



Density Matters: A Complexity Dichotomy of Deleting Edges to Bound Subgraph Density

Matthias Bentert  

Technische Universität Berlin, Germany
University of Bergen, Norway

Tom-Lukas Breitkopf  

Algorithmics and Computational Complexity, Technische Universität Berlin, Germany

Vincent Froese  

Algorithmics and Computational Complexity, Technische Universität Berlin, Germany

Anton Herrmann  

Algorithmics and Computational Complexity, Technische Universität Berlin, Germany

André Nichterlein  

Algorithmics and Computational Complexity, Technische Universität Berlin, Germany

Abstract

We study τ -BOUNDED-DENSITY EDGE DELETION (τ -BDED), where given an undirected graph G , the task is to remove as few edges as possible to obtain a graph G' where no subgraph of G' has density more than τ . The density of a (sub)graph is the number of edges divided by the number of vertices. This problem was recently introduced and shown to be NP-hard for $\tau \in \{2/3, 3/4, 1+1/25\}$, but polynomial-time solvable for $\tau \in \{0, 1/2, 1\}$ [Bazgan et al., JCSS 2025]. We provide a complete dichotomy with respect to the target density τ :

1. If $2\tau \in \mathbb{N}$ (half-integral target density) or $\tau < 2/3$, then τ -BDED is polynomial-time solvable.
2. Otherwise, τ -BDED is NP-hard.

We complement the NP-hardness with fixed-parameter tractability with respect to the treewidth of G . Moreover, for integral target density $\tau \in \mathbb{N}$, we show τ -BDED to be solvable in randomized $O(m^{1+o(1)})$ time. Our algorithmic results are based on a reduction to a new general flow problem on restricted networks that, depending on τ , can be solved via MAXIMUM s - t -FLOW or GENERAL FACTORS. We believe this connection between these variants of flow and matching to be of independent interest.

2012 ACM Subject Classification Theory of computation \rightarrow Parameterized complexity and exact algorithms; Theory of computation \rightarrow Problems, reductions and completeness; Mathematics of computing \rightarrow Matchings and factors; Mathematics of computing \rightarrow Network flows

Keywords and phrases Transshipment, Maximum Flow, General Factors, Matching, Graph Modification Problem

Digital Object Identifier 10.4230/LIPIcs.STACS.2026.12

Related Version *Full Version:* <https://arxiv.org/abs/2601.03129>

Funding The research leading to these results has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No. 819416) and the German Research Foundation (DFG, grant 1-5003036: SPP 2378 – ReNO-2).

1 Introduction

Finding a densest subgraph in a given undirected graph is an “evergreen research topic” [16]. The density of a (sub)graph is the number of edges divided by the number of vertices, which is equal to half the average degree. A graph with n vertices can thus have a density up



© Matthias Bentert, Tom-Lukas Breitkopf, Vincent Froese, Anton Herrmann, and André Nichterlein;
licensed under Creative Commons License CC-BY 4.0

43rd International Symposium on Theoretical Aspects of Computer Science (STACS 2026).

Editors: Meena Mahajan, Florin Manea, Annabelle McIver, and Nguyễn Kim Thăng

Article No. 12; pp. 12:1–12:20



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



to $(n - 1)/2$; with this maximum being reached if and only if the graph is a clique. While a reader unfamiliar with the problem might suspect it to be NP-hard – like `CLIQUE` – it can in fact be solved in polynomial time: This was first discovered 1979 and published 1982 by Picard and Queyranne [19] with an algorithm based on maximum-flow computations (cf. Lanciano et al. [16]). An improved algorithm due to Goldberg [12] also relies on a maximum s - t -flow.

In this work, we study the following closely related problem called τ -BOUNDED-DENSITY EDGE DELETION (τ -BDED) where $\tau \geq 0$: Given an undirected graph $G = (V, E)$, delete as few edges as possible such that the densest subgraph of the remaining graph has density at most τ . This problem was recently introduced by Bazgan et al. [3], who observed that for $\tau < 1$ the remaining graph has to be a forest where each tree can have at most $\lfloor 1/(1 - \tau) \rfloor$ vertices. Thus, τ -BDED generalizes the polynomial-time solvable MAXIMUM CARDINALITY MATCHING ($\tau = 1/2$) and the NP-hard PERFECT P_3 -PACKING ($\tau = 2/3$). For $\tau > 1$, not much is known about the complexity of τ -BDED beyond NP-hardness for τ slightly above 1 ($\tau = 1 + 1/c$ with c being any constant larger than 24 [3]).

