Can You Link Up With Treewidth?

Radu Curticapean
University of Regensburg and IT University of Copenhagen

Simon Döring
Max Planck Institute for Informatics and Saarland University

Daniel Neuen
University of Regensburg and Max Planck Institute for Informatics

Jiaheng Wang
University of Regensburg

Abstract

A central result of Marx [ToC '10] proves that there are k-vertex graphs H of maximum degree 3 such that $n^{o(k/\log k)}$ time algorithms for detecting colorful H-subgraphs would refute the Exponential-Time Hypothesis (ETH). This result is widely used to obtain almost-tight conditional lower bounds for parameterized problems under ETH.

Our first contribution is a new and fully self-contained proof of this result that further simplifies a recent work by Karthik et al. [SOSA 2024]. Towards this end, we introduce a novel graph parameter, the linkage capacity $\gamma(H)$, and show with an elementary proof that detecting colorful H-subgraphs in time $n^{o(\gamma(H))}$ refutes ETH. Then, we use a simple construction of communication networks credited to Beneš to obtain k-vertex graphs of maximum degree 3 and linkage capacity $\Omega(k/\log k)$, avoiding the use of expander graphs. We also show that every graph H of treewidth t has linkage capacity $\Omega(t/\log t)$, thus recovering the stronger result of Marx [ToC '10] with a simplified proof.

Additionally, we obtain new tight lower bounds for certain types of patterns by analyzing their linkage capacity. For example, we prove that almost all k-vertex graphs of polynomial average degree $\Omega(k^{\beta})$ for some $\beta > 0$ have linkage capacity $\Theta(k)$, which implies tight lower bounds for such patterns H. As an application of these results, we also obtain tight lower bounds for counting small induced subgraphs having a certain property Φ, improving bounds from [Roth et al., FOCS 2020].

The research is funded by the European Union (ERC, CountHom, 101077083). Views and opinions expressed are those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

The title was found with the help of a popular LLM. We thank Cornelius Brand for pointing out a connection to extension complexity.

1 Introduction

Over the past two decades, it has been discovered that complexity assumptions about *exponential-time* problems imply far-reaching lower bounds for *polynomial-time* [13, 67, 68] and *parameterized* [26, 56] problems. Among the first such results, it was shown that the Exponential-Time Hypothesis (ETH) about the Boolean satisfiability problem implies an $n^{\Omega(k)}$-time lower bound for the seemingly unrelated parameterized problem CLIQUE of detecting k-cliques in n-vertex graphs [16, 17]. This lower bound solidifies the status of CLIQUE as a canonical hard problem in parameterized complexity.

Ideally, a reduction from CLIQUE to some target problem would also transfer the $n^{\Omega(k)}$-time lower bound under ETH from CLIQUE to the target problem. However, reductions from CLIQUE often require k gadgets that encode the vertices of a k-clique, as well as $\Theta(k^2)$ additional gadgets to verify the edges between all pairs of encoded vertices. As each gadget typically increases the parameter by $O(1)$, an instance for CLIQUE is then transformed into an instance of the target problem with a parameter value of $\Theta(k^2)$ (see, e.g., [26, Section 13.6.3]). This in turn means that only $n^{o(\sqrt{\ell})}$-time algorithms can be ruled out under ETH for a target problem with parameter ℓ.

Tighter lower bounds could be obtained if we could reduce from a similar subgraph problem, but for k-vertex patterns H with only $O(k)$ rather than $\Theta(k^2)$ edges. More specifically, for a fixed graph H, let COLSUB(H) be the problem of detecting H-subgraph copies in graphs G with vertex-colors from $V(H)$ such that every $v \in V(H)$ is mapped into color v in G. (This problem can equivalently be interpreted as a *constraint satisfaction problem* with variables x_v for $v \in V(H)$ and arity-2 relations R_e for $e \in E(H)$. The domain of x_v is the set of v-colored vertices in G.) Known parameterized hardness results based on the problem CLIQUE can often be modified to use COLSUB(H) as the reduction source. This is useful, because a seminal result by Marx [59, Corollary 6.1] (very recently shown with a simpler proof [50, Theorem 1.3]) shows that COLSUB(H) is indeed hard under ETH for graphs H of maximum degree 3, albeit not with an entirely tight lower bound:

Theorem 1.1 ([50, 59])**.** *Assuming ETH, there exists a fixed constant $\alpha > 0$ and an infinite sequence of graphs $H_1, H_2, \ldots$ such that, for all $k \in \mathbb{N}$, the graph H_k has k vertices and maximum degree 3, and* COLSUB(H) *does not admit an $O(n^{\alpha \cdot k/\log k})$-time algorithm.*

This theorem has become a standard tool to prove almost-tight lower bounds along the lines of the above reduction scheme, and it has been applied to numerous parameterized problems from a diverse range of areas [1, 5, 8, 9, 10, 11, 12, 14, 18, 19, 20, 21, 22, 23, 25, 27, 34, 35, 36, 39, 42, 45, 49, 52, 57, 60, 62, 64].

1.1 Main Concept: Linkage Capacity

In this paper, we provide a new perspective on the seminal Theorem 1.1, which allows us to simplify its proof significantly (even its more recent version [50]) and derive several new results. Our interpretation hinges upon a new graph parameter, the *linkage capacity* $\gamma(H)$ of a graph H. Roughly speaking, this parameter measures how well vertices of H can be connected by vertex-disjoint paths on specified endpoint pairs.

Known Lower Bound for the Clique Problem To explain our ideas, let us first sketch the classical $n^{\Omega(k)}$-time lower bound for CLIQUE under ETH (see, e.g., [26, Theorem 14.21]) and then describe our modifications. The original proof is as follows: It is known that, assuming ETH, the 3-COLORING problem cannot be solved in $2^{o(n)}$ time for n-vertex graphs G with maximum degree 4. If G can be transformed into an equivalent instance X of CLIQUE with approximately $3^{n/k}$ vertices, then an $n^{o(k)}$-time algorithm for CLIQUE would imply a $2^{o(n)}$-time algorithm for the 3-COLORING problem, contradicting ETH.

To transform G into X, the vertex set $V(G)$ is divided equitably into blocks $V_1, \dots, V_k$. The vertices of X correspond to the 3-colorings of these blocks, and two vertices in X are connected by an edge if their colorings are compatible, meaning they come from different blocks and together form a proper coloring. This way, the k-cliques K in this "compatiblity graph" X correspond bijectively to valid 3-colorings of G: Indeed, the vertices of K provide a valid coloring for each block, and the presence of edges between all $u, v \in V(K)$ in X ensures that the union of these partial colorings is a valid coloring of the entire graph G.

From Cliques to General Subgraphs To show hardness of ColSub(H) with general k-vertex patterns H, we adapt the lower bound for Clique. First, consider the favorable scenario that the vertices of an input graph G for 3-Coloring can be split equitably into blocks $V_1, \dots, V_k$, corresponding to the k vertices of H, such that the edges of G "respect" H: Every edge of G is contained within one block or between blocks V_i and V_j with $ij \in E(H)$. In this scenario, not all pairs of partial 3-colorings need to be checked for compatibility. Indeed, it suffices to check this between blocks V_i and V_j with $ij \in E(H)$, since no other edges could lead to an incompatibility.

In general however, we cannot assume that an n-vertex graph G of maximum degree 4 can be split equitably such that its edges respect H. To address this, we "re-route" the edges in G along paths on new vertices (that are placed in the old blocks) and edges that *do* respect H. While this eventually yields a graph G' in which all edges indeed respect H, it may be possible that most edges are routed on paths of length $\Omega(k)$, thus increasing the block size from n/k back to n. Even if routing via short paths is possible, it may be possible that a few blocks are hit disproportionally often, leading to the same problem. Both issues would render a fast algorithm for ColSub(H) useless for the purpose of obtaining a (too) fast algorithm for 3-Coloring.

Linkage Capacity The crucial observation is that many patterns H enable a simultaneous "batch-rerouting" of batches with $\Omega(k)$ edges in G; adding all paths for any such a batch to G increases each block size only by 1. Moreover, as also observed in [50, Theorem 4.2], it is sufficient to consider batches that are matchings, since G has maximum degree 4 and thus admits a 5-edge-coloring, i.e., a partition of its edges into 5 matchings.

The *linkage capacity* $\gamma(H)$ allows us to precisely quantify how well H supports batch-rerouting of matchings by vertex-disjoint paths. To define it, first let the *blowup* $H \otimes J_t$ for $t \in \mathbb{N}$ be H with every vertex copied to t clones that form a clique; this is essentially the maximal graph with block size t whose edges respect H. See also Figure 1. Second, call a set X in a graph H' *matching-linked* if, for every matching M with vertex-set X, there exist disjoint u-v-paths in H' realizing the edges $uv \in M$. Then the linkage capacity $\gamma(H)$ of a graph H is the largest $c > 0$ such that $H \otimes J_t$ contains a matching-linked set X of size $\lfloor ct \rfloor$; this is finite, and we even have $\gamma(H) \le k$, as $|X| \le |V(H \otimes J_t)| = kt$.

Following the reduction sketch from 3-Coloring given above, and using large matching-linked sets in blowups $H \otimes J_t$ to accommodate the vertices of a 3-Coloring instance G, we establish a conditional lower bound on the complexity of ColSub(H) based on $\gamma(H)$.

Theorem 1.2. *Assuming ETH, there exist fixed constants $\alpha, \gamma_0 > 0$ such that no fixed graph H with $\gamma(H) \ge \gamma_0$ admits an $O(n^{\alpha \cdot \gamma(H)})$-time algorithm for* ColSub(H).

It remains to determine when H has large linkage capacity. For example, if H itself admits a large matching-linked set, then this translates to its blowups, thus establishing high $\gamma(H)$. This is however only a sufficient criterion, even though most of our lower bounds are based on it. As we investigate in Section 6, the linkage capacity is related to certain fractional multicommodity flow problems whose relevance in the context of lower bounds for ColSub(H) under ETH was already identified before [50, 59]. Linkage capacity however is a much more elementary and more applicable concept. In particular, the restriction to matchings allows us to connect it to known results on routing with specified terminal pairs in order to obtain lower bounds on $\gamma(H)$. This in turn allows us to prove new results under ETH without much technical effort.

1.2 Applications of Linkage Capacity

With Theorem 1.2 in hand, we show lower bounds on the complexity of the colorful H-subgraph problem via the linkage capacity $\gamma(H)$. For this, we enlist the help of communication network theory [54, 6], random graph theory [15], linear programming [37, 55], and classical results on connectivity via vertex-disjoint paths from graph theory [58, 66].

A Fully Self-Contained Proof of Theorem 1.1 Our first application of Theorem 1.2 is a significantly simplified and self-contained[1] proof of the seminal Theorem 1.1. The original proof of this theorem by Marx [59] uses highly nontrivial arguments regarding multicommodity flows as a black box [37]. Even a very recent simplification [50] still requires the construction of expander graphs and routing algorithms for such graphs, both of which are highly nontrivial [2, 55].

By approaching the problem through linkage capacity, we observe that expansion is *not* required to obtain Theorem 1.1. Instead, we can rely on a very simple construction of telecommunication networks, credited to a 1964 paper by Beneš [6], then employed at Bell Labs: A *Beneš network* contains $s = 2^\ell$ input and output vertices, and $k = O(s \log s)$ vertices in total. For every pairing of inputs to outputs, the network guarantees private data streams (i.e., vertex-disjoint paths) connecting each input to its specified output. Both the network construction and routing therein are elementary divide-and-conquer arguments that feature in undergraduate introduction courses to discrete mathematics [54]. A minuscule augmentation of this construction gives us k-vertex graphs of maximum degree 4 and linkage capacity $\Omega(k/\log k)$. Combined with Theorem 1.2, this gives a novel proof of Theorem 1.1.

We recently found that graphs with large matching-linked sets have been used in communication and extension complexity: A paper by Göös, Jain, and Watson [41, Section 3.3] briefly mentions "bounded-degree butterfly graphs" from an unpublished manuscript on pebble games by Nordström [61, Proposition 5.2] as an alternative to expanders; this alternative construction turns out to be precisely that of Beneš.

Tight Lower Bounds for Dense Graphs Alon and Marx [4, Theorem 1.4] argue that the logarithmic slack in Theorem 1.1 cannot be overcome by current approaches—including ours. This holds even for patterns H of constant *average* rather than *maximum* degree. More modestly, one can ask for "just slightly" dense k-vertex patterns H such that ColSub(H) requires $n^{\Omega(k)}$ time under ETH.

Indeed, Alon and Marx [4, Theorem 1.5(2)] showed that, for every $\delta > 0$, certain specifically constructed patterns S with average degree $O(k^\delta)$ enjoy strong embeddability properties that entail $n^{\Omega(k)}$-time lower bounds on the colorful S-subgraph problem [4, Theorem 1.8]. For some problems of interest however, e.g., for counting induced k-vertex patterns [24, 32, 64], one can only reduce from the colorful H-subgraph problem for *some* (say, adversarially chosen) dense pattern H, which may not necessarily be a graph S constructed by Alon and Marx. This imposes a bottleneck towards tight lower bounds for such problems.

One partial remedy lies in using large clique minors (see, e.g., [64]). Kostochka [53] showed that *every* graph H of average degree d contains a K_q-minor with $q = \Omega(d/\sqrt{\log d})$. Given a K_q-minor in H, a straightforward reduction yields an $n^{\Omega(q)}$-time lower bound on the colorful H-subgraph problem under ETH. This implies that *every* pattern H of linear average degree $\Omega(k)$

[1] We give a self-contained proof starting from the known result that, under ETH, the 3-Coloring problem requires $2^{\Omega(n)}$ time on 4-regular graphs with n vertices. This can be shown easily from ETH together with the sparsification lemma.

requires an exponent of $\Omega(k/\sqrt{\log k})$ for the colorful H-subgraph problem. While this improves upon the lower bound from Theorem 1.1, a slack of $\Omega(\sqrt{\log k})$ remains.

Using linkage capacity, we eliminate this slack and obtain a tight lower bound for dense patterns: Combining two textbook results [30], we show that every pattern H of average degree d has linkage capacity $\Omega(d)$.[2] Theorem 1.2 then immediately yields:

Theorem 1.3. *There is a constant $\alpha > 0$ such that, for every graph H with average degree d, the existence of an $O(n^{\alpha \cdot d})$ time algorithm for* ColSub(H) *would refute ETH.*

This theorem covers the "worst case", i.e., patterns H of fixed average degree d that are adversarially chosen so as to minimize $\gamma(H)$. In particular, for linear average degree, an $n^{\Omega(k)}$ bound under ETH follows. This implies new tight lower bounds for very general classes of induced pattern counting problems [24, 64] (see Section 7 for details).

