21. *The following (deletion) operations preserve a directed multigraph's acyclicity:*

- Remove isolated vertex. *Remove a vertex with in- and out-degree 0.*
- Remove pendant. *Remove a vertex with in-degree 1 and out-degree 0, or in-degree 0 and out-degree 1.*
- Smooth vertex. *Remove a vertex with in-degree and out-degree 1, then add an arc from its unique in-neighbour to its unique out-neighbour.*

Proof. The first two operations do not add any arcs, and therefore cannot introduce any cycles. As for the third, let H be the directed multigraph obtained by smoothing some node u in G . If G is acyclic, any cycle in H would have to use the freshly added arc (x, y) . But such a cycle would imply a path from y to x in H , and therefore also in G . There would then be a cycle $x \rightarrow u \rightarrow y \rightarrow \dots x$ in G , which would contradict G 's acyclicity. Therefore, if G is acyclic, then so is H . □

Observation 1 is a helpful lemma whose proof is immediate.

Observation 1. *The following (deletion) operations cannot increase any of the remaining vertices' in- or out-degrees:*

- Remove isolated node. *Remove any node with in-degree and out-degree 0.*
- Remove pendant. *Remove any node with in-degree 1 and out-degree 0, or with in-degree 0 and out-degree 1.*
- Vertex smoothing. *Remove a node with in- and out-degree both exactly 1, then add an arc connecting its unique in-neighbour to its unique out-neighbour.*

We are now ready to prove Theorem 19.

Proof of Theorem 19. Consider the following recursive algorithm **PlanarSet**, which takes as input a directed acyclic graph G where each vertex has in-degree at most 2:

$\text{PlanarSet}(G)$:

1. If G has order 3 or less, output $\emptyset$
2. If every node has out-degree at most 1, output $\emptyset$
 - // A finite DAG where every node has out-degree at most 1 is indeed a forest.
 - // On the face of it, each of its connected components can have at most as many arcs ...
 - // ... as it has vertices, but since every component contains a sink, there must be *fewer* ...
 - // ... arcs than vertices in each component. Each component is therefore a tree.
3. Otherwise, consider a topological ordering of G , and let u be the last node with out-degree at least 2. Consider the multigraph H obtained by starting with $G \setminus u$ then recursively applying the following rules until none are possible: $\triangleright$ This process terminates in polynomial time as each of the four rules strictly decreases the number of vertices.
 - remove any node with in-degree and out-degree 0
 - remove any node with in-degree 1 and out-degree 0
 - remove any node with in-degree 0 and out-degree 1
 - smooth any node with in-degree and out-degree exactly 1
 - // Since this operation introduces an arc, it may create parallel arc.
 - // Describing this operation succinctly is the motivation for considering directed multigraphs.
4. Let G' denote the simple di-graph obtained by merging any and all parallel arcs in directed multigraph H , and output $\{u\} \cup \text{PlanarSet}(G')$

Let us now show, by induction on $n \in \mathbb{N}$, that for every order- n DAG G with in-degree ≤ 2 , $\text{PlanarSet}(G)$ correctly returns a planarising set for G of size at most $\lfloor n/3 \rfloor$.