Our main result is a complete complexity dichotomy for τ -BDED:

► **Theorem 1.1.** *If $2\tau \in \mathbb{N}$ or $\tau < 2/3$, then τ -BOUNDED-DENSITY EDGE DELETION is polynomial-time solvable, otherwise it is NP-hard. Moreover, if $\tau \in \mathbb{N}$, then it is solvable in randomized $O(m^{1+o(1)})$ time.*

Additionally, we show fixed-parameter tractability for the parameter treewidth. To this end, note that for a graph of treewidth tw a tree decomposition of width tw can be computed in time $2^{O(tw^2)}n^{O(1)}$ [15] and a tree decomposition of width $2tw + 1$ in time $2^{O(tw)}n$ [14].

► **Theorem 1.2.** *τ -BOUNDED-DENSITY EDGE DELETION is fixed-parameter tractable with respect to treewidth. Given a tree decomposition of width tw with $\tau = \frac{a}{b}$, it can be solved in time $(\max\{a, b\} + 1)^{tw+1}n^{O(1)}$.*

On the technical side, we follow the line of Picard and Queyranne [19] and Goldberg [12]: We achieve our algorithmic results by efficiently reducing to a generalized flow problem (which we call 2-LAYER GENERAL TRANSSHIPMENT (2GT)). For integral target density $\tau \in \mathbb{N}$, 2GT can be solved with a simple maximum s - t -flow computation in a linear-size network with small integer capacities (at most τ). This allows employing recent algorithms for computing maximum flows [7] to solve τ -BDED with $\tau \in \mathbb{N}$ in randomized $O(m^{1+o(1)})$ time on graphs with m edges. For $\tau = 1/2$, our flow approach to solve the 2GT instance fails. The reason is that we rely on a maximum flow being integral and some arcs having capacity τ . In such cases, a simple scaling of capacities yields integer values. However, our flow-problem then gets an additional constraint that some arcs either transport a full flow (equal to the capacity) or no flow at all. These “all-or-nothing” restrictions render the MAXIMUM FLOW problem NP-hard in general, even if all arcs have integer capacities of at most two.¹ The restricted network structure (being bipartite with all arcs directed towards one partition) in our flow problem, however, allows a reduction to GENERAL FACTORS, which is a generalization of MAXIMUM CARDINALITY MATCHING. Note that this reverses the standard approach of solving MAXIMUM CARDINALITY MATCHING in bipartite graphs by a reduction to MAXIMUM

¹ While we did not find this statement in the scientific literature, an NP-hardness proof was given by Neal Young on stackexchange: <https://cstheory.stackexchange.com/questions/51243/maximum-flow-with-parity-requirement-on-certain-edges>. The full proof is included in the full version of this paper.

FLOW [21]. Leveraging known algorithmic results for GENERAL FACTORS, we obtain a polynomial-time algorithm for half-integral target density τ and fixed-parameter tractability with respect to treewidth.

Related Work. We refer to the survey of Lanciano et al. [16] for an overview over DENSEST SUBGRAPH and various approaches to solve it and its variants.

To the best of our knowledge, so far only Bazgan et al. [3] investigated BOUNDED-DENSITY EDGE DELETION (BDED), the problem variant of τ -BDED where the target density is part of the input and can thus depend on n . Bazgan et al. used the allowed dependency on n to prove W[1]-hardness of BDED with respect to treewidth [3] (the proof requires a target density $\tau = 1 - 1/o(n)$). Considering the relation between the choice of τ and the computational complexity of the problem, they obtained the following results: BDED is NP-hard in case $\tau \in [\frac{2}{3}, \frac{3}{4}]$ or $\tau = 1 + \frac{1}{c}$, where c is any constant larger than 24, but BDED becomes polynomial-time solvable in case $\tau < \frac{2}{3}$ or $\tau \in [1 - \frac{1}{n}, 1]$.