In the "average case", much lower density turns out to be sufficient for an $n^{\Omega(k)}$ bound. Indeed, known results on routing in random graphs [15] imply directly that *almost all* k-vertex graphs H with average degree $d \in \Omega(k^{\beta})$ for constant $\beta > 0$ have linkage capacity $\Theta(k)$. Observe that the average degree is that of the specifically constructed patterns S by Alon and Marx [4]; we show that not only specific patterns, but *almost all* patterns of polynomial average degree have an $n^{\Omega(k)}$ bound for ColSub(H).

More generally, we show that the linkage capacity of the Erdős-Rényi random graph $\mathcal{G}(k,p)$ for non-degenerate probabilities p is $\Omega(k/\rho)$, where $\rho = \log(k)/\log(kp)$ is the typical distance between vertices in $\mathcal{G}(k,p)$ [7, 51]. We obtain the following general lower bound:

Theorem 1.4. *There is a constant $\alpha > 0$ such that for every constant $\varepsilon > 0$ and every $p \geq (1+\epsilon)\log k/k$, the following holds: With high probability, an Erdős-Rényi random graph $H \sim \mathcal{G}(k,p)$ is such that an $O(n^{\alpha \cdot k/\rho})$ time algorithm for* ColSub(H) *would refute ETH. Here, we write $\rho = \log(k)/\log(kp)$ for the typical distance in $G(k,p)$.*

Note that ρ is the logarithm of k in the base of the average degree kp; this captures the time needed to concurrently explore all k vertices in a process that branches into kp random vertices from each vertex. It is intuitively clear that the linkage capacity should be at most $O(k/\rho)$: Almost all vertex pairs u, v in a random graph require u-v-paths of length ρ, so we cannot connect more than k/ρ vertex pairs without exhausting k vertices. The bound from [15] shows that, with high probability, $\Omega(k/\rho)$ vertex pairs *can* be connected.

Linkage Capacity and Treewidth Besides studying the interplay between the linkage capacity $\gamma(H)$ and $|V(H)|$, we also consider its connection to the treewidth $\mathrm{tw}(H)$ of a graph H. For example, it is easily shown that the ℓ-by-ℓ grid $\boxplus_\ell$ has treewidth ℓ and linkage capacity $\Omega(\ell)$, see Lemma 3.9 and Figure 1. Thus, we have the linear relationship $\gamma(\boxplus) \in \Omega(\mathrm{tw}(\boxplus))$ for every grid $\boxplus$. Standard bidimensionality arguments [29] imply such a linear relationship for every class of graphs $\mathcal{H}$ excluding a fixed minor, thus yielding tight lower bounds under ETH parameterized by the *treewidth* $\mathrm{tw}(H)$ rather than $|V(H)|$. However, we stress that, since patterns $H \in \mathcal{H}$ excluding a fixed minor satisfy $|V(H)| \in \Omega(\mathrm{tw}(H)^2)$, such patterns are not amenable for results similar to Theorem 1.1.

In Section 6, we drop the assumption of excluding a fixed minor and establish the bound of $\gamma(H) = \Omega(t/\log t)$ for general graphs H of treewidth t. This recovers a general lower bound by Marx [59] with a more transparent proof—albeit the same approximate min-cut/max-flow theorem for multicommodity flows [37, 55] still appears as black box.

[2] This lower bound is asymptotically tight, as worst-case examples like $K_{d,s-d}$ have linkage capacity at most d. Indeed, a linkage with $d+1$ paths in $K_{d,s-d}$ would in particular imply a matching with $d+1$ edges, which clearly does not exist in $K_{d,s-d}$.

2 Preliminaries

We write $\mathbb{N} = \{1, 2, 3, \dots\}$ for the natural numbers. For $n \in \mathbb{N}$, we write $[n] := \{1, 2, \dots, n\}$. All logarithms are natural unless specified otherwise.

2.1 Basic Definitions

We use standard graph notation [30]. Graphs are finite and undirected, and we write uv for edges between u and v. A *path* from u to v is a sequence $P = (u = w_0, w_1, \dots, w_\ell = v)$ of distinct vertices such that consecutive vertices are adjacent. Slightly abusing notation, we also interpret P as a path from v to u. For a graph G and $X \subseteq V(G)$, we write $G[X]$ for the subgraph induced by X and $G - X \coloneqq G[V(G) \setminus X]$ for the result of deleting X from G.

A *colored graph* is a triple $G = (V, E, c)$ where $c\colon V(G) \to C$ is a not necessarily proper *(vertex-)coloring* of the vertices. We say G is *canonically colored* if c is the identity mapping and we write G^{can} for the canonically colored version of G.

Given a "pattern" graph H and "host" graph G, we write $\#\mathrm{Sub}(H \to G)$ for the number of subgraphs of G that are isomorphic to H. If H and G are colored, only subgraphs preserving the coloring are counted. For a fixed graph H, the problem #ColSub(H) takes as input a colored graph $G = (V, E, c)$ with $c\colon V(G) \to V(H)$, and asks to compute $\#\mathrm{Sub}(H^{\mathsf{can}} \to G)$, while its decision version ColSub(H) asks whether $\#\mathrm{Sub}(H^{\mathsf{can}} \to G) \geq 1$.

To analyze the complexity of ColSub(H), we rely on several tools.

Definition 2.1 (Blowup)**.** Given a graph H and an integer $t \geq 1$, the *blowup graph* $H \otimes J_t$ contains the vertices $v^{(i)}$ for all $v \in V(H)$ and $i \in [t]$, and edges

$$\{u^{(i)}v^{(j)} \mid uv \in E(H),\, i, j \in [t]\} \;\cup\; \{u^{(i)}u^{(j)} \mid u \in V(H),\, i \neq j \in [t]\}.$$

Remark 2.2. Marx [59] uses the notation $H^{(t)}$ instead of $H \otimes J_t$. We choose $H \otimes J_t$ since there is no exponential increase in size, but we are rather taking a tensor product of H and the $(t \times t)$ all-ones matrix (usually denoted by J_t) and then turn cloned vertices into cliques.

A *multigraph* M is a graph that allows parallel edges with the same endpoints, but no self-loops. The *degree* $\deg_M(v)$ of a vertex $v \in V(M)$ is the number of edges incident to v, taking multiplicities into account. The *average degree* of M is $d(M) \coloneqq 2|E(M)|/|V(M)|$.

A *matching* in M is set $M' \subseteq E(M)$ of pairwise vertex-disjoint edges. Slightly abusing notation, we regularly interpret M' again as a graph with edge set M' and vertices for all endpoints in M'. Given a vertex set X, a matching *on* X is a matching whose endpoints are all contained in X. In particular, given a multigraph M, a matching on some vertex set $X \subseteq V(M)$ need not be a matching in M.

A *q-edge coloring* of M is a partition of $E(M)$ into q matchings. The *edge-chromatic number* of M, denoted by $\chi'(M)$, is the minimum number q such that M admits a q-edge coloring. A theorem by Shannon [65] provides an upper bound on the edge-chromatic number in terms of the maximum degree, though a looser bound with factor 2 can be achieved by a straightforward greedy algorithm.

Theorem 2.3 ([65])**.** *Every multigraph M of maximum degree d has $\chi'(M) \leq \lfloor \frac{3}{2} d \rfloor$.*

2.2 Linkages

Our hardness proofs for ColSub(H) crucially rely on linkages in graphs.

Definition 2.4 (Linkage and congestion)**.** Given a graph H and a multigraph M with vertex set $X \subseteq V(H)$, an *M-linkage* in H is a collection of paths $Q = (P_{uv})_{uv \in E(M)}$ such that P_{uv} has endpoints u and v. For $r \in \mathbb{N}$, we say that Q is *r-congested* if, for all $w \in V(H)$, at most r paths $P_{uv} \in Q$ contain w. If $r = 1$, we call Q an *uncongested* M-linkage.

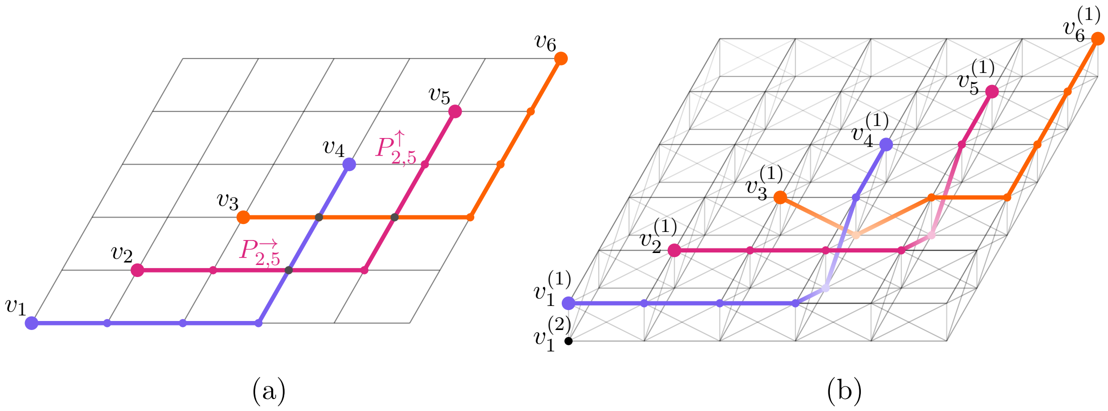

Figure 1: (a) The grid graph $\boxplus_6$. Thick paths depict a 2-congested M-linkage, where $M = \{v_1v_4, v_2v_5, v_3v_6\}$ is a matching on the diagonal vertices. (b) The blowup graph $\boxplus_6 \otimes J_2$, and an uncongested M-linkage obtained from the 2-congested M-linkage in $\boxplus_6$.

Observe that, if Q is an uncongested M-linkage, then M is necessarily a matching. We note that we commonly work with uncongested M-linkages. More precisely, we usually require uncongested M-linkages in blowups of graphs H. Towards this end, it is often convenient to "project" M back to the base graph H. Let H be a graph and let M be a multigraph with $V(M) \subseteq V(H \otimes J_q)$. We define the *$H$-projection* of M to be the multigraph $\pi(M)$ with vertex set $V(\pi(M)) = \{v \mid v^{(i)} \in V(M)\}$ and edge set $E(\pi(M)) = \{\{vw \mid v^{(i)}w^{(j)} \in E(M), v \neq w\}\}$.

Lemma 2.5. *Let H be a graph and let $q \in \mathbb{N}$. Also let M be a matching with $V(M) \subseteq V(H \otimes J_q)$. If there is a q-congested $\pi(M)$-linkage in H, then there is an uncongested M-linkage in $H \otimes J_{2q}$.*

Note that a blowup of order $2q$ rather than q is needed in Lemma 2.5. This is needed since we do not allow self-loops in the projection, i.e., the projection of M ignores edges contained in the same block $\{v^{(i)} \mid i \in [q]\}$. For technical reasons, we decide to handle those edges separately at the cost of losing a factor of two.

Proof. Let Q be a q-congested $\pi(M)$-linkage in H. We obtain an uncongested M-linkage Q' in $H \otimes J_{2q}$ as follows. For a vertex $w \in V(H)$, let $P_1, \ldots, P_t$ (if any exists) be all paths in Q that contain w as an internal vertex. We have $t \leq q$ by definition. We replace w in P_i with the vertex $w^{(q+i)}$ from the blowup graph $H \otimes J_{2q}$. Also, all endpoints of the paths are replaced in the natural way, i.e., if P has endpoints u and v, and uv is the "projection" of $u^{(i)}v^{(j)}$ in M, then P gets endpoints $u^{(i)}$ and $v^{(j)}$. Finally, for each edge $v^{(i)}v^{(j)} \in E(M)$, we add the path $(v^{(i)}, v^{(j)})$ to Q'. By the definition of the blowup graph, the resulting collection Q' is an uncongested M-linkage in $H \otimes J_{2q}$. □

The following example illustrates the notion of linkages and the interplay between congestion and blowups; see Figure 1.

Example 2.6 (Grid graph)**.** Write $\boxplus_\ell$ for the grid graph on vertex set $[\ell] \times [\ell]$. For every matching M on the set of diagonal vertices $v_i = (i, i)$ for $i \in [\ell]$, we observe that $\boxplus_\ell$ contains a 2-congested M-linkage $Q(M) = \{P_{uv}\}_{uv \in M}$. This 2-congested linkage induces an uncongested M-linkage in $\boxplus_\ell \otimes J_2$ via Lemma 2.5. See also Figure 1.

More specifically, given an edge from $u = (a, a)$ to $v = (b, b)$ for $a < b$, we define P_{uv} as the concatenation of the path $P^{\rightarrow}_{uv}$ on vertices $u = (a, a), \ldots, (b, a)$ and the path $P^{\uparrow}_{uv}$ on vertices $(b, a), \ldots, (b, b) = v$. The paths $P^{\rightarrow}_{uv}$ for $uv \in M$ are vertex-disjoint (as distict paths have distinct y-coordinates), and so are the paths $P^{\uparrow}_{uv}$ for $uv \in M$ (having distinct x-coordinates), so $Q(M)$ is indeed a 2-congested M-linkage.

3 Lower Bounds from Linkage Capacity

The Exponential-Time Hypothesis ETH [43] postulates the existence of a constant $\alpha > 0$ such that no $O(2^{\alpha n})$ time algorithm can decide, on input a 3-CNF formulas φ with n variables, whether φ admits a satisfying assignment. Its *a priori* weaker counting version #ETH postulates the same lower bound for *counting* the satisfying assignments of φ [28]. For both hypotheses, the sparsification lemma [44, 28] rules out such algorithms even under the additional condition that every variable in φ appears in at most C clauses, for some constant $C \in \mathbb{N}$. By a standard reduction, lower bounds follow for the problem 3-Coloring of deciding whether an input graph G admits a proper vertex-coloring with 3 colors where no adjacent vertices receive the same color; see for example [56, Theorem 3.2].

Theorem 3.1. *Assuming ETH, there is a constant $\beta > 0$ such that* 3-Coloring *cannot be solved in time $O(2^{\beta \cdot n})$ for n-vertex input graphs G of maximum degree* 4*. The same holds for* #3-Coloring *under the counting variant of ETH.*

This theorem is the foundation for the lower bounds shown in this paper.