- *Initialisation.* If $n \leq 3$, $\text{PlanarSet}(G)$ outputs $\emptyset$. Since every DAG on three vertices or fewer is planar, this is correct.
- *Induction step.* Let $n \geq 4$, and assume the induction hypothesis is true up to order $n - 1$. Let G be an order- n DAG of fan-in at most 2. If all of G 's vertices have out-degree at most 1, then G is a forest, and therefore planar. $\text{PlanarSet}(G)$ outputs $\emptyset$, which is correct. Otherwise, let u be the last node with out-degree at least 2. Let H and G' be defined as in the algorithm. Observe that:
 - G' is *acyclic*. $G \setminus u$ is an (induced) subgraph of G , which is acyclic, so $G \setminus u$ is acyclic. By Lemma 21, each of the steps used to turn $G \setminus u$ into H preserves acyclicity, so H is acyclic. In turn, since G' is a subgraph of H , G' is acyclic.
 - G' has *in-degree at most 2*. By Observation 1, G' inherits this property from G .
 - G' has *order at most $n - 1$* . G' is obtained by removing at least one vertex from G (namely u), and then potentially adding and removing arcs, but never adding any vertices. Therefore G' has order at most $n - 1$.
 - $\text{PlanarSet}(G')$ is a *planarising set for H* . Since G' is a DAG and has order at most $n - 1$, we can apply the induction hypothesis and conclude that $\text{PlanarSet}(G')$ is a planarising set for G' . By construction, $H \setminus \text{PlanarSet}(G')$ can be obtained from $G' \setminus \text{PlanarSet}(G')$ by replacing arcs with any finite number of parallel arcs. Since this operation preserves planarity, $\text{PlanarSet}(G')$ is a planarising set for H .
 - $\text{PlanarSet}(G') \sqcup \{u\}$ is a *planarising set for G* . Observe that each of the algorithm's four elimination rules (which were used to turn $G \setminus u$ into H) can be inverted by applying one of the three addition rules of Lemma 20. Therefore, by Lemma 20, since $\text{PlanarSet}(G')$ is a planarising set for H , it is also a planarising set for $G \setminus u$. It follows that $\text{PlanarSet}(G') \sqcup \{u\}$ is a planarising set for G .
 - G' has *order at most $n - 3$* . Recall that u has out-degree at least 2. Let us denote x and y two distinct out-neighbours of u . Since the graph has in-degree at most 2, x and y each have in-degree

at most 1 in $G \setminus u$. But because we defined u to be the last node with out-degree at least 2, and because they appear later than u in any topological ordering, x and y must also have out-degree at most 1. But since the algorithm's four elimination rules can never increase the remaining nodes' in- or out-degrees (Observation 1), x and y will always have in- and out-degree at most 1, until they are eliminated. This guarantees they will both eventually be eliminated. Thence G' contains neither u , x , nor y , and has order at most $n - 3$.

Given all these facts, we are ready to conclude. By the induction hypothesis, since G' has order at most $n - 3$, $|\text{PlanarSet}(G')| \leq \lfloor (n - 3)/3 \rfloor = \lfloor n/3 \rfloor - 1$. Therefore G admits a planarising set (specifically $\text{PlanarSet}(G') \sqcup \{u\}$) of size at most $\lfloor n/3 \rfloor$.

In order to conclude properly, we need to (1) solve the terminology mismatch between Theorem 19 and the guarantees of `PlanarSet`, and (2) bound the latter's running time.

- *From PlanarSet to Theorem 19.* Let $G = (V, E)$ be an undirected 2-degenerate graph of order n . Being 2-degenerate, G admits a 2-elimination ordering $v_1, \dots, v_n$ of V , i.e. one in which every v_i has at most two neighbours among $v_1, \dots, v_{i-1}$: it suffices to repeatedly remove a vertex of degree at most 2 from the graph induced by the vertices not yet removed (such a vertex exists at every step, precisely because every induced subgraph of a 2-degenerate graph is itself 2-degenerate), and to read the removal order backwards. Orienting every edge of G from its earlier endpoint to its later one therefore yields a directed graph $\vec{G}$ in which every vertex has in-degree at most 2, and which is acyclic since each of its arcs respects the ordering $v_1, \dots, v_n$. Running `PlanarSet` on $\vec{G}$ produces, by the induction above, a set $S \subseteq V$ with $|S| \leq \lfloor n/3 \rfloor$ such that $\vec{G} \setminus S$ is planar. Since planarity of a directed multigraph is by definition planarity of its underlying undirected multigraph, and since orienting the edges of G leaves that underlying graph unchanged, the induced subgraph $G \setminus S$ is planar as well. Finally, as n is an integer,

$$|V \setminus S| = n - |S| \geq n - \left\lfloor \frac{n}{3} \right\rfloor = \left\lceil n - \frac{n}{3} \right\rceil = \left\lceil \frac{2n}{3} \right\rceil,$$

so $G \setminus S$ is indeed a planar induced subgraph of G of order at least $\lceil 2n/3 \rceil$.