The vertex deletion variant of BDED is called BOUNDED-DENSITY VERTEX DELETION (BDVD) and generalizes VERTEX COVER ($\tau = 0$) and FEEDBACK VERTEX SET ($\tau = 1 - 1/n$). It was also introduced by Bazgan et al. [3], who showed the following results: BDVD is NP-hard for any constant target density. Both BDED and BDVD are fixed-parameter tractable with respect to the vertex cover number, but W[1]- respectively W[2]-hard with respect to the solution size, that is, the number of edges, respectively vertices, to remove.

Chandrasekaran et al. [5] studied the approximability of BDVD. Interestingly, they (also) focus on integral target density $\tau \geq 2$ and provide a polynomial-time $O(\log n)$ -approximation that cannot be improved to a factor of $o(\log n)$ assuming $P \neq NP$. Their approximation algorithm indeed solves more general problems like SUPERMODULAR DENSITY DELETION and MATROID FEEDBACK VERTEX SET.

Organization. After introducing basic notions in Section 2, we provide an easy flow-construction for integral target density τ as warm-up in Section 3. We show our main result on the generalized flow variant in Section 4 and complement our findings with NP-hardness results for τ -BDED in Section 5 for the remaining target densities. We conclude in Section 6.

2 Preliminaries

We set $\mathbb{N} := \{0, 1, \dots\}$ to include 0 and denote by $\mathbb{Q}_0^+ := \{q \in \mathbb{Q} \mid q \geq 0\}$ ($\mathbb{Q}_0^- := \{q \in \mathbb{Q} \mid q \leq 0\}$) the set of non-negative (non-positive) rational numbers. For $n \in \mathbb{N}$, we define $[n] := \{1, 2, \dots, n\}$ with $[0] := \emptyset$. We use \mathbb{H} to denote the set of all non-negative half-integral rational numbers, that is, $\mathbb{H} = \{q \in \mathbb{Q} \mid 2q \in \mathbb{N}\}$.

Graph & Digraphs. All graphs in this work are simple and unweighted. Let G be an undirected graph and D be a digraph. We denote the set of vertices of G and D by $V(G)$ and $V(D)$ and the set of edges of G and arcs of D by $E(G)$ and $A(D)$, respectively. For $C \in \{G, D\}$, we denote with n_C and m_C the number of vertices and edges (arcs) of C . We denote the degree of a vertex $v \in V(G)$ by $\deg_G(v) := |N(v)|$, where $N(v) := \{u \in V(G) \mid \{u, v\} \in E(G)\}$. Similarly, we denote the in-degree (out-degree) of a vertex $v \in V(D)$ by $\deg_D^-(v) := |N^-(v)|$ ($\deg_D^+(v) := |N^+(v)|$), where $N^-(v) := \{u \in V(D) \mid (u, v) \in A(D)\}$ ($N^+(v) := \{u \in V(D) \mid (v, u) \in A(D)\}$). If the considered graph is clear from the context, then we drop the subscripts. We write $H \subseteq G$ to denote that H is a subgraph of G . For a subset $E' \subseteq E(G)$, we write $G - E'$ for the subgraph of G obtained by deleting the edges in E' . We denote by K_n the complete graph on n vertices (also called a clique of order n).

12:4 Density Matters: A Complexity Dichotomy of τ -Bounded-Density Edge Deletion

The *density* of G is $\rho(G) := m/n$, where the density of the empty graph is defined as zero. We denote by $\rho^*(G)$ the density of the densest subgraph of G , that is, $\rho^*(G) := \max_{H \subseteq G} \rho(H)$. A graph G is *balanced* if $\rho(H) \leq \rho(G)$ for every subgraph $H \subseteq G$; this implies $\rho^*(G) = \rho(G)$. We study the following problem, which is defined for every rational $\tau \geq 0$:

τ -BOUNDED-DENSITY EDGE DELETION (τ -BDED)

Input: An undirected graph G and an integer $k \in \mathbb{N}$.

Question: Is there a subset $F \subseteq E(G)$ with $|F| \leq k$ such that $\rho^*(G - F) \leq \tau$?

Throughout the rest of this work we assume G to be connected and hence have $n_G - 1 \leq m_G$ and $O(m_G + n_G) = O(m_G)$. Clearly, if G is not connected we can obtain a solution by solving each connected component individually and isolated vertices can be ignored.