3.1 Instances That Fit into Blowups

It will prove useful for us to generalize 3-Coloring slightly, by allowing edges to either enforce equality or disequality of their endpoint colors. Since "equality edges" can be contracted without changing the number of valid assignments, we obtain an immediate way to simulate edges in a 3-Coloring instance by paths.

Definition 3.2. Given a graph $G = (V, E)$ with a partition $E = E_{=} \cup E_{\neq}$, a *proper* 3*-assignment* is a function $a : V \to [3]$ such that $a(u) = a(v)$ for all $uv \in E_{=}$, while $a(u) \neq a(v)$ for all $uv \in E_{\neq}$. The problem 3-Assignment asks to determine the existence of a proper assignment on input $(G, E_{=}, E_{\neq})$, while #3-Assignment asks to count them.

It is possible to convert instances G for 3-Assignment into instances X for ColSub(H). Moreover, if G fits into a moderately small blowup of H, then X is only of moderately exponential size. This can be shown by a simple "split-and-list" reduction that follows the sketch given in the introduction for Clique.

Lemma 3.3. *Let H be a fixed k-vertex graph with canonical vertex-coloring. Given a subgraph G of $H \circledast J_t$ for $t \in \mathbb{N}$, a colored graph X on $k \cdot 3^t$ vertices can be computed in $9^t \cdot \mathrm{poly}(k, t)$ time such that $\#\mathrm{Sub}(H \to X)$ equals the number of proper* 3*-assignments in G.*

Proof of Lemma 3.3. Suppose $V(H) = [k]$ and consider the partition of $V(G)$ into $V_w = \{w^{(1)}, \ldots, w^{(t)}\}$ for $w \in [k]$. Define X_w for $w \in [k]$ as the set of all proper 3-assignments to $G[V_w]$. For $w, w' \in [k]$, we call two 3-assignments $a \in X_w$ and $a' \in X_{w'}$ *compatible* if their union is a proper 3-assignment of $G[V_w \cup V_{w'}]$. Let us define

$$A_X := \{(a_1, \ldots, a_k) \in X_1 \times \ldots \times X_k \mid \forall ww' \in E(H) : a_w \text{ and } a_{w'} \text{ are compatible}\}$$
$$A_G := \{a \colon V(G) \to [3] \mid a \text{ is proper 3-assignment of } G\}$$

We observe that the map $a \mapsto (a_1, \ldots, a_k)$ from A_G to A_X, where a_w is the restriction of a to V_w, is a bijection. Indeed, in the image of $a \in A_G$ under this map, a_w and $a_{w'}$ are compatible for all $w, w' \in E(H)$. Conversely, given $(a_1, \ldots, a_k) \in A_X$, recall that every edge $uv \in E(G)$ satisfies $u \in V_w$ and $v \in V_{w'}$ for some $w, w' \in [k]$ with (a) $w = w'$, or (b) $ww' \in E(H)$. In case (a), since a_w is proper, the endpoints of uv receive a proper assignment under a. In case (b), because the union of a_w and $a_{w'}$ is a proper 3-assignment of $G[V_w \cup V_{w'}]$, the endpoints of uv receive a proper assignment under a. Thus $a \in A_G$.

Finally, the graph X is defined on vertices $\bigcup_{w\in[k]} X_w$, where each vertex in X_w is colored by $w \in [k]$. An edge is present between $a \in X_w$ and $b \in X_{w'}$ if and only if $ww' \in E(H)$ and a and b are compatible. The (colored) subgraphs S of X isomorphic to H correspond to tuples in A_X. Indeed, $V(S)$ corresponds to a tuple $(a_1, \ldots, a_k) \in X_1 \times \ldots \times X_k$, and the presence of edges of H in S implies that a_w and $a_{w'}$ are compatible for $ww' \in E(H)$. Since $|X_w| \leq 3^t$ for all $w \in [k]$, the graph X can be computed by brute-force in $9^t \cdot \mathrm{poly}(k,t)$ time. □

3.2 The Linkage Capacity of a Graph

First, we need to define a term for vertex sets X in graphs that can be paired up arbitrarily via paths in H. This resembles Diestel's [30] notion of *linkedness*, see Section 5.1, which however requires this property for the entire graph H.

Definition 3.4 (Matching-linked set)**.** Given a graph H, we say that $X \subseteq V(H)$ is *matching-linked* if H contains an uncongested M-linkage for every matching M on vertex set X.

Remark 3.5. We stress that the condition in the definition is crucially required even if M is *not* contained in $E(H)$: Only the endpoints of M need to be contained in $V(H)$.

A simple edge-coloring argument, also used in Lemma 3.11, shows that large matching-linked sets X in blowups $H \otimes J_t$ suffice to embed graphs G of maximum degree Δ into $H \otimes J_{2\Delta\cdot t}$. Thus, large matching-linked sets X *in blowups of* H are a useful "resource" attainable from H that allows us to use Lemma 3.3. Not all such sets X however need to originate from matching-linked sets in H itself. Consider a set X in H that *just* fails to be matching-linked, as in Example 2.6, in the sense that X still admits M-linkages of congestion 2 in H. Such M-linkages then induce uncongested M-linkages in $H \otimes J_2$. As our goal is to embed a 3-Coloring instance G into a moderately large blowup of H, such a constant-factor loss would be acceptable. This flexibility is captured by the *linkage capacity*, which measures the maximum size of matching-linked sets in blowups of H relative to the blowup order.

Definition 3.6 (Linkage capacity)**.** The *linkage capacity* $\gamma(H)$ is the supremum over $c > 0$ such that $H \otimes J_t$ contains a matching-linked set X with $|X| = \lfloor ct \rfloor$ for all large enough $t \in \mathbb{N}$.

Every graph H trivially satisfies $1 \leq \gamma(H) \leq |V(H)|$. We show below that large matching-linked sets in H lift into blowups, establishing high linkage capacity $\gamma(H)$—but as mentioned above, even a matching-linked set in a small blowup of H suffices.

Lemma 3.7. *Let H be a graph and suppose $H \otimes J_q$ for $q \in \mathbb{N}$ contains a matching-linked set X. Then $\gamma(H) \geq \frac{1}{3} \cdot |X|/q$.*

Proof. The proof rests on the following claim.

Claim 3.8. *Let H' be a graph and suppose $X' \subseteq V(H')$ is a matching-linked set. Then $X'_t := \{v^{(i)} \mid v \in X', 1 \leq i \leq t/3\}$ is matching-linked in $H' \otimes J_t$ for every $t \in \mathbb{N}$.*

Proof. Let M be a matching on X'_t and let $M' := \pi(M)$ be the H'-projection of M. Observe that M' has maximum degree at most $t/3$, so $\chi'(M') \leq t/2$ by Theorem 2.3. Hence, the multigraph M can be partitioned into $r \leq t/2$ matchings $M_1, \ldots M_r$ on X'. Since X' is a matching-linked set, for every M_i, there is an uncongested M_i-linkage Q_i in H'. So $Q := \bigcup_{i\in[r]} Q_i$ is a r-congested M'-linkage in H', and there is an uncongested M-linkage in $H' \otimes J_t$ by Lemma 2.5. ⌟

Let $c < \frac{1}{3} \cdot |X|/q$. Then there is $t_0 \in \mathbb{N}$ and $c' < \frac{1}{3} \cdot |X|$ such that

$$c' \cdot \lfloor t/q \rfloor \geq c \cdot t \tag{1}$$

for all $t \geq t_0$. Also, there is some $\ell_0 \in \mathbb{N}$ such that

$$|X| \cdot \lfloor \ell/3 \rfloor \geq \lfloor c'\ell \rfloor \tag{2}$$

for all $\ell \geq \ell_0$. Now let $t \geq \max(t_0, \ell_0 \cdot q)$ and let $\ell := \lfloor t/q \rfloor \geq \ell_0$ and consider the graph $H' := H \otimes J_q$. By Claim 3.8 the graph $H' \otimes J_\ell$ contains a matching-linked set X'_ℓ with

$$|X'_\ell| \geq |X| \cdot \lfloor \ell/3 \rfloor \overset{(1)}{\geq} \lfloor c'\ell \rfloor \overset{(2)}{\geq} \lfloor ct \rfloor.$$

Since $(H \otimes J_q) \otimes J_\ell$ is a subgraph of $H \otimes J_t$, we conclude that $H \otimes J_t$ contains a matching-linked set of size $\lfloor ct \rfloor$, thus proving the lemma. □

As a concrete example, let us use Example 2.6 to bound the linkage capacity of grids.

Lemma 3.9. *For the ℓ-by-ℓ grid graph $\boxplus_\ell$, we have $\gamma(\boxplus_\ell) \geq (\ell - 1)/6$.*

Proof. Let $V(\boxplus_\ell) = [\ell]^2$ and $\ell' \in \{\ell - 1, \ell\}$ be even. By Example 2.6, the set $X := \{(i,i)^{(1)} \mid i \in [\ell']\}$ is matching-linked in the blowup $\boxplus_\ell \otimes J_2$. The lemma follows by invoking Lemma 3.7. □

3.3 Fitting Instances into Blowups via Linkage Capacity

Having introduced linkage capacity and its key properties, we now use it to embed graphs G into blowups $H \otimes J_t$ with $t = O(n/\gamma(H))$. If we can show that $\gamma(H)$ is large, then ETH implies strong lower bounds for ColSub(H) via Theorem 3.1 and Lemma 3.3.

A minor constructivity issue arises: Some techniques for lower-bounding $\gamma(H)$ do not necessarily yield efficient algorithms for finding linkages in blowups of H. Thus, it could *a priori* be possible for G to embed into $H \otimes J_t$, yet we cannot efficiently find an embedding. This concern is resolved by known algorithms for graphs of bounded neighborhood diversity [40, Theorem 3.7], or through a self-contained argument given in Appendix A:

Theorem 3.10. *Let $f(k) = 3k^{k+2}$. Given a k-vertex graph H and $t \geq 2$ as input, a matching-linked set X of maximum size in $H \otimes J_t$ can be found in $O(t^{f(k)})$ time. Given additionally a matching-linked set X in $H \otimes J_t$ and a matching M with vertex set X, an M-linkage in $H \otimes J_t$ can be found in $O(t^{f(k)})$ time.*

We can now turn to our main lemma. In the following, given graphs G and G' without loops or multi-edges, a G-linkage $Q = (P_{uv})_{uv \in E(G)}$ in G' is a *topological G-minor model* in G' if paths P_{uv} and $P_{u'v'}$ for $uv, u'v' \in E(G)$ in Q intersect only at endpoints. In particular, such intersections can occur only if uv and $u'v'$ share a common vertex. We refer to the subgraph of G' induced by Q as the *image* of Q.

Lemma 3.11. *Let H be a fixed k-vertex graph and let $f(k) = 3k^{k+2}$. Then there is an $O(n^{f(k)})$ time algorithm that, given an instance G for* 3-Coloring *with n vertices and maximum degree 4, outputs an instance for* 3-Assignment *with graph G' such that*

1. *G' is the image of a topological G-minor in $H \otimes J_t$ for $t = 8\lceil n/\gamma(H) \rceil$, and*
2. *the proper 3-colorings of G correspond bijectively to the proper 3-assignments of G'.*

Proof. Let $t' = \lceil n/\gamma(H) \cdot 8/7 \rceil$. Definition 3.6 implies that, if t' is large enough then $H \otimes J_{t'}$ contains a matching-linked set X of size n.[3] In $O(n^{f(k)})$ time, Theorem 3.10 finds such a set X. Fix $V(G) = X$ in the following.

The straightforward greedy algorithm yields a 7-edge-coloring $E(G) = M_1 \cup \ldots \cup M_7$ in time $O(n)$. As X is matching-linked, the graph $H \otimes J_{t'}$ contains an uncongested M_i-linkage Q_i for every individual $i \in [7]$. Each linkage can be found $O(n^{f(k)})$ time via Theorem 3.10. These linkages together induce a topological G-minor model in $H \otimes J_{7t'}$ (see Figure 2): Consider $V(H \otimes J_{7t'})$ to be partitioned into 7 layers such that layer $i \in [7]$ contains the vertices $v^{(j)}$ with

[3] If t' is not large enough, then n is bounded by a function of H. We can then compute the number q of 3-colorings of G in constant time and output a dummy instance G' with q proper 3-assignments.

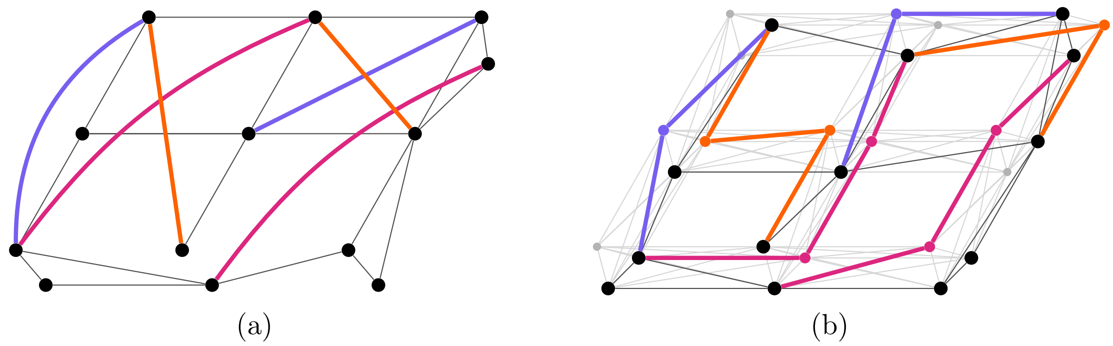

Figure 2: (a) A graph G that fails to be embedded into $\boxplus_3$ due to the colored edges, which are partitioned into three matchings. (b) An embedding of G into the blowup $\boxplus_3 \otimes J_3$ as a topological minor, where each colored edge gets routed via new vertices from the blowup.

$v \in V(H)$ and $j \in (t'-1)i + [t']$. By placing non-endpoint vertices from the linkages $Q_1, \ldots, Q_7$ into different layers and keeping all endpoints in the first layer, we obtain a topological G-minor model Q in $H \otimes J_{7t'}$. Let us write G' for the image of Q.

We finalize the construction of the 3-Assignment instance by specifying a partition of $E(G')$ into $E_=$ and $E_{\neq}$: For each path P_{uv} in Q, place one arbitrary edge into $E_{\neq}$ and all other edges into $E_=$. Then the proper 3-assignments to G' correspond to the proper 3-colorings of G, since contracting all equality edges in G' (which does not change the number of 3-assignments) yields an isomorphic copy of G on disequality edges. □

Combining the above, the proof of Theorem 1.2 is complete.