- *Running time.* Since $\vec{G}$ has in-degree at most 2, it has at most $2n$ arcs; furthermore none of the operations performed by `PlanarSet` ever increases the number of arcs (smoothing replaces two arcs by one, while the three other elimination rules and the merging of parallel arcs only delete arcs), so every multigraph arising during the computation has at most $2n$ arcs. A single call to `PlanarSet` computes a topological ordering and locates u in time $O(n)$; it then performs at most n elimination steps, since each of the four rules deletes a vertex, and an applicable rule can be located in time $O(n)$ by scanning the current vertex set; merging parallel arcs costs $O(n \log n)$. A single call therefore costs $O(n^2)$. As each recursive call is made on a graph of order at most $n - 3$, at most $n/3$ calls are made in total, and `PlanarSet` runs in time $O(n^3)$. Since finding a 2-elimination ordering of G can be done in polynomial time, a planar induced subgraph of order at least $\lceil 2n/3 \rceil$ can indeed be found in polynomial time.

□

5 A trade-off between the size and the fan-in of circuits

In this section, we combine the modularisation bound of Theorem 19, the depth-reduction bound of Corollary 11, and the [CCMN24] template recalled in Section 2, to obtain the desired trade-off.

Corollary 22 (*k-path transversals in 2-degenerate graphs*). *Let $k = \omega(1)$. Every 2-degenerate graph $G = (V, E)$ of order n admits a k -path hitting set of size $n/3 + O(n/\sqrt{k}) = n/3 + o(n)$.*

Proof. By Theorem 19, there exists $X \subseteq V$ with $|X| \leq \lfloor n/3 \rfloor$ such that $G \setminus X$ is planar. By Corollary 11, the planar graph $G \setminus X$ admits a k -path hitting set Y of size $O(|V \setminus X|/\sqrt{k}) = O(n/\sqrt{k})$. Since $(G \setminus X) \setminus Y = G \setminus (X \cup Y)$ contains no k -path, the set $X \cup Y$ is a k -path hitting set for G , and $|X \cup Y| \leq \lfloor n/3 \rfloor + O(n/\sqrt{k})$. □

Because every directed k -path is in particular a k -path of the underlying undirected graph, the same bound holds for *directed* k -path transversals.

Corollary 23 (Directed k -path transversals in fan-in 2 DAGs). *Let $k = \omega(1)$. Every DAG $G = (V, E)$ of order n whose vertices all have in-degree at most 2 admits a directed k -path hitting set of size $n/3 + O(n/\sqrt{k})$.*

Proof. The undirected graph underlying G is 2-degenerate: any of its subgraphs is the underlying graph of a sub-DAG of G , whose last vertex in any topological ordering has out-degree 0 and in-degree at most 2, hence degree at most 2. Corollary 22 therefore provides an (undirected) k -path hitting set S of the announced size. Removing S leaves an undirected graph with no k -path, and hence a directed graph with no directed k -path, since every directed k -path of $G \setminus S$ would be a k -path of the graph underlying $G \setminus S$. $\square$

Plugging Corollary 23 into the template of Charbit, Couteau, Meyer, and Naserasr [CCMN24] (recalled in Section 2) yields the circuit-complexity statement announced in the introduction.