Fractional Orientations. A crucial tool in our subsequent proofs and approaches are orientations, which follow from a dual LP formulation by Charikar [6, Section 3.1] to compute the density of a densest subgraph. An *orientation* of an undirected graph G assigns a direction to each edge $\{u, v\}$, that is, $\{u, v\}$ is replaced either by (u, v) or by (v, u) . Charikar [6] observed that the maximum indegree in any orientation of G is an upper bound for $\rho^*(G)$. This is also true for fractional orientations where we assign some part of an edge $\{u, v\}$ to u and the rest to v . More formally, a *fractional orientation* ϕ assigns to each edge $e = \{u, v\}$ two non-negative rationals $\phi(e)^u$ and $\phi(e)^v$ with $\phi(e)^u + \phi(e)^v = 1$. We denote with

$$\deg_{\phi}^{-}(v) := \sum_{u \in N(v)} \phi(\{u, v\})^v \quad \text{and} \quad \Delta_{\phi}^{-} := \max_{v \in V(G)} \deg_{\phi}^{-}(v)$$

the *indegree* of a vertex $v \in V(G)$ and the *maximum indegree* of G with respect to the fractional orientation ϕ . From now on we write $\phi(uv)^v$ as a short-hand for $\phi(\{u, v\})^v$.

We restate the aforementioned property as a lemma together with a stricter formulation.

► **Lemma 2.1.** *For every graph G and fractional orientation ϕ , it holds that $\rho^*(G) \leq \Delta_{\phi}^{-}$.*

► **Lemma 2.2** (Charikar [6]). *For every graph G , there is a fractional orientation ϕ with $\rho^*(G) = \Delta_{\phi}^{-}$.*

It follows from the work of Charikar [6] and the fact that LPs can be solved in polynomial time [21], that an orientation ϕ with $\rho^*(G) = \Delta_{\phi}^{-}$ can be computed in polynomial time.

Flows and Factors. Studied for well over half a century MAXIMUM FLOW is still subject of active research [7]. We refer to standard text books for an overview on flows [1, 21] and subsequently recall very briefly only the most relevant notions and concepts. In the basic setting, we are given a directed graph D forming the flow network with two distinguished vertices $s, t \in V(D)$ and arcs that are equipped with capacities $\text{cap}: A(D) \rightarrow \mathbb{Q}_0^+$. A flow $f: A(D) \rightarrow \mathbb{Q}_0^+$ assigns to each arc $a \in A(D)$ a non-negative amount of flow $f(a) \leq \text{cap}(a)$ such that for each vertex $v \in V(D) \setminus \{s, t\}$ the amount $f_{\Sigma}^{-}(v) := \sum_{u \in N^{-}(v)} f(u, v)$ of flow entering v is equal to the amount $f_{\Sigma}^{+}(v) := \sum_{u \in N^{+}(v)} f(v, u)$ of flow leaving v . Here, and throughout the paper, we write $f(a, b)$ as a short-hand for $f((a, b))$. The excess of the flow at v is $f_{\Sigma}(v) := f_{\Sigma}^{+}(v) - f_{\Sigma}^{-}(v)$, that is, s and t are the only vertices with non-zero excess. The value of a flow is defined as $\text{val}(f) = f_{\Sigma}(s) = -f_{\Sigma}(t)$.

Many extensions on the basic MAXIMUM FLOW have been studied, including costs on arcs, multiple sources and sinks, or lower bounds on the flow amount for the arcs (to name just a few). Computing maximum flows respecting all these restrictions can be done in

polynomial time [1, 21]. Computing a maximum flow that is even on *all* arcs (while still respecting capacities) is still doable in polynomial time (and seems a common exercise given the online search results). However, restricting a subset $F \subsetneq A(D)$ of arcs to have either flow 0 or 2 is already sufficient to make the problem NP-hard (see full version of the paper).

In contrast to MAXIMUM FLOW, the following MATCHING generalization is still polynomial-time solvable if restricted to even (or odd) numbers.

GENERAL FACTORS

Input: An undirected graph G , a collection of sets $\mathcal{B} = \{B_v \subseteq \{0, 1, \dots, \deg_G(v)\} \mid v \in V(G)\}$.

Question: Is there a subset $F \subseteq E(G)$, such that $\deg_{(V(G), F)}(v) \in B_v$ for every vertex $v \in V(G)$?