Proof of Theorem 1.2. By Theorem 3.1, ETH implies a constant $\beta > 0$ such that no $O(2^{\beta \cdot n})$-time algorithm solves 3-Coloring on n-vertex graphs G of maximum degree 4. We set $\gamma_0 = 26/\beta$ and $\alpha = \beta/10$ and derive a contradiction from an $O(s^{\alpha \cdot \gamma})$-time algorithm for ColSub(H) on s-vertex input graphs, where H is any fixed graph with $\gamma = \gamma(H) \geq \gamma_0$.

In the following, let G be an instance for 3-Coloring of maximum degree 4. In time $O(n^{f(k)})$, Lemma 3.11 computes from G an equivalent instance for 3-Assignment with a graph $G' \subseteq H \otimes J_t$ for $t = 8\lceil n/\gamma \rceil$. In time $9^t \cdot \mathrm{poly}(k,t)$, Lemma 3.3 then yields a graph X with $|V(X)| \leq k \cdot 3^t$ such that 3-assignments in G' correspond to colorful H-copies in X. The overall running time to construct X is

$$O(n^{f(k)}) + 9^{8\lceil n/\gamma \rceil} \cdot \mathrm{poly}(k,t) \;=\; O(2^{26n/\gamma}) \;=\; O(2^{\beta \cdot n}). \tag{3}$$

In the last step, we use $\gamma \geq \gamma_0 = 26/\beta$. Then use the assumed $O(s^{\beta/10 \cdot \gamma})$-time algorithm for ColSub(H) on s-vertex input graphs. Its running time on the graph X constructed before, with $s \leq k \cdot 3^t$ vertices, is

$$O((k \cdot 3^t)^{\beta/10 \cdot \gamma}) \;=\; O(3^{5\lceil n/\gamma \rceil \cdot \beta/10 \cdot \gamma}) \;=\; O(3^{\beta/2 \cdot n}) \;=\; O(2^{\beta \cdot n}). \tag{4}$$

Combining Equation (3) and Equation (4), we conclude that both constructing X and solving ColSub(H) on X can be achieved in overall time $O(2^{\beta \cdot n})$. This contradicts Theorem 3.1, also for the counting version. □

4 Switching Networks

In this section, we consider a construction by Beneš [6] that yields k-vertex graphs with degree 4 and a linkage capacity of $\Omega(k/\log k)$. In particular, this allows us to complete the fully self-contained proof of Theorem 1.1.

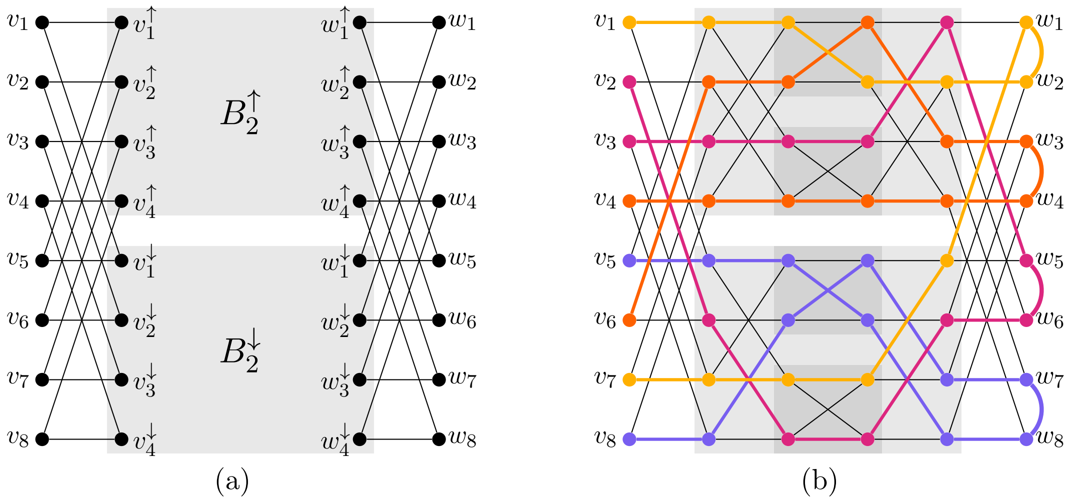

Figure 3: (a) Recursive construction of Beneš network B_3 with 8 inputs and 8 outputs from two copies of B_2. (b) The augmented Beneš network $\breve{B}_3$ is obtained by adding a matching to the outputs of B_3, shown as curved edges. Thick paths indicate an M-linkage in $\breve{B}_3$ for the matching $M = \{v_1v_7, v_2v_3, v_4v_6, v_5v_8\}$ on the input vertices.

The Beneš network B_ℓ for $\ell \in \mathbb{N}$ has 2^ℓ distinguished input and output vertices. In our terms, for every matching M between the inputs and outputs, the network B_ℓ admits an uncongested M-linkage. By "short-circuiting" the outputs, we obtain an *augmented Beneš network* $\breve{B}_\ell$, which allows routing paths from inputs back to inputs. In our terms, the inputs form a matching-linked set, since every matching M on the inputs admits an uncongested M-linkage in $\breve{B}_\ell$.

Algorithm 1 Construct plain Beneš networks

procedure BENES(ℓ) returns B_ℓ with $s = 2^\ell$ in-/outputs
  if $\ell = 1$ **then return** $K_{2,2}$ with in-/outputs v_i, w_i for $i \in [2]$
  else
    $B^\uparrow \leftarrow$ BENES($\ell - 1$) with in-/outputs $v_i^\uparrow, w_i^\uparrow$ for $i \in [s/2]$
    $B^\downarrow \leftarrow$ BENES($\ell - 1$) with in-/outputs $v_i^\downarrow, w_i^\downarrow$ for $i \in [s/2]$
    $B \leftarrow$ vertex-disjoint union of $B^\uparrow$ and $B^\downarrow$
    for $i \in [s/2]$ **do** add to B
      all four edges between $\{v_i, v_{i+s/2}\}$ and $\{v_i^\uparrow, v_i^\downarrow\}$,
      all four edges between $\{w_i, w_{i+s/2}\}$ and $\{w_i^\uparrow, w_i^\downarrow\}$
    return B with in-/outputs v_i, w_i for $i \in [s]$

Definition 4.1 (Beneš networks)**.** The plain *Beneš network* B_ℓ for $\ell \in \mathbb{N}$ is the graph with distinguished inputs v_i and outputs w_i, for $i \in [s]$ with $s = 2^\ell$, returned by BENES(ℓ) in Algorithm 1. The *augmented Beneš network* $\breve{B}_\ell$ is obtained from B_ℓ by adding an edge between outputs w_{2i-1} and w_{2i}, for each $i \in [s/2]$.

Both B_ℓ and $\breve{B}_\ell$ clearly have maximum degree 4. Let $T(s)$ for $s = 2^\ell$ count the vertices in the s-input Beneš network B_ℓ or $\breve{B}_\ell$. By construction, we have $T(s) = 2 \cdot T(s/2) + 2s$, and thus $T(s) = 2s \log_2 s$. Beneš networks are designed to admit uncongested linkages between the inputs and outputs [6]:

Theorem 4.2. *For $\ell \in \mathbb{N}$, the set V of inputs in $\breve{B}_\ell$ is matching-linked, with $|V| = s = 2^\ell$. Moreover, given as input $\ell \in \mathbb{N}$ and a matching M on V, an uncongested M-linkage in $\breve{B}_\ell$ can be computed in $O(s \log s)$ time.*

A proof is given in Appendix B for completeness. With Lemma 3.7, we obtain:

Corollary 4.3. *For $s = 2^\ell$, we have $\gamma(\check{B}_\ell) \geq s/3$.*

By combining Theorem 1.2 and Corollary 4.3, we can give an elementary proof of Theorem 1.1 (in a slightly modified form; see Remark 4.5).

Theorem 4.4. *Assuming ETH, there exists a fixed constant $\alpha > 0$ and an infinite sequence of graphs $H_1, H_2, \ldots$ such that, for all $k \in \mathbb{N}$, the graph H_k has k vertices and maximum degree 4, and* ColSub(H_k) *does not admit an $O(n^{\alpha \cdot k/\log k})$-time algorithm.*

Proof. For every $k \in \mathbb{N}$, pick $\ell \in \mathbb{N}$ maximal such that $|V(\check{B}_\ell)| \leq k$. Let H_k be obtained from $\check{B}_\ell$ by adding isolated vertices until the number of vertices is k. Since $|V(\check{B}_\ell)| = 2^{\ell+1}\ell$, we conclude that $2^{\ell+1}\ell \leq k < 2^{\ell+2}(\ell+1)$ which implies that $k/\log_2 k < 2^{\ell+2}$. So

$$\gamma(H_k) \geq \gamma(\check{B}_\ell) \geq 2^\ell/3 > \frac{1}{12} \cdot k/\log_2 k$$

by Corollary 4.3. Now the theorem follows from Theorem 1.2. □

Remark 4.5. Observe that Theorem 1.1 provides a sequence of graphs of maximum degree 3 whereas Theorem 4.4 "only" guarantees maximum degree 4. However, the augmented Beneš networks $\check{B}_\ell$ can easily be modified to have maximum degree 3 by replacing each vertex with an edge, so $\succ\!\!\bullet\!\!\prec$ becomes $\succ\!\bullet\!\!-\!\!\bullet\!\prec$, and all other relevant properties remain the same.

For readers familiar with expander graphs, let us also remark that the Beneš network B_ℓ with $s = 2^\ell$ does *not* have constant expansion, as witnessed by its "upper half" U that contains the vertices of $B^{\uparrow}_{\ell-1}$ and all inputs and outputs with indices $i \in [s/2]$: We have $|U| = s \log_2 s$, but the $2s$ neighbors of U are all contained in the first two and last two columns of B_ℓ. This also holds for the augmented $\check{B}_\ell$.

Universality of augmented Beneš networks

As an independent point of interest, let us remark that blowups of Beneš networks are universal for bounded-degree graphs with respect to topological minor containment: Mimicking the proof of Lemma 3.11, every n-vertex graph of maximum degree Δ can be found as a topological minor in the 2Δ-blowup of an augmented Beneš network with n inputs.

For comparison, every graph that contains every n-vertex graph of maximum degree Δ as a *subgraph* must necessarily have $\Omega(n^{2-2/\Delta})$ edges [3]. Under the more relaxed notion of universality via topological minor containment, Beneš networks show that universal graphs with only $O(\Delta^2 \cdot n \log n)$ vertices and edges are achievable.

Theorem 4.6. *For every $n, \ell \in \mathbb{N}$, every graph G of maximum degree Δ and $n \leq 2^\ell$ vertices is a topological minor of $\check{B}_\ell \otimes J_{2\Delta-1}$. Moreover, on input G, a topological G-minor model in $\check{B}_\ell \otimes J_{2\Delta-1}$ can be computed in polynomial time.*

Proof. By Theorem 4.2, the inputs in $\check{B}_\ell$ form a matching-linked set X of size $s = 2^\ell \geq n$. View $V(G) \subseteq X$ and decompose $E(G)$ into $2\Delta - 1$ matchings via the greedy edge-coloring algorithm. For each matching M, use Theorem 4.2 to find an M-linkage Q in $\check{B}_\ell$ and place the internal vertices of Q in a private layer of $\check{B}_\ell \otimes J_{2\Delta}$, as in the proof of Lemma 3.11. The union of the linkages constructed this way is a topological G-minor model in $\check{B}_\ell \otimes J_{2\Delta}$. □

5 Patterns of Superlinear Density

We turn our attention to dense patterns, i.e., k-vertex patterns H of average degree $d(H) \in \omega(1)$. Unlike the sparse setting discussed earlier, a linkage capacity of $\Theta(k)$ is achievable in the dense case, which implies tight lower bounds for ColSub(H) under ETH.

5.1 Worst Case

We show that, for every graph H, the average degree $d(H) = 2|E(H)|/|V(H)|$ is a lower bound on the linkage capacity of H, up to a constant factor. First, we use Mader's Theorem [58, Corollary 1] to extract a highly connected subgraph from H. A graph H is *ℓ-connected* if $|V(H)| > \ell$ and $H - X$ is connected for every set $X \subseteq V(H)$ with $|X| < \ell$.

Theorem 5.1 (see [30, Theorem 1.4.3])**.** *Every graph H with $d(H) \geq 4\ell$ contains a $(\ell + 1)$-connected subgraph H' with $d(H') > d(H) - 2\ell$.*

Second, within the scope of this subsection only, we say that a graph H is *ℓ-globally linked* if $|V(H)| \geq 2\ell$ and each set $X \subseteq V(H)$ of size at most 2ℓ is matching-linked in H. (In graph theory, this notion is usually just called *ℓ-linked*—see, e.g., [30]. In our paper, we refer to it as *ℓ-globally linked* to distinguish it from our previous definitions of linkedness.) This definition implies in particular that H contains a matching-linked set X with $|X| \geq 2\ell$. Thomas and Wollan [66, Corollary 1.2] show that high connectivity implies high global linkedness.

Theorem 5.2 (see [30, Theorem 3.5.3])**.** *Let H be a graph and $\ell \in \mathbb{N}$. If H is 2ℓ-connected and $d(H) \geq 16\ell$, then H is ℓ-globally linked.*

Together, these two theorems imply a lower bound on the linkage capacity that is linear in the average degree.

Lemma 5.3. *For every graph H, we have $\gamma(H) \geq d(H)/48$.*

Proof. Theorem 5.1 yields a $\lceil d(H)/4 \rceil$-connected subgraph H' of H with $d(H') > d(H)/2$. Theorem 5.2 shows that H' is $\lceil d(H)/32 \rceil$-globally linked and thus contains a matching-linked set X of size at least $d(H)/16$. Then Lemma 3.7 shows that $\gamma(H) \geq \gamma(H') \geq d(H)/48$, where the first inequality uses that H' is subgraph of H. □

Now, Theorem 1.3 follows from Theorem 1.2 and Lemma 5.3.

5.2 Average Case

To show the hardness in the average case, we consider the linkage capacity of the Erdős-Rényi random graph. Let $\mathcal{G}(k, p)$ denote the distribution over k-vertex graphs where each edge is included independently with probability p. We need the following theorem adapted from [15], where "with high probability" refers to a probability tending to 1 for $k \to \infty$.