Corollary 24 (A size/fan-in trade-off for circuits). *Let $k = \omega(1)$. Any function computable by a fan-in 2 circuit with n inputs, m outputs, and s internal gates is computable by a circuit with n inputs, m outputs, and $s/3 + O(s/\sqrt{k}) = s/3 + o(s)$ internal gates, each of fan-in at most 2^k .*

Proof. Let C be such a circuit and let G be the DAG obtained from C by deleting the vertices corresponding to its inputs and its outputs. Then G has order s , and every one of its vertices has in-degree at most 2 because C has fan-in 2. By Corollary 23, G admits a directed k -path hitting set of size $\frac{s}{3}(1 + O(1/\sqrt{k}))$; selecting these gates, together with all the inputs and outputs of C , meets the two requirements of the template of Section 2, which therefore returns a circuit of the announced size and fan-in. $\square$

AI Disclosure. We used Large Language Models Opus 4.8, Fable 5, and GPT 5.6 Sol in order to explore the literature downstream of [HKRS97, NOdM12] in search of instances of results needing a significant correction due to the bug we uncovered. Proving the fixed version of the lemma (namely Theorem 6)—which we would consider a routine exercise—was obtained (and then reworked into a more “human-readable” form by the author) by an LLM with instructions to fix the proof of [NOdM12] according to the following methodology:

1. Apply the observations of paragraph “The issue with [NOdM12].” of Section 3.1.1 to fix [NOdM12, Theorem 16.5]
2. Translate the proof from a statement about weak hyperfiniteness (as considered in [NOdM12, Theorem 16.5]) to a statement about strict divisions (as considered in [HKRS97, Lemma 2.1]).

The author verified the correctness and originality of all content including references.

References

- [ADOS22] Damiano Abram, Ivan Damgård, Claudio Orlandi, and Peter Scholl. An algebraic framework for silent preprocessing with trustless setup and active security. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part IV*, volume 13510 of *LNCS*, pages 421–452. Springer, Cham, August 2022.
- [BCM22] Elette Boyle, Geoffroy Couteau, and Pierre Meyer. Sublinear secure computation from new assumptions. In Eike Kiltz and Vinod Vaikuntanathan, editors, *TCC 2022, Part II*, volume 13748 of *LNCS*, pages 121–150. Springer, Cham, November 2022.
- [BCM23] Elette Boyle, Geoffroy Couteau, and Pierre Meyer. Sublinear-communication secure multi-party computation does not require FHE. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part II*, volume 14005 of *LNCS*, pages 159–189. Springer, Cham, April 2023.
- [BDBS14] Mieczysław Borowiecki, Ewa Drgas-Burchardt, and Elżbieta Sidorowicz. A feedback vertex set of 2-degenerate graphs. *Theoretical Computer Science*, 557:50–58, 2014.
- [Ber42] Joseph Bertrand. Règles sur la convergence des séries. *Journal de Mathématiques Pures et Appliquées*, 7:35–54, 1842.