We call a solution F a *factor* of (G, \mathcal{B}) . GENERAL FACTORS with $B_v = \{1\}$ for each vertex v is exactly PERFECT MATCHING. Cornuéjols [8] showed that GENERAL FACTORS is polynomial-time solvable if the maximum gap in any B_v is of length at most one. Here, an integer subset $B \subseteq \mathbb{Z}$ has a gap of length $d \geq 1$ if there exists an integer $p \in B$ with $p + 1, \dots, p + d \notin B$ but $p + d + 1 \in B$. For example, the set $\{0, 2, 3, 5\}$ has gaps of length one as certified by $p = 0$ and $p = 3$. We denote by *max-gap* the maximum gap length over all sets $B_v \in \mathcal{B}$. Note that a max-gap of one is the boundary of polynomial-time solvability: There is a straightforward reduction from the NP-hard EXACT COVER BY 3-SETS [11, problem SP2] to GENERAL FACTORS where every vertex v has allowed degrees either $B_v = \{0, 3\}$ (the vertex represents a set) or $B_v = \{1\}$ (the vertex represents an element) [8]. Clearly, the constructed instance has max-gap two.

Recently, Marx et al. [17] provided an algorithm for GENERAL FACTORS running in $(M + 1)^{3 \text{tw}}$ time on graphs of treewidth tw , where M is the largest integer appearing in any set B_v .


3 Warm Up: Integral Target Density

Consider a τ -BOUNDED-DENSITY EDGE DELETION instance $\mathcal{I} = (G, k)$ with integral target density τ . We transform \mathcal{I} into a MAXIMUM FLOW instance with $O(m_G)$ vertices and arcs and small integral capacities. As a maximum flow in such networks can be computed in randomized $O(m^{1+o(1)})$ time [7], we get a near-linear time algorithm for τ -BDED if $\tau \in \mathbb{N}$.

The flow shall correspond to a fractional orientation. That is, one unit of flow should leave each edge $\{u, v\}$; a part of it flowing to u , the rest to v . Moreover, at each vertex v a total flow of at most τ is allowed to arrive. This idea easily translates to the following construction of the flow network D (see Figure 1 for a visualization of the construction):

- Set the vertex set of D as $V(D) = V(G) \cup E(G) \cup \{s, t\}$. This means for every vertex in G as well as every edge in G , there is a vertex in D .
- For every $e = \{u, v\} \in E(G)$ add an arc with capacity 1 to D going from source s to vertex e .
- For every $v \in V(G)$ add an arc with capacity τ to D going from v to sink t .
- For every edge $e = \{u, v\} \in E(G)$ add two arcs with capacity 1 to D , one going from vertex e to vertex u and one going from e to v .

Consider a maximum s - t -flow f in the constructed network D . Since all capacities are integral, it follows that f is integral [1, 21]. It is not hard to see that we have $\text{val}(f) = m_G$ (that is, all arcs leaving s transport a flow of 1) if and only if there is a fractional orientation ϕ in G with $\Delta_\phi^- \leq \tau$; the flow on the arcs in $E(G) \times V(G)$ corresponds to ϕ . However, a stronger statement holds: There is a solution $F \subseteq E(G)$ for \mathcal{I} of size $|F| = k$ if and only if f has value $m_G - k$. The correspondence between f and F becomes clear at the arcs



■ **Figure 1** Example transformation of a τ -BDED instance with $\tau \in \mathbb{N}$ (on the left) to a flow instance (on the right). All arcs on the same layer of the flow network have the same capacity as indicated at the leftmost arc.

incident to s : For an edge $e \in E(G)$ we have $f(s, e) = 0$ if and only if $e \in F$. We defer the formal proof of this statement to the next section (see Lemma 4.1), where we provide a transformation to a more general flow-variant that covers also non-integral target densities.

Note that the flow network D has size $O(m_G + n_G) = O(m_G)$ and can be constructed in linear time. Moreover, the largest reasonable value for τ is $\binom{n_G}{2}$. Using fast existing maximum-flow algorithms for polynomially bounded integral capacities [7], we obtain the following:

► **Theorem 3.1.** *There is a randomized algorithm solving τ -BOUNDED-DENSITY EDGE DELETION with integral target density $\tau \in \mathbb{N}$ in time $O(m^{1+o(1)})$.*

A further consequence of the flow-construction is that if $\tau \in \mathbb{N}$, then there exists a fractional orientation that is in fact integral. That is, each edge is assigned to either of its two endpoints. This was previously observed by Chandrasekaran et al. [5].