Theorem 5.4. *Let $\varepsilon > 0$ be a constant. For all $p \geq (1 + \varepsilon) \log(k)/k$ the following holds: With high probability, for a random graph $H \sim \mathcal{G}(k, p)$, every matching M on $V(H)$ can be partitioned into $r = O(\log k/ \log kp)$ matchings $M_1, \ldots, M_r$ such that H contains an uncongested M_i-linkage for all $i \in [r]$.*

The original theorem statement and proof in [15, Corollary 1.1] are concerned with the fixed-sized random graph model $G(k, m)$, and only deals with even k. But on the other hand, they give a stronger statement concerning the algorithmic efficiency of finding the desired partition, that it can be obtained with high probability by a random partition. In Appendix C, we give a proof of the version stated here.

The last theorem can be used to find large matching-linked sets inside a proper blowup of a random graph, which implies a high linkage capacity by Lemma 3.7.

Lemma 5.5. *Let $\varepsilon > 0$ be a constant. For all $p \geq (1 + \varepsilon) \log(k)/k$, the linkage capacity of $H \sim \mathcal{G}(k, p)$ is at least $\Omega(\frac{k \log(kp)}{\log k})$ with high probability.*

Proof. Let r be the bound specified in Theorem 5.4 and consider the blowup graph $H \otimes J_{r+1}$. Let $X := \{v^{(1)} \mid v \in V(H)\}$. We show that X is matching-linked in $H \otimes J_{2r}$ with high probability, and the lemma then follows using Lemma 3.7.

Let M' be a matching on X. By the definition of X, its H-projection, $M := \pi(M')$, is also a matching on $V(H)$. We invoke Theorem 5.4 on the graph H with respect to the matching M to obtain a partition $M_1, \ldots, M_r$, such that, for all $i \in [r]$ there is an uncongested M_i-linkage Q_i in H. Then $Q = \bigcup_{i\in[r]} Q_i$ is an r-congested M-linkage in H. So there is an uncongested M'-linkage in $H \otimes J_{2r}$ by Lemma 2.5. □

Now, Theorem 1.4 follows from Theorem 1.2 and Lemma 5.5.

6 Large-Treewidth Patterns and Concurrent Flows

In this section, we relate the linkage capacity of a graph to its treewidth. Towards this end, we first connect the linkage capacity to certain (fractional) multicommodity flows, and afterward rely on existing connections between such flows and treewidth [59, Section 3.1].

More specifically, given a graph H and $W \subseteq V(H)$, we consider the following multicommodity flow problem. For every pair $(u, v) \in W^2$, there is a distinct commodity uv that can be sent in arbitrary fractional amounts along different paths from u to v in H. The goal is to determine whether all pairs (u, v) can concurrently send an ϵ amount of uv to each other, while the total flow through every vertex $w \in V(H)$ is at most some globally fixed capacity C. Formally, this is captured by the following LP:

Definition 6.1. Let H be a graph. For $u, v \in V(H)$, write $\mathcal{P}_H(uv)$ for the set of paths from u to v in H; the set $\mathcal{P}_H(uv)$ for $u = v$ contains only the path (u). Given $W \subseteq V(H)$, the concurrent flow LP (for H and W) with vertex capacity $C > 0$ asks to

$$\begin{aligned} &\text{maximize } \varepsilon \\ &\text{subject to} \quad \sum_{p\in\mathcal{P}_H(uv)} x_p \geq \varepsilon \quad \forall u, v \in W \\ &\sum_{u,v\in W}\ \sum_{p\in\mathcal{P}_H(uv):\, w\in p} x_p \leq C \quad \forall w \in V(G) \\ &x_p \geq 0 \quad \forall u, v \in W, p \in \mathcal{P}_H(uv). \end{aligned}$$

We write $\varepsilon(H, W)$ to denote the optimal LP value for capacity $C = 1$.

While an optimal solution for $C = 1$ may assign fractional values to the variables x_p, every solution can be scaled to an integral solution, increasing the required capacity and the optimal LP value by the same factor. This integral solution then induces a congested model of the multigraph $K_{t,q}$ in H, where $t := |W|$ and $q \in \mathbb{N}$ is suitably chosen, and $K_{t,q}$ has t vertices and contains each possible (undirected) edge with multiplicity q.

Lemma 6.2. *Let H be a graph and $W \subseteq V(H)$ be a set of size t. Then there is some $D \in \mathbb{N}$ such that $q := D \cdot \varepsilon(H, W)$ is an integer and H contains a D-congested $K_{t,q}$-linkage, where we set $V(K_{t,q}) = W$.*

Proof. Let D be the common denominator of the values for all x_p in a (rational) optimal solution of the concurrent flow LP for H and W with capacity $C = 1$. Scaling all values by D yields an integral solution of value $q := D \cdot \varepsilon(H, W)$ for the LP with capacity D. Now, consider the multiset Q which, for every distinct $u, v \in W$, contains every path $p \in \mathcal{P}_H(uv)$ with multiplicity x_p. Then Q is a D-congested $K_{t,q}$-linkage where $V(K_{t,q}) = W$. □

Using this congested $K_{t,q}$-linkage, we will establish lower bounds on the linkage capacity of H. The following lemma will be useful, as it allows us to route arbitrary multigraphs of bounded degree via short paths in this $K_{t,q}$.

Lemma 6.3. *Let $q \in \mathbb{N}$ and let M be a multigraph with $V(M) = [t]$ and maximum degree at most qt. Then there is an M-linkage $Q = (P_{uv})_{uv \in E(M)}$ in K_t such that every edge $e \in E(K_t)$ appears in at most $18q$ paths in Q.*

Proof. If $t \leq 12$ the statement trivially holds by choosing $P_{uv} = (u, v)$ for every $uv \in E(M)$. So in the remainder of the proof, we assume that $t > 12$.

First observe that $|E(M)| \leq qt^2/2$ since $\deg_M(v) \leq qt$ for all $v \in V(M)$. For every $e = uv \in E(M)$ we set $P_{uv} = (u, x_e, v)$ for some suitable *middle vertex* $x_e \in V(K_t) \setminus \{u, v\}$. We construct the paths one by one in a greedy fashion. Suppose $\mathcal{P}$ is the collection of paths constructed so far. We ensure that

(a) every edge of K_t appears in at most $18q$ paths in $\mathcal{P}$, and

(b) every vertex is the middle vertex on at most qt paths in $\mathcal{P}$.

Now consider an edge $e = uv \in E(M)$ that is not yet covered by $\mathcal{P}$. We argue that there is some $x \in V(K_t) \setminus \{u, v\}$ such that $\mathcal{P} \cup \{(u, x, v)\}$ still satisfies Conditions (a) and (b).

Since $|E(M)| \leq qt^2/2$, there are at most $t/2$ vertices that are the middle vertex of exactly qt paths in $\mathcal{P}$ (i.e., they cannot be selected as a middle vertex). Also, the total number of edges incident to u used in $\mathcal{P}$ is at most $3qt$ since $\deg_M(u) \leq qt$ and u is the middle vertex of at most qt paths. So there are most $t/6$ vertices $x \in V(G) \setminus \{u\}$ such that the edge ux is *full*, i.e., ux already appears in $18q$ paths in $\mathcal{P}$. Similarly, there are most $t/6$ vertices $x \in V(G) \setminus \{v\}$ such that the edge vx is full. Since

$$\frac{t}{2} + 2 \cdot \left(1 + \frac{t}{6}\right) = \frac{5}{6}t + 2 < t,$$

there exists at least one $x_e \in V(K_t) \setminus \{u, v\}$ such that $\mathcal{P} \cup \{(u, x_e, v)\}$ still satisfies Conditions (a) and (b). □

We can conclude that a large value of the concurrent flow LP implies large linkage capacity.

Theorem 6.4. *Let H be a graph and $W \subseteq V(H)$. Then $\gamma(H) \geq \varepsilon(H, W) \cdot |W|^2/108$.*

Proof of Theorem 6.4. Let $D \in \mathbb{N}$ be the integer obtained from Lemma 6.2 and set $D' := 18 \cdot D$. Let $q := D \cdot \varepsilon(H, W)$ which, by Lemma 6.2, is an integer. Observe that $\varepsilon(H, W) \leq 1/|W|$, so $D \geq q \cdot |W|$. Finally, let $s := q \cdot |W|$.

Consider the graph $H \otimes J_{2D'}$ and let

$$X := \{w^{(i)} \mid w \in W, i \in [s]\}.$$

We show that X is matching-linked in $H \otimes J_{2D'}$. Let M be a matching on X. Let $\widehat{M} := \pi(M)$ be the H-projection of M. Observe that $\deg_{\widehat{M}}(w) \leq s = q \cdot |W|$ for all $w \in W$.

Lemma 6.3 finds an $\widehat{M}$-linkage $Q = (P_{uv})_{uv \in E(\widehat{M})}$ in $K_{t,q}$ (where $V(K_{t,q}) = W$) such that every edge of $K_{t,q}$ appears in at most 18 of those paths. Moreover, by Lemma 6.2, the graph H contains a D-congested $K_{t,q}$-linkage $Q' = (P'_{uv})_{uv \in E(K_{t,q})}$. We construct a D'-congested $\widehat{M}$-linkage $\widehat{Q} = (\widehat{P}_{uv})_{uv \in E(\widehat{M})}$ in H as follows. For every $uv \in E(\widehat{M})$ we obtain $\widehat{P}_{uv}$ from P_{uv} by substituting P'_e for every edge e appearing on P_{uv}. Clearly, $\widehat{Q}$ is D'-congested since $D' = 18 \cdot D$. So there is an uncongested M-linkage in $H \otimes J_{2D'}$ by Lemma 2.5.

Overall, we get that X is matching-linked in $H \otimes J_{2D'}$. So

$$\gamma(G) \geq \frac{1}{3} \cdot \frac{|X|}{2 \cdot D'} = \frac{1}{3} \cdot \frac{s \cdot |W|}{36 \cdot D} = \frac{1}{108} \cdot \frac{q \cdot |W|^2}{D} = \frac{1}{108} \cdot \varepsilon(H, W) \cdot |W|^2$$

by Lemma 3.7. □

To bound the linkage capacity by the treewidth, we combine Theorem 6.4 with the following lemma that is (implicitly) shown by Marx [59].

Lemma 6.5 ([59])**.** *Let H be a graph of treewidth t. Then there is a set $W \subseteq V(H)$ such that $|W| = t$ and $\varepsilon(H, W) = \Omega(1/(t \log t))$.*

The basic idea to prove the lemma is to consider a set $W \subseteq V(H)$ of size t that does not admit balanced separators; large treewidth guarantees such a set (see [59, Lemma 3.2]). Then, using results from [37, 55], we obtain a bound on the optimal value of the dual LP, which gives $\varepsilon(H, W) = \Omega(1/(t \log t))$ (see [59, Proof of Lemma 3.6]).

Corollary 6.6. *Let H be a graph of treewidth t. Then $\gamma(H) = \Omega(t/\log t)$.*

In particular, combining Theorem 1.2 and Corollary 6.6 allows us to recover the complexity lower bounds on $\textsc{ColSub}(H)$ proved in [59].

7 Implications for Counting Small Induced Subgraphs

We conclude with an application of our lower bounds for the complexity of counting induced k-vertex subgraphs. A *k-vertex graph invariant* Φ is an isomorphism-invariant map from k-vertex graphs H to some ring. We consider Φ to be fixed and wish to sum $\Phi(G[X])$ over all k-vertex subsets X of an input graph G to count, e.g., the planar or Hamiltonian induced k-vertex subgraphs of G. Formally, for a k-vertex graph invariant Φ, the problem $\#\textsc{IndSub}(\Phi)$ takes as input a graph G, and asks to compute

$$\#\mathrm{IndSub}(\Phi \to G) \coloneqq \sum_{X \subseteq V(G)} \Phi(G[X]).$$

This problem was first studied in its parameterized version (where k is part of the input) by Jerrum and Meeks [46, 47, 48] and received significant attention in recent years [23, 24, 31, 32, 33, 38, 63, 64].

To determine the complexity of $\#\textsc{IndSub}(\Phi)$, recent works usually analyze the *alternating enumerator* to build a generic reduction from $\#\textsc{ColSub}(H)$. Formally, the *alternating enumerator* of a graph invariant Φ on a graph H is defined as[4]

$$\widehat{\Phi}(H) = (-1)^{|E(H)|} \sum_{S \subseteq E(H)} (-1)^{|S|}\, \Phi(H[S]),$$

where $H[S]$ has vertex set $V(H)$ and edge set S. Now, suppose H is a k-vertex graph with $\widehat{\Phi}(H) \neq 0$. Then the problem $\#\textsc{ColSub}(H)$ can be reduced to $\#\textsc{IndSub}(\Phi)$ in polynomial time (see, e.g., [24, Lemmas 3.3 & A.3]). Hence, building on Theorem 1.2, we also obtain new lower bounds for $\#\textsc{IndSub}(\Phi)$ assuming $\widehat{\Phi}(H) \neq 0$ for suitable graphs H. In particular, we obtain the following result via Lemma 5.3 which improves over the corresponding lower bound in [24, Theorem 3.5(a)] (see also [31] and [32, Lemma 2.2]). We give a proof in Appendix D.

Theorem 7.1. *There is a universal constant $\alpha_{\mathrm{IND}} > 0$ and an integer $N_0 \geq 1$ such that for all numbers $k, \ell \geq 1$, the following holds: If Φ is a k-vertex graph invariant and there exists a graph H with $\widehat{\Phi}(H) \neq 0$ and $E(H) \geq k \cdot \ell \geq N_0$, then $\#\textsc{IndSub}(\Phi)$ cannot be solved in time $O(n^{\alpha_{\mathrm{IND}} \cdot \ell})$ unless ETH fails.*

As pointed out in Section 1.2, the weaker version, which only rules out an exponent of $\alpha_{\mathrm{IND}} \cdot \ell/\sqrt{\log \ell}$, has been used to derive various lower bounds for specific types of invariants in [24, 32, 64]. All these lower bounds are improved by our new results. Let us give one concrete example, which improves over [24, Corollary 5.2]. For a k-vertex graph invariant Φ, we write $\mathrm{supp}(\Phi)$ for the set of all graphs H with $V(H) = [k]$ and $\Phi(H) \neq 0$.

[4] The precise formula is not relevant here, but we still give it for completeness.