- [BGI16] Elette Boyle, Niv Gilboa, and Yuval Ishai. Breaking the circuit size barrier for secure computation under DDH. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part I*, volume 9814 of *LNCS*, pages 509–539. Springer, Berlin, Heidelberg, August 2016.
- [BGW88] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). In *20th ACM STOC*, pages 1–10. ACM Press, May 1988.
- [BJT22] Csilla Bujtás, Marko Jakovac, and Zsolt Tuza. The k -path vertex cover: General bounds and chordal graphs. *Networks*, 80(1):63–76, 2022.
- [BKKS11] Boštjan Brešar, František Kardoš, Ján Katrenič, and Gabriel Semanišin. Minimum k -path vertex cover. *Discrete Applied Mathematics*, 159(12):1189–1195, 2011.
- [BKS19] Elette Boyle, Lisa Kohl, and Peter Scholl. Homomorphic secret sharing from lattices without FHE. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part II*, volume 11477 of *LNCS*, pages 3–33. Springer, Cham, May 2019.
- [BLMZ17] Fabrice Benhamouda, Tancrede Lepoint, Claire Mathieu, and Hang Zhou. Optimization of bootstrapping in circuits. In Philip N. Klein, editor, *28th SODA*, pages 2423–2433. ACM-SIAM, January 2017.
- [CCD88] David Chaum, Claude Crépeau, and Ivan Damgård. Multiparty unconditionally secure protocols (extended abstract). In *20th ACM STOC*, pages 11–19. ACM Press, May 1988.
- [CCMN24] Pierre Charbit, Geoffroy Couteau, Pierre Meyer, and Reza Naserasr. A note on low-communication secure multiparty computation via circuit depth-reduction. In Elette Boyle and Mohammad Mahmoody, editors, *TCC 2024, Part IV*, volume 15367 of *LNCS*, pages 167–199. Springer, Cham, December 2024.
- [CK89] Benny Chor and Eyal Kushilevitz. A zero-one law for Boolean privacy (extended abstract). In *21st ACM STOC*, pages 62–72. ACM Press, May 1989.
- [CM21] Geoffroy Couteau and Pierre Meyer. Breaking the circuit size barrier for secure computation under quasi-polynomial LPN. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part II*, volume 12697 of *LNCS*, pages 842–870. Springer, Cham, October 2021.
- [COS⁺22] Ilaria Chillotti, Emmanuela Orsini, Peter Scholl, Nigel P. Smart, and Barry van Leeuwen. Scooby: Improved multi-party homomorphic secret sharing based on FHE. In Clemente Galdi and Stanislaw Jarecki, editors, *SCN 22*, volume 13409 of *LNCS*, pages 540–563. Springer, Cham, September 2022.
- [Cou19] Geoffroy Couteau. A note on the communication complexity of multiparty computation in the correlated randomness model. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part II*, volume 11477 of *LNCS*, pages 473–503. Springer, Cham, May 2019.
- [CSS09] Artur Czumaj, Asaf Shapira, and Christian Sohler. Testing hereditary properties of nonexpanding bounded-degree graphs. *SIAM Journal on Computing*, 38(6):2499–2510, 2009.
- [DIJL23] Quang Dao, Yuval Ishai, Aayush Jain, and Huijia Lin. Multi-party homomorphic secret sharing and sublinear MPC from sparse LPN. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part II*, volume 14082 of *LNCS*, pages 315–348. Springer, Cham, August 2023.
- [DN19] Zdeněk Dvořák and Sergey Norin. Treewidth of graphs with balanced separations. *Journal of Combinatorial Theory, Series B*, 137:137–144, 2019.
- [EF01] Keith Edwards and Graham Farr. Fragmentability of graphs. *Journal of Combinatorial Theory, Series B*, 82(1):30–37, 2001.