Corollary 7.2. *For every $0 < \varepsilon < 1$ there are $N_0, \delta > 0$ such that the following holds. Let $k \geq N_0$ and let Φ be a k-vertex graph invariant with $1 \leq |\mathrm{supp}(\Phi)| \leq (2-\varepsilon)^{\binom{k}{2}}$. Then no algorithm solves $\#\textsc{IndSub}(\Phi)$ in time $O(n^{\delta \cdot k})$ unless ETH fails.*

We stress that the exponent in the lower bound of Corollary 7.2 is asymptotically optimal.

Proof. Let $0 < \varepsilon < 1$. By [24, Theorem 5.1] there is some $\delta' > 0$ such that for every $k \geq 1$ and every k-vertex graph invariant Φ satisfying the condition of the theorem, there is a k-vertex graph H such that $\widehat{\Phi}(H) \neq 0$ and $|E(H)| \geq \delta' \cdot \binom{k}{2}$. Therefore, the theorem follows from Theorem 7.1 by setting $\delta := \frac{1}{3} \cdot \delta' \cdot \alpha_{\mathrm{IND}}$ and choosing $\ell := \frac{1}{3} \cdot \delta' \cdot k$. $\square$

References

[1] Akanksha Agrawal, Fahad Panolan, Saket Saurabh, and Meirav Zehavi. Simultaneous feedback edge set: A parameterized perspective. *Algorithmica*, 83(2):753–774, 2021. doi:10.1007/S00453-020-00773-9.

[2] Noga Alon. Explicit expanders of every degree and size. *Comb.*, 41(4):447–463, 2021. doi:10.1007/S00493-020-4429-X.

[3] Noga Alon, Michael R. Capalbo, Yoshiharu Kohayakawa, Vojtech Rödl, Andrzej Rucinski, and Endre Szemerédi. Universality and tolerance. In *41st Annual Symposium on Foundations of Computer Science, FOCS 2000, 12-14 November 2000, Redondo Beach, California, USA*, pages 14–21. IEEE Computer Society, 2000. doi:10.1109/SFCS.2000.892007.

[4] Noga Alon and Dániel Marx. Sparse balanced partitions and the complexity of subgraph problems. *SIAM J. Discret. Math.*, 25(2):631–644, 2011. doi:10.1137/100812653.

[5] Saeed A. Amiri, Stephan Kreutzer, Dániel Marx, and Roman Rabinovich. Routing with congestion in acyclic digraphs. *Inf. Process. Lett.*, 151, 2019. doi:10.1016/J.IPL.2019.105836.

[6] Václav E. Beneš. Permutation groups, complexes, and rearrangeable connecting networks. *Bell System Tech. J.*, 43(4):1619–1640, 1964. doi:10.1002/j.1538-7305.1964.tb04102.x.

[7] Béla Bollobás. The diameter of random graphs. *Trans. Amer. Math. Soc.*, 267(1):41–52, 1981. doi:10.2307/1998567.

[8] Édouard Bonnet, Sergio Cabello, Bojan Mohar, and Hebert Pérez-Rosés. The inverse voronoi problem in graphs I: hardness. *Algorithmica*, 82(10):3018–3040, 2020. doi:10.1007/S00453-020-00716-4.

[9] Edouard Bonnet, Panos Giannopoulos, and Michael Lampis. On the parameterized complexity of red-blue points separation. *J. Comput. Geom.*, 10(1):181–206, 2019. doi:10.20382/JOCG.V10I1A7.

[10] Édouard Bonnet, Yoichi Iwata, Bart M. P. Jansen, and Lukasz Kowalik. Fine-grained complexity of k-OPT in bounded-degree graphs for solving TSP. In Michael A. Bender, Ola Svensson, and Grzegorz Herman, editors, *27th Annual European Symposium on Algorithms, ESA 2019, September 9-11, 2019, Munich/Garching, Germany*, volume 144 of *LIPIcs*, pages 23:1–23:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICS.ESA.2019.23.

[11] Édouard Bonnet and Tillmann Miltzow. Parameterized hardness of art gallery problems. *ACM Trans. Algorithms*, 16(4):42:1–42:23, 2020. doi:10.1145/3398684.

[12] Édouard Bonnet and Florian Sikora. The graph motif problem parameterized by the structure of the input graph. *Discret. Appl. Math.*, 231:78–94, 2017. doi:10.1016/J.DAM.2016.11.016.

[13] Karl Bringmann. Fine-grained complexity theory (tutorial). In Rolf Niedermeier and Christophe Paul, editors, *36th International Symposium on Theoretical Aspects of Computer Science, STACS 2019, March 13-16, 2019, Berlin, Germany*, volume 126 of *LIPIcs*, pages 4:1–4:7. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICS.STACS.2019.4.

[14] Karl Bringmann, László Kozma, Shay Moran, and N. S. Narayanaswamy. Hitting set for hypergraphs of low VC-dimension. In Piotr Sankowski and Christos D. Zaroliagis, editors, *24th Annual European Symposium on Algorithms, ESA 2016, August 22-24, 2016, Aarhus, Denmark*, volume 57 of *LIPIcs*, pages 23:1–23:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPICS.ESA.2016.23.

[15] Andrei Z. Broder, Alan M. Frieze, Stephen Suen, and Eli Upfal. An efficient algorithm for the vertex-disjoint paths problem in random graphs. In Éva Tardos, editor, *Proceedings of the Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, 28-30 January 1996, Atlanta, Georgia, USA*, pages 261–268. ACM/SIAM, 1996. URL: http://dl.acm.org/citation.cfm?id=313852.314072.

[16] Jianer Chen, Benny Chor, Mike Fellows, Xiuzhen Huang, David W. Juedes, Iyad A. Kanj, and Ge Xia. Tight lower bounds for certain parameterized NP-hard problems. *Inf. Comput.*, 201(2):216–231, 2005. doi:10.1016/J.IC.2005.05.001.

[17] Jianer Chen, Xiuzhen Huang, Iyad A. Kanj, and Ge Xia. Strong computational lower bounds via parameterized complexity. *J. Comput. Syst. Sci.*, 72(8):1346–1367, 2006. doi:10.1016/J.JCSS.2006.04.007.

[18] Rajesh Chitnis, Andreas E. Feldmann, and Pasin Manurangsi. Parameterized approximation algorithms for bidirected Steiner network problems. *ACM Trans. Algorithms*, 17(2):12:1–12:68, 2021. doi:10.1145/3447584.

[19] Rajesh Hemant Chitnis, Andreas E. Feldmann, Mohammad T. Hajiaghayi, and Dániel Marx. Tight bounds for planar strongly connected Steiner subgraph with fixed number of terminals (and extensions). *SIAM J. Comput.*, 49(2):318–364, 2020. doi:10.1137/18M122371X.

[20] Vincent Cohen-Addad, Éric C. de Verdière, Dániel Marx, and Arnaud de Mesmay. Almost tight lower bounds for hard cutting problems in embedded graphs. *J. ACM*, 68(4):30:1–30:26, 2021. doi:10.1145/3450704.

[21] Jason Crampton, Robert Crowston, Gregory Z. Gutin, Mark Jones, and Maadapuzhi S. Ramanujan. Fixed-parameter tractability of workflow satisfiability in the presence of seniority constraints. In Michael R. Fellows, Xuehou Tan, and Binhai Zhu, editors, *Frontiers in Algorithmics* and *Algorithmic Aspects in Information and Management, Third Joint International Conference, FAW-AAIM 2013, Dalian, China, June 26-28, 2013. Proceedings*, volume 7924 of *Lecture Notes in Computer Science*, pages 198–209. Springer, 2013. doi:10.1007/978-3-642-38756-2_21.

[22] Radu Curticapean, Holger Dell, and Thore Husfeldt. Modular counting of subgraphs: Matchings, matching-splittable graphs, and paths. In Petra Mutzel, Rasmus Pagh, and Grzegorz Herman, editors, *29th Annual European Symposium on Algorithms, ESA 2021, September 6-8, 2021, Lisbon, Portugal (Virtual Conference)*, volume 204 of *LIPIcs*, pages 34:1–34:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPICS.ESA.2021.34.

[23] Radu Curticapean, Holger Dell, and Dániel Marx. Homomorphisms are a good basis for counting small subgraphs. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 210–223. ACM, 2017. doi:10.1145/3055399.3055502.

[24] Radu Curticapean and Daniel Neuen. Counting small induced subgraphs: Hardness via fourier analysis. *CoRR*, abs/2407.07051, 2024. arXiv:2407.07051.

[25] Radu Curticapean and Mingji Xia. Parameterizing the permanent: Genus, apices, minors, evaluation mod 2k. In Venkatesan Guruswami, editor, *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 994–1009. IEEE Computer Society, 2015. doi:10.1109/FOCS.2015.65.

[26] Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015. doi:10.1007/978-3-319-21275-3.

[27] Argyrios Deligkas, Eduard Eiben, and Tiger-Lily Goldsmith. Parameterized complexity of hotelling-downs with party nominees. In Luc De Raedt, editor, *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, pages 244–250. ijcai.org, 2022. doi:10.24963/IJCAI.2022/35.

[28] Holger Dell, Thore Husfeldt, Dániel Marx, Nina Taslaman, and Martin Wahlen. Exponential time complexity of the permanent and the tutte polynomial. *ACM Trans. Algorithms*, 10(4):21:1–21:32, 2014. doi:10.1145/2635812.

[29] Erik D. Demaine and MohammadTaghi Hajiaghayi. Linearity of grid minors in treewidth with applications through bidimensionality. *Comb.*, 28(1):19–36, 2008. doi:10.1007/S00493-008-2140-4.

[30] Reinhard Diestel. *Graph Theory*. Springer Berlin, 5 edition, 2017. doi:10.1007/978-3-662-53622-3.

[31] Julian Dörfler, Marc Roth, Johannes Schmitt, and Philip Wellnitz. Counting induced subgraphs: An algebraic approach to #W[1]-hardness. *Algorithmica*, 84(2):379–404, 2022. doi:10.1007/s00453-021-00894-9.

[32] Simon Döring, Dániel Marx, and Philip Wellnitz. Counting small induced subgraphs with edge-monotone properties. In Bojan Mohar, Igor Shinkar, and Ryan O'Donnell, editors, *Proceedings of the 56th Annual ACM Symposium on Theory of Computing, STOC 2024, Vancouver, BC, Canada, June 24-28, 2024*, pages 1517–1525. ACM, 2024. doi:10.1145/3618260.3649644.

[33] Simon Döring, Dániel Marx, and Philip Wellnitz. From graph properties to graph parameters: Tight bounds for counting on small subgraphs. *CoRR*, abs/2407.06801, 2024. arXiv:2407.06801, doi:10.48550/ARXIV.2407.06801.

[34] Eduard Eiben, Gregory Z. Gutin, Philip R. Neary, Clément Rambaud, Magnus Wahlström, and Anders Yeo. Preference swaps for the stable matching problem. *Theor. Comput. Sci.*, 940(Part):222–230, 2023. doi:10.1016/J.TCS.2022.11.003.

[35] Eduard Eiben, Dusan Knop, Fahad Panolan, and Ondrej Suchý. Complexity of the Steiner network problem with respect to the number of terminals. In Rolf Niedermeier and Christophe Paul, editors, *36th International Symposium on Theoretical Aspects of Computer Science, STACS 2019, March 13-16, 2019, Berlin, Germany*, volume 126 of *LIPIcs*, pages 25:1–25:17.

Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICS.STACS.2019.25.

[36] David Eppstein and Daniel Lokshtanov. The parameterized complexity of finding point sets with hereditary properties. In Christophe Paul and Michal Pilipczuk, editors, *13th International Symposium on Parameterized and Exact Computation, IPEC 2018, August 20-24, 2018, Helsinki, Finland*, volume 115 of *LIPIcs*, pages 11:1–11:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPICS.IPEC.2018.11.

[37] Uriel Feige, MohammadTaghi Hajiaghayi, and James R. Lee. Improved approximation algorithms for minimum weight vertex separators. *SIAM J. Comput.*, 38(2):629–657, 2008. doi:10.1137/05064299X.

[38] Jacob Focke and Marc Roth. Counting small induced subgraphs with hereditary properties. *SIAM J. Comput.*, 53(2):189–220, 2024. doi:10.1137/22m1512211.

[39] Fedor V. Fomin, Fahad Panolan, Maadapuzhi S. Ramanujan, and Saket Saurabh. On the optimality of pseudo-polynomial algorithms for integer programming. *Math. Program.*, 198(1):561–593, 2023. doi:10.1007/S10107-022-01783-X.

[40] Robert Ganian. Using neighborhood diversity to solve hard problems. *CoRR*, abs/1201.3091, 2012. arXiv:1201.3091.

[41] Mika Göös, Rahul Jain, and Thomas Watson. Extension complexity of independent set polytopes. *SIAM J. Comput.*, 47(1):241–269, 2018. doi:10.1137/16M109884X.

[42] Jiong Guo, Sepp Hartung, Rolf Niedermeier, and Ondrej Suchý. The parameterized complexity of local search for TSP, more refined. *Algorithmica*, 67(1):89–110, 2013. doi:10.1007/S00453-012-9685-8.

[43] Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-SAT. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001. doi:10.1006/JCSS.2000.1727.

[44] Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.*, 63(4):512–530, 2001. doi:10.1006/JCSS.2001.1774.

[45] Klaus Jansen, Stefan Kratsch, Dániel Marx, and Ildikó Schlotter. Bin packing with fixed number of bins revisited. *J. Comput. Syst. Sci.*, 79(1):39–49, 2013. doi:10.1016/J.JCSS.2012.04.004.

[46] Mark Jerrum and Kitty Meeks. The parameterised complexity of counting connected subgraphs and graph motifs. *J. Comput. Syst. Sci.*, 81(4):702–716, 2015. doi:10.1016/j.jcss.2014.11.015.

[47] Mark Jerrum and Kitty Meeks. Some hard families of parameterized counting problems. *ACM Trans. Comput. Theory*, 7(3):11:1–11:18, 2015. doi:10.1145/2786017.

[48] Mark Jerrum and Kitty Meeks. The parameterised complexity of counting even and odd induced subgraphs. *Comb.*, 37(5):965–990, 2017. doi:10.1007/s00493-016-3338-5.

[49] Mark Jones, Daniel Lokshtanov, Maadapuzhi S. Ramanujan, Saket Saurabh, and Ondrej Suchý. Parameterized complexity of directed Steiner tree on sparse graphs. *SIAM J. Discret. Math.*, 31(2):1294–1327, 2017. doi:10.1137/15M103618X.

[50] Karthik C. Srikanta, Dániel Marx, Marcin Pilipczuk, and Uéverton S. Souza. Conditional lower bounds for sparse parameterized 2-csp: A streamlined proof. In Merav Parter and Seth Pettie, editors, *2024 Symposium on Simplicity in Algorithms, SOSA 2024, Alexandria, VA, USA, January 8-10, 2024*, pages 383–395. SIAM, 2024. doi:10.1137/1.9781611977936.35.

[51] Victor Klee and David Larman. Diameters of random graphs. *Canadian J. Math.*, 33(3):618–640, 1981. doi:10.4153/CJM-1981-050-1.

[52] Dusan Knop, Simon Schierreich, and Ondrej Suchý. Balancing the spread of two opinions in sparse social networks (student abstract). In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, pages 12987–12988. AAAI Press, 2022. doi:10.1609/AAAI.V36I11.21630.

[53] Alexandr V. Kostochka. Lower bound of the hadwiger number of graphs by their average degree. *Comb.*, 4(4):307–316, 1984. doi:10.1007/BF02579141.

[54] Eric Lehman, F. Thomson Leighton, and Albert R. Meyer. Beneš Network, June 30 2021. [Online; accessed 2024-09-26]. URL: https://eng.libretexts.org/@go/page/48364.

[55] Frank T. Leighton and Satish Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *J. ACM*, 46(6):787–832, 1999. doi:10.1145/331524.331526.

[56] Daniel Lokshtanov, Dániel Marx, and Saket Saurabh. Lower bounds based on the exponential time hypothesis. *Bull. EATCS*, 105:41–72, 2011.

[57] Daniel Lokshtanov, Maadapuzhi S. Ramanujan, Saket Saurabh, and Meirav Zehavi. Parameterized complexity and approximability of directed odd cycle transversal. In Shuchi Chawla, editor, *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 2181–2200. SIAM, 2020. doi:10.1137/1.9781611975994.134.

[58] Wolfgang Mader. Existenz n-fach zusammenhängender Teilgraphen in Graphen genügend großer Kantendichte. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, 37:86–97, 1972. doi:10.1007/BF02993903.

[59] Dániel Marx. Can you beat treewidth? *Theory Comput.*, 6(1):85–112, 2010. doi:10.4086/toc.2010.v006a005.

[60] Jesper Nederlof and Céline M. F. Swennenhuis. On the fine-grained parameterized complexity of partial scheduling to minimize the makespan. *Algorithmica*, 84(8):2309–2334, 2022. doi:10.1007/S00453-022-00970-8.

[61] Jakob Nordström. New wine into old wineskins: A survey of some pebbling classics with supplemental results. Technical report, KTH Royal Institute of Technology, 2015.

[62] Marcin Pilipczuk and Magnus Wahlström. Directed multicut is W[1]-hard, even for four terminal pairs. *ACM Trans. Comput. Theory*, 10(3):13:1–13:18, 2018. doi:10.1145/3201775.

[63] Marc Roth and Johannes Schmitt. Counting induced subgraphs: A topological approach to #W[1]-hardness. *Algorithmica*, 82(8):2267–2291, 2020. doi:10.1007/s00453-020-00676-9.

[64] Marc Roth, Johannes Schmitt, and Philip Wellnitz. Counting small induced subgraphs satisfying monotone properties. In Sandy Irani, editor, *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 1356–1367. IEEE, 2020. doi:10.1109/FOCS46700.2020.00128.

[65] Claude E. Shannon. A theorem on coloring the lines of a network. *J. Math. Physics*, 28(1-4):148–152, 1949. doi:10.1002/sapm1949281148.

[66] Robin Thomas and Paul Wollan. An improved linear edge bound for graph linkages. *Eur. J. Comb.*, 26(3-4):309–324, 2005. doi:10.1016/J.EJC.2004.02.013.

[67] Virginia V. Williams. Hardness of easy problems: Basing hardness on popular conjectures such as the strong exponential time hypothesis (invited talk). In Thore Husfeldt and Iyad A. Kanj, editors, *10th International Symposium on Parameterized and Exact Computation, IPEC 2015, September 16-18, 2015, Patras, Greece*, volume 43 of *LIPIcs*, pages 17–29. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2015. doi:10.4230/LIPICS.IPEC.2015.17.

[68] Virginia V. Williams. On some fine-grained questions in algorithms and complexity. In *Proceedings of the International Congress of Mathematicians—Rio de Janeiro 2018. Vol. IV. Invited lectures*, pages 3447–3487. World Sci. Publ., Hackensack, NJ, 2018. doi:10.1142/9789813272880_0188.

A Linkage Capacity

In this section we prove Theorem 3.10.

Proof of Theorem 3.10. We first prove the following claim.

Claim A.1. *There exists an algorithm $\mathbb{A}$ that given a vertex set X of $H \otimes J_t$ and a matching M with vertex set X, checks if there is an uncongested M-linkage in $H \otimes J_t$ in time $O(t^{f(k)})$.*

Proof. Our algorithm works in two steps. In the first step, we select a collection of paths in H. In the second step, we embed the paths into $H \otimes J_t$. First, for each pair of vertices $u, v \in H$, we write $\mathcal{P}_{u,v}$ for the set of simple paths from u to v in H. We write $\mathcal{P}$ for the union of all $\mathcal{P}_{u,v}$. Note that $|\mathcal{P}| \leq k^{2+k}$ and that we can enumerate all elements in $\mathcal{P}$ in time $O(k^{3+k})$.

In the next step, we consider all collections Q of $\mathcal{P}$ with multiplicity at most t.[5] We check if Q defines an uncongested M-linkage in $H \otimes J_t$ using the following subroutine.

Let $V := V(H \otimes J_t)$ and $N := M$. For an edge $v^{(i)}u^{(j)} \in N$, check if there is a path $P = (v, a_2, \ldots, a_{s-1}, u)$ in Q. If so, check if there is a vertex $a_i^{(b_i)}$ in V for each i. This defines a path $P' = (v^{(i)}, a_1^{(b_1)}, \ldots, a_{s-1}^{(b_{s-1})}, u^j)$ in $H \otimes J_t$ that goes from $v^{(i)}$ to $u^{(j)}$ since P is a path in H. Lastly, remove P from Q, $v^{(i)}u^{(j)}$ from N, and the vertices of P' from V. We repeat this process until either we abort, or N is empty. If N is empty, then Q' defines a valid uncongested M-linkage since for each edge $v^{(i)}u^{(j)} \in M$ there is a path in Q from $v^{(i)}$ to $u^{(j)}$, and each vertex appears in at most one path. This subroutine runs in time $O(t \cdot k^3)$ since there are at most $t \cdot k$ many edges in M, we can find a path P in Q in time $O(\log(k^{2+k}))$ using some basic data structure, and we can embed P in G in time $O(k)$ in G.

We apply this subroutine to all collections Q of $\mathcal{P}$ with multiplicity at most t. If we find an uncongested M-linkage this way, we return that X contains an uncongested M-linkage. Otherwise, we return false. Since there are at most $(t+1)^{k^{2+k}}$ many possible collections, the algorithm runs in $O((t+1)^{k^{2+k}} \cdot tk^3)$ which is in $O(t^{f(k)})$ for $f(k) = 3k^{2+k}$. ⌟

To find a matching-linked set X of maximum size, we first observe that each vertex set $X \subseteq H \otimes J_t$ corresponds to a multiset X' of H where X' contains the vertex v with multiplicity $|\{v^{(i)} : i \in [t]\}|$. Thus X' has multiplicity at most t.

It is easy to see that if two vertex sets X and Y define the same multiset X' then there is an automorphism in $H \otimes J_t$ that maps a vertex $v^{(i)}$ only to a vertex of the form $v^{(j)}$ and that maps X to Y. Given a matching M on X and a matching N on Y, this automorphism also maps each uncongested M-linkage to an uncongested N-linkage. Thus X is a matching-linked set if and only if Y is a matching-linked set. So, we only have to consider the $(t+1)^k$ many multisets of H' with multiplicity at most t.

Further, we apply the same idea to matchings M in $H \otimes J_t$. For each edge $v^{(i)}u^{(j)} \in M$, we add vu to a multiset M'. Observe that M' has multiplicity at most t. It is now easy to see that if two matchings M and N correspond to the same multiset M' then $H \otimes J_t$ contains a M-linkage if and only if $H \otimes J_t$ contains a N-linkage. Thus we only have to consider the $(t+1)^{k^2}$ many multisets M' with multiplicity at most t.

Lastly, to find the maximal linked set we first iterate over all multisets X' of $V(H)$ with multiplicity at most t. Each X' defines a vertex set $X = \{v^{(i)} : 1 \leq i \leq d(X'_v)\}$ where $d(X'_v)$ is the multiplicity of the vertex v in X'. To check if X is matching-linked, we iterate over all multisets M' with multiplicity at most t. Each set M' defines a matching M in the following way. Let $V := V(H \otimes J_t)$. For each $vu \in M'$, we add $v^{(i)}u^{(j)}$ to M where $i := \min(s : v^{(s)} \in V)$ and $j := \min(s : u^{(s)} \in V)$.[6] If not possible, we just ignore M' and consider the next multiset. Then remove $v^{(i)}$ and $u^{(j)}$ from V and continue with the next element in the multiset M'.

[5] This means that each path P appears at most t times in Q

[6] if $v = u$, then we add $v^{(i)}v^{(i+1)}$ to M.

This way, we obtain a matching M that corresponds to M'. We check M is a matching on X. If not, we check the next multiset M'. Otherwise, we use algorithm $\mathbb{A}$ from Claim A.1 to check if X is a M-linkage. If this is not the case, then X is not matching-linked. After considering all $(t+1)^{k^2}$ many multiset M', we know if X is matching-linked. By applying this procedure to all multisets X', we find the largest matching-linked set X in time

$$O(\underbrace{(t+1)^k}_{\text{vertex sets}} \cdot \underbrace{(t+1)^{k^2}}_{\text{matchings}} \cdot \underbrace{(t+1)^{k^2+k} \cdot tk^3}_{\text{Algorithm } \mathbb{A}}),$$

which is in $O(t^{f(k)})$ for $f(k) = 3k^{2+k}$. □

B Linkages in Beneš Networks

We prove Theorem 4.2. First, we consider linkages between inputs and outputs in plain Beneš networks. Algorithm 2 guarantees the existence of such linkages and constructs them efficiently, as shown in Lemma B.1. Linkages between inputs and outputs in plain Beneš networks then readily imply linkages among inputs in augmented Beneš networks.

Lemma B.1 ([6])**.** *Given as input $\ell \in \mathbb{N}$ and a perfect matching M between the $s = 2^\ell$ inputs and outputs of B_ℓ, the procedure* BenesLink(ℓ, M) *in Algorithm 2 computes an M-linkage in B_ℓ in time $O(s \log s)$.*

Proof. We first validate Line 20 by proving that the resulting graph D is always bipartite. Construct two graphs D_1 and D_2 containing only the edges added on Line 18 and Line 19 respectively. Trivially, D_1 is a perfect matching. We show that D_2 is also a perfect matching, which implies that the graph D contains only even cycles and is hence bipartite. For each $i \in [s]$, we let $r(i) := i + s/2$ where addition tacitly wraps into range $[s]$, so $r(r(i)) = i$. Let $\pi \in S_r$ be the permutation specified the perfect matching M; that is, $v_i w_{\pi(i)} \in M$. By the construction of Line 19, an edge ab is in D_2 if and only if $a = \pi^{-1} r \pi(b)$. Because $(\pi^{-1} r \pi)(\pi^{-1} r \pi) = \mathtt{id}_s$ is the identity mapping, the mapping $\pi^{-1} r \pi$ defines the perfect matching D_2.

We then prove by induction that the collection of paths Q returned by BenesLink(ℓ, M) indeed forms an uncongested M linkage. The base case B_1 is easy to verify. Suppose this is true for $B_{\ell-1}$, meaning $Q^\uparrow$ in Line 8 is a $M^\uparrow$-linkage in $B^\uparrow$ (and the same for $Q^\downarrow$ in Line 10). First, the new edges $v_i L(i)$ and $R(j) w_j$ exist due the construction of B_ℓ, so Q is indeed a collection of paths. To show that Q is uncongested, note that each vertex $v_c^\uparrow$ and $v_c^\downarrow$ is adjacent only to v_c and $v_{c+s/2}$, so the constraints introduced in Line 18 ensure that the mapping L is injective. Therefore, all vertices $v_c^\uparrow$ and $v_c^\downarrow$ appears in exactly one path in Q. For the same reason, all vertices $w_c^\uparrow$ and $w_c^\downarrow$ appears in exactly one path in Q, too. Together with the induction hypothesis, the set Q is uncongested.

Finally, we analyze the running time $T(s)$ of BenesLink(ℓ, M) for $s = 2^\ell$ in the word RAM model. Because the mapping $\pi^{-1} r \pi$ is easy to compute, Line 20 can be implemented in $\Theta(s)$ time by a DFS. All the other operations except for the recursive calls cost $\Theta(s)$ time in total. We obtain the recurrence $T(s) = 2T(s/2) + \Theta(s)$, with $T(s) = \Theta(s \log s)$ by the Master Theorem. □

Proof of Theorem 4.2. Let $\ell \in \mathbb{N}$ and $s = 2^\ell$, and denote the inputs and outputs of the augmented network $\check{B}_\ell$ by $V = \{v_1, \ldots, v_s\}$ and $W = \{w_1, \ldots, w_s\}$. Note that these sets also exist in the plain network B_ℓ. Given a perfect matching $\check{M} = \{e_1, \ldots, e_{s/2}\}$ on vertex set V, we construct an $\check{M}$-linkage in $\check{B}_\ell$ in time $O(s \log s)$.