- [EM94] Keith Edwards and Colin McDiarmid. New upper bounds on harmonious colorings. *Journal of Graph Theory*, 18(3):257–267, 1994.
- [Fed87] Greg N. Frederickson. Fast algorithms for shortest paths in planar graphs, with applications. *SIAM Journal on Computing*, 16(6):1004–1022, 1987.
- [Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 169–178. ACM Press, May / June 2009.
- [GLL⁺19] Anupam Gupta, Euiwoong Lee, Jason Li, Pasin Manurangsi, and Michał Włodarczyk. Losing tree-width by separating subsets. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA ’19, page 1731–1749, USA, 2019. Society for Industrial and Applied Mathematics.
- [GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In Alfred Aho, editor, *19th ACM STOC*, pages 218–229. ACM Press, May 1987.
- [HKRS97] Monika R Henzinger, Philip Klein, Satish Rao, and Sairam Subramanian. Faster shortest-path algorithms for planar graphs. *Journal of Computer and System Sciences*, 55(1):3–23, 1997.
- [IKM⁺13] Yuval Ishai, Eyal Kushilevitz, Sigurd Meldgaard, Claudio Orlandi, and Anat Paskin-Cherniavsky. On the power of correlated randomness in secure computation. In Amit Sahai, editor, *TCC 2013*, volume 7785 of *LNCS*, pages 600–620. Springer, Berlin, Heidelberg, March 2013.
- [KKR25] Sina Kalantarzadeh, Nathan Klein, and Victor Reis. A randomized rounding approach for DAG edge deletion, 2025. APPROX 2025.
- [KPPS17] Sreyash Kenkre, Vinayaka Pandit, Manish Purohit, and Rishi Saket. On the approximability of digraph ordering. *Algorithmica*, 78(4):1182–1205, 2017.
- [KRRS94] Philip Klein, Satish Rao, Monika Rauch, and Sairam Subramanian. Faster shortest-path algorithms for planar graphs. In *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing*, STOC ’94, page 27–37, New York, NY, USA, 1994. Association for Computing Machinery.
- [Kus89] Eyal Kushilevitz. Privacy and communication complexity. In *30th FOCS*, pages 416–421. IEEE Computer Society Press, October / November 1989.
- [KW15] Nathan Klein and Tom Wexler. On the approximability of DAG edge deletion, 2015.
- [LT79] Richard J. Lipton and Robert Endre Tarjan. A separator theorem for planar graphs. *SIAM Journal on Applied Mathematics*, 36(2):177–189, 1979.
- [MORS25] Pierre Meyer, Claudio Orlandi, Lawrence Roy, and Peter Scholl. Silent circuit relinearisation: Sublinear-size (boolean and arithmetic) garbled circuits from DCR. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part IV*, volume 16003 of *LNCS*, pages 426–458. Springer, Cham, August 2025.
- [MS18] Guy Moshkovitz and Asaf Shapira. Decomposing a graph into expanding subgraphs. *Random Structures & Algorithms*, 52(1):158–178, 2018. Theorem and corollary numbering in the present paper follows the arXiv version, <https://arxiv.org/abs/1502.00409>.
- [NOdM10] Jaroslav Nešetřil and Patrice Ossona de Mendez. Sparse Combinatorial Structures: Classification and Applications. In *Proceedings of the International Congress of Mathematicians 2010 (ICM 2010)*, volume IV, pages 2502–2529, Hyderabad, India, August 2010. World Scientific.
- [NOdM12] Jaroslav Nešetřil and Patrice Ossona de Mendez. *Sparsity: Graphs, Structures, and Algorithms*, volume 28 of *Algorithms and Combinatorics*. Springer Heidelberg, 2012.

- [OSY21] Claudio Orlandi, Peter Scholl, and Sophia Yakoubov. The rise of paillier: Homomorphic secret sharing and public-key silent OT. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part I*, volume 12696 of *LNCS*, pages 678–708. Springer, Cham, October 2021.
- [PRS94] D. Paik, S. Reddy, and S. Sahni. Deleting vertices to bound path length. *IEEE Transactions on Computers*, 43(9):1091–1096, 1994.
- [PV16] Marie Paindavoine and Bastien Vialla. Minimizing the number of bootstrappings in fully homomorphic encryption. In Orr Dunkelman and Liam Keliher, editors, *SAC 2015*, volume 9566 of *LNCS*, pages 25–43. Springer, Cham, August 2016.
- [RB89] Tal Rabin and Michael Ben-Or. Verifiable secret sharing and multiparty protocols with honest majority (extended abstract). In *21st ACM STOC*, pages 73–85. ACM Press, May 1989.
- [RS86] Neil Robertson and P.D Seymour. Graph minors. ii. algorithmic aspects of tree-width. *Journal of Algorithms*, 7(3):309–322, 1986.
- [RS21] Lawrence Roy and Jaspal Singh. Large message homomorphic secret sharing from DCR and applications. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part III*, volume 12827 of *LNCS*, pages 687–717, Virtual Event, August 2021. Springer, Cham.
- [Sve12] Ola Svensson. Hardness of vertex deletion and project scheduling. In *Approximation, Randomization, and Combinatorial Optimization (APPROX/RANDOM)*, volume 7408 of *Lecture Notes in Computer Science*, pages 301–312. Springer, 2012.
- [Tu22] Jianhua Tu. A survey on the k -path vertex cover problem. *Axioms*, 11(5):191, 2022.
- [Yao86] Andrew Chi-Chih Yao. How to generate and exchange secrets (extended abstract). In *27th FOCS*, pages 162–167. IEEE Computer Society Press, October 1986.