First, define from $\check{M}$ a perfect matching M between V and W to be used in the plain network: For $i \in [s/2]$, write $e_i = ab$ and include edges $a w_{2i-1}$ and $b w_{2i}$ into M. Lemma B.1 finds an M-linkage Q in the plain network B_ℓ in time $O(s \log s)$. We construct from Q an $\check{M}$-linkage $\check{Q}$ in $\check{B}_\ell$: For $ab \in \check{M}$, writing w and w' for the unique partners of a and b in M, the concatenated

Algorithm 2 Routing from inputs to outputs in plain Beneš networks

procedure BenesLink(ℓ, M) computes an M-linkage in B_ℓ
  if $\ell = 1$ **then**
    return M, interpreted as linkage
  else
    $k \leftarrow 2^\ell$
    $F, L, R \leftarrow$ResolveConflict(s, M)
      ▷ *match inputs/outputs of* B_ℓ *with inputs/outputs of* $B^{\uparrow}_{\ell-1}$ *or* $B^{\downarrow}_{\ell-1}$
      ▷ *here,* $F : [s] \to \{\uparrow, \downarrow\}$
      ▷ $L : [s] \to \{v^{\uparrow}_1, \cdots, v^{\uparrow}_{s/2}, v^{\downarrow}_1, \cdots, v^{\downarrow}_{s/2}\}$ *maps to an input in* $B^{\uparrow}_{\ell-1}$ *or* $B^{\downarrow}_{\ell-1}$
      ▷ $R : [s] \to \{w^{\uparrow}_1, \cdots, w^{\uparrow}_{s/2}, w^{\downarrow}_1, \cdots, w^{\downarrow}_{s/2}\}$ *maps to an output in* $B^{\uparrow}_{\ell-1}$ *or* $B^{\downarrow}_{\ell-1}$
    $M^{\uparrow} \leftarrow \{L(i)R(j) \mid F(i) = \uparrow, v_i w_j \in M\}$
    $Q^{\uparrow} = \{P^{\uparrow}_{v^{\uparrow} w^{\uparrow}}\} \leftarrow$ BenesLink($\ell - 1$, $M^{\uparrow}$) ▷ *find linkage in* $B^{\uparrow}_{\ell-1}$
    $M^{\downarrow} \leftarrow \{L(i)R(j) \mid F(i) = \downarrow, v_i w_j \in M\}$
    $Q^{\downarrow} = \{P^{\downarrow}_{v^{\downarrow} w^{\downarrow}}\} \leftarrow$ BenesLink($\ell - 1$, $M^{\downarrow}$) ▷ *find linkage in* $B^{\downarrow}_{\ell-1}$
    for $i \in [k]$ **do** let j be the unique index such that $v_i w_j \in M$
      Let $P_{v_i w_j} \leftarrow v_i L(i) \circ P^{F(i)}_{L(i)R(j)} \circ R(j) w_j$ and add to Q
    return Q

procedure ResolveConflict(s, M) outputs three mappings F, L and R
  Define $t(i) := \begin{cases} i, & i \le s/2; \\ i - s/2, & i > s/2. \end{cases}$
  Let $D = ([s], \emptyset)$. ▷ *undirected* 2*-regular conflict graph*
  for $i \in [s/2]$ **do** add edges ab to D ▷ *adding two conflicts per iteration*
    with $a = i$ and $b = i + s/2$ ▷ *conflicting paths at inputs* i *and* $i + s/2$
    with unique a, b such that $v_a w_i \in M$ and $v_b w_{i+s/2} \in M$
      ▷ *conflicting paths at outputs* i *and* $i + s/2$*, translating to inputs* a *and* b
  Compute a proper 2-coloring F of D using colors $\{\uparrow, \downarrow\}$
  Construct L by letting $L(i) \leftarrow v^{F(i)}_{t(i)}$.
  Construct R by letting $R(j) \leftarrow w^{F(i)}_{t(j)}$ where i the unique index with $v_i w_j \in M$.

path $\check{P}_{ab} = P_{aw}\, w\, w'\, P_{w'b}$ exists in $\check{B}_\ell$. Because all paths $\check{P}_{ab}$ exist in $\check{B}_\ell$ and are vertex-disjoint, it follows that $\check{Q} = \{\check{P}_{ab} \mid ab \in \check{M}\}$ is an $\check{M}$-linkage in $\check{B}_\ell$. □

C Linkages in Random Graphs

We use $\mathcal{G}(k, p)$ for the Erdős-Rényi random graph model with edge probability p on k vertices, and $G(k, m)$ for the uniform distribution over all graphs with k vertices and m edges.

An *equipartition* of a set M into r parts is a partition $M_1, \cdots, M_r$ such that $|M_i - M_j| \le 1$ for all $i, j \in [r]$.

Theorem C.1 ([15, Corollary 1.1])**.** *Let* $\varepsilon > 0$ *be a constant and* $d(k) \ge (1+\varepsilon)\log(k)$*. With high probability, for random* $H \sim G(k, m)$ *with even* k *and* $m = k \cdot d(k)/2$ *and any perfect matching* M *on vertices* $[k]$*, with high probability a random equipartition of* M *into* $r = \lceil \beta \log(k)/\log(d) \rceil$ *matchings* $M_1, \ldots, M_r$ *satisfies that* H *contains an* M_i*-linkage for all* $i \in [r]$*.*

We first transfer the above result from the $G(k, m)$ model to the $\mathcal{G}(k, p)$ model.

Corollary C.2. *Let* $\varepsilon' > 0$ *be a constant and* $p \ge (1 + \varepsilon')\log(k)/k$*. With high probability, for random* $H \sim \mathcal{G}(k, p)$ *with even* k *and any perfect matching* M *on vertices* $[k]$*, with high probability* *a random equipartition of M into $r = \lceil \beta \log(k)/\log(kp) \rceil$ matchings $M_1, \ldots, M_r$ satisfies that H contains an M_i-linkage for all $i \in [r]$.*

Proof. Let $\mathcal{P}$ be the graph property specified in Theorem C.1 that $G(k, m)$ satisfies with high probability, and $\overline{\mathcal{P}}$ be the negation of $\mathcal{P}$. We show that, in the setting of this corollary, any graph drawn from the $\mathcal{G}(k, p)$ model satisfies $\mathcal{P}$ with high probability.

Note that $\mathcal{P}$ is a monotone increasing property, meaning if H satisfies $\mathcal{P}$ then $H + e$ satisfies $\mathcal{P}$ too. Therefore, $\overline{\mathcal{P}}$ is monotone decreasing. By coupling, it holds that

$$\Pr_{H \sim G(k,m_1)}[H \in \overline{\mathcal{P}}] \leq \Pr_{H \sim G(k,m_2)}[H \in \overline{\mathcal{P}}] \qquad \text{for } m_1 \geq m_2. \tag{5}$$

Take $\varepsilon = \varepsilon'/3$, and set $m^* := \frac{1+\varepsilon}{1+2\varepsilon} p \cdot \binom{k}{2}$. Draw a graph $H \sim \mathcal{G}(k, p)$. By the law of total probability, we have

$$\Pr_{H \sim \mathcal{G}(k,p)}[H \in \overline{\mathcal{P}}] \leq \underbrace{\Pr_{H \sim \mathcal{G}(k,p)}[|E(H)| < m^*]}_{①} + \underbrace{\Pr_{H \sim \mathcal{G}(k,p)}[H \in \overline{\mathcal{P}} \wedge |E(H)| \geq m^*]}_{②}. \tag{6}$$

We bound term ① by a standard Chernoff bound

$$① = \Pr\left[\mathrm{Bin}\left(\binom{k}{2}, p\right) < \left(1 - \frac{\varepsilon}{1+2\varepsilon}\right)\binom{k}{2}p\right] \leq \exp\left\{-\binom{k}{2}p \cdot \frac{\varepsilon^2}{2(1+2\varepsilon)^2}\right\} = O(k^{-k}).$$

We bound term ② by the following.

$$\begin{aligned}
② &= \sum_{t=m^*}^{\binom{k}{2}} \Pr_{H \sim \mathcal{G}(k,p)}\left[H \in \overline{\mathcal{P}} \;\middle|\; |E(H)| = t\right] \cdot \Pr_{H \sim \mathcal{G}(k,p)}[|E(H)| = t] \\
&= \sum_{t=m^*}^{\binom{k}{2}} \Pr_{H' \sim G(k,t)}[H' \in \overline{\mathcal{P}}] \cdot \Pr_{H \sim \mathcal{G}(k,p)}[|E(H)| = t] \\
&\leq \sum_{t=m^*}^{\binom{k}{2}} \Pr_{H' \sim G(k,m^*)}[H' \in \overline{\mathcal{P}}] \cdot \Pr_{H \sim \mathcal{G}(k,p)}[|E(H)| = t] \qquad \text{(by (5))} \\
&= \Pr_{H' \sim G(k,m^*)}[H' \in \overline{\mathcal{P}}] \cdot \sum_{t=m^*}^{\binom{k}{2}} \Pr_{H \sim \mathcal{G}(k,p)}[|E(H)| = t] \leq \Pr_{H' \sim G(k,m^*)}[H' \in \overline{\mathcal{P}}].
\end{aligned}$$

Because

$$\frac{2m^*}{k} \geq \left(\frac{1+3\varepsilon}{1+2\varepsilon} \cdot \frac{k-1}{k}\right) \cdot (1+\varepsilon)\log k,$$

we invoke Theorem C.1 with constant ε for large enough k, to see that ② is also negligible. □

Proof of Theorem 5.4. Assume k is even. We extend the matching M to a perfect matching M' by pairing the unmatched vertices. By Corollary C.2, a random equipartition of M' into $M'_1, \cdots, M'_r$ satisfies the desired property. We then drop the edges in $M' \setminus M$ from this partition to obtain $M_1, \cdots, M_r$ with the desired property.

Assume k is odd. Given the matching M, we find an arbitrary unmatched vertex $w \in V(H)$. The induced subgraph $H - w$ is subject to the uniform distribution $\mathcal{G}(k-1, p)$. In the regime of Corollary C.2, if $p = p(k) \geq (1+\varepsilon')\log(k)/k$, then $p(k-1) \geq (1+\varepsilon'')\log(k)/k$ for some other constant $\varepsilon'' > 0$. Therefore, we can invoke Corollary C.2 again, and the rest of the argument is the same as the even k case. □

D Counting Small Induced Subgraphs

We give a proof of Theorem 7.1, relying on [24, Lemma 3.3 & A.3], stated below.

Lemma D.1 ([24, Lemma 3.3])**.** *Given a k-vertex graph invariant Φ, for $k \geq 1$, we have*

$$\#\mathrm{IndSub}(\Phi \to \star) = \sum_{H} \widehat{\Phi}(H) \cdot \#\mathrm{Sub}(H \to \star), \tag{7}$$

where H ranges over all unlabelled k-vertex graphs.

Let $G = (V, E, c)$ be a colored graph where $c\colon V(G) \to C$. We define $G^\circ = (V, E)$ to be the uncolored version of G. For $i \in C$, we write $V_i(G)$ for the vertices of color i, and for $i, j \in C$, we write $E_{ij}(G)$ for the edges in G with one endpoint of color i and another of color j. For $X \subseteq C$ and $Y \subseteq \binom{C}{2}$, let $G_{\setminus X,Y}$ be the graph obtained from G by deleting all vertices with colors from X and all edges whose endpoints have a color pair from Y, i.e,

$$G_{\setminus X,Y} = \left(V \setminus \bigcup_{i\in X} V_i,\ E \setminus \bigcup_{ij\in Y} E_{ij} \right).$$

The following observation is immediate.

Observation D.2. *For graphs H and G with canonically colored H, we have*

$$\#\mathrm{Sub}(H \to G) = \#\mathrm{Sub}(H \to G_{\setminus\emptyset,\overline{E(H)}}).$$

Also, we write $G \cong H$ to denote that two (colored or uncolored) graphs G, H are isomorphic.

Lemma D.3 ([24, Lemma A.3])**.** *Let $k \in \mathbb{N}$ and let the following be given:*

- *Numbers $\alpha_1, \ldots, \alpha_s \in \mathbb{Q}$ and pairwise non-isomorphic uncolored graphs $H_1, \ldots, H_s$ with $|V(H_i)| = k$ for all $i \in [s]$, which define the graph invariant*

$$f(\star) := \sum_{i=1}^{s} \alpha_i \cdot \#\mathrm{Sub}(H_i \to \star),$$

- *a canonically colored graph H with $V(H) = [k]$ and $H^\circ \cong H_b$ for some $b \in [s]$, and*
- *a colored graph G with coloring $c : V(G) \to [k]$ satisfying $E_{ij}(G) = \emptyset$ for $ij \notin E(H)$.*

Then we have

$$\alpha_b \cdot \#\mathrm{Sub}(H \to G) = \sum_{\substack{X\subseteq V(H)\\ Y\subseteq E(H)}} (-1)^{|X|+|Y|} f(G^\circ_{\setminus X,Y}).$$

With these tools at our disposal, we are ready to prove Theorem 7.1.

Proof of Theorem 7.1. Let Φ be a k-vertex graph invariant and suppose H is a graph with $\widehat{\Phi}(H) \neq 0$ and $E(H) \geq k \cdot \ell \geq N_0$. Without loss of generality assume $V(H) = [k]$. We give an algorithm for #ColSub(H) that uses an algorithm for #IndSub(Φ) as a subroutine.

Let G be the input graph for the problem #ColSub(H). We may assume that $G = G_{\setminus\emptyset,\overline{E(H)}}$ by Observation D.2. We wish to determine $\#\mathrm{Sub}(H^{\mathsf{can}} \to G)$. By Lemma D.1 we have

$$\#\mathrm{IndSub}(\Phi \to \star) = \sum_{F} \widehat{\Phi}(F) \cdot \#\mathrm{Sub}(F \to \star) =: f(\star),$$

where F ranges over all k-vertex graphs. Invoking Lemma D.3, we obtain that

$$\widehat{\Phi}(H) \cdot \#\mathrm{Sub}(H^{\mathsf{can}} \to G) = \sum_{\substack{X \subseteq V(H) \\ Y \subseteq E(H)}} (-1)^{|X|+|Y|} f(G^{\circ}_{\setminus X,Y}). \tag{8}$$

So we can compute $\#\mathrm{Sub}(H^{\mathsf{can}} \to G)$ by evaluating the right-hand side of (8) and dividing by $\widehat{\Phi}(H) \neq 0$. Note that all relevant values $f(G^{\circ}_{\setminus X,Y})$ can be obtained by the oracle calls $\#\mathrm{IndSub}(\Phi \to G^{\circ}_{\setminus X,Y})$ without parameter increase in overall time $2^{|V(H)|+|E(H)|} \cdot n^{O(1)}$. The value $\widehat{\Phi}(H)$ can be computed by brute-force by evaluating Φ on $2^{O(k^2)}$ many k-vertex graphs. Hence, an $O(n^{\beta \cdot \ell})$ algorithm for #IndSub(Φ) gives an $O(n^{c \cdot \beta \cdot \ell})$ for #ColSub(H) for some suitable fixed constant c. Now, the theorem follows from Theorem 1.3 □