Note that for $\tau = 1/2$ the capacities in D are no longer integral. Therefore, an integral maximum flow is not guaranteed. However, τ -BDED for $\tau = 1/2$ corresponds to MAXIMUM CARDINALITY MATCHING. Hence, for our particular network D there is a polynomial-time algorithm that computes among all flows that are integral on the arcs leaving s a maximum one. Note that a linear-time reduction from MAXIMUM CARDINALITY MATCHING to FLOW is not known (so far). Thus, a polynomial-time algorithm for τ -BDED with half-integral capacities cannot rely on such a simple reduction as provided above. We present the solution to this issue in the next section.

4 Algorithmic Results

In this section, we provide efficient algorithms for τ -BDED with certain target densities. More precisely, we provide polynomial-time algorithms for integral and half-integral target densities, while for arbitrary constant target densities, we prove fixed-parameter tractability with respect to the treewidth of the input graph. Our approach builds on the overall idea of Section 3, that is, we reduce τ -BDED to some kind of flow problem. In particular, we introduce the 2-LAYER GENERAL TRANSSHIPMENT (2GT) problem, which is, informally spoken, a flow problem on graphs with multiple source and sink vertices, where every edge goes directly from a source to a sink and every vertex has a set of allowed excess values. Afterwards, we translate this problem to GENERAL FACTORS. Applying known results for GENERAL FACTORS, we finally deduce algorithmic results for 2GT, which, on the one hand, are of interest on their own and, on the other hand, can then be translated to τ -BDED.

4.1 Transshipment Formulation

To describe τ -BDED more generally, we switch to transshipment problems, that is, a flow-variant where we have multiple sources and sinks with specified amounts of flow that are provided resp. consumed by the sources and sinks [21]. Consider a digraph D . For a function $b: V(D) \rightarrow \mathbb{Q}$, a b -transshipment is a flow f with $f_\Sigma(v) = b(v)$ for all vertices $v \in V(D)$. Note that $\sum_{v \in V} b(v) = 0$ is necessary for a b -transshipment to exist. Adapting notation from GENERAL FACTORS, we specify a set collection $\mathcal{B} = \{B_v \subseteq \mathbb{Q} \mid v \in V(D)\}$ of allowed excess values for each vertex. A \mathcal{B} -transshipment is then a (flow) function $f: A(D) \rightarrow \mathbb{Q}_0^+$ with $f_\Sigma(v) \in B_v$ for all $v \in V(D)$. Note that if also arc capacities $\text{cap}: A(D) \rightarrow \mathbb{Q}_0^+$ are provided, then we have the usual additional requirement $f(a) \leq \text{cap}(a)$ for each arc $a \in A(D)$. We call an instance $(D, \text{cap}, \mathcal{B})$ *integral* if $\text{cap}: A(D) \rightarrow \mathbb{N}$ and $B_v \subseteq \mathbb{Z}$ for all $v \in V(D)$. We call a digraph D *2-layered*, if the underlying undirected graph is bipartite with parts U and W and all arcs go from U to W , that is, $A(D) \subseteq U \times W$. For a 2-layered network D with vertex partitions U and W and flow f , we define the *value* $\text{val}(f) = \sum_{u \in U} f_\Sigma(u) = \sum_{w \in W} -f_\Sigma(w)$. Note that in a 2-layered transshipment instance, we can assume that the vertices in U are sources (non-negative excess) and the vertices in W are sinks (non-positive excess). With these notions, we can define our transshipment problem:

2-LAYER GENERAL TRANSSHIPMENT (2GT)

Input: A 2-layered network D with parts U and W , arc capacities $\text{cap}: A(D) \rightarrow \mathbb{Q}_0^+$, and a collection $\mathcal{B} = \{B_v \subseteq \mathbb{Q}_0^+ \mid v \in U\} \cup \{B_v \subseteq \mathbb{Q}_0^- \mid v \in W\}$ of sets.

Task: Find a \mathcal{B} -transshipment f in D which maximizes $\text{val}(f)$.

For the rest of this work we assume the underlying undirected graph of D to be connected and therefore have $n_D - 1 \leq m_D$ and $O(m_D + n_D) = O(m_D)$. Clearly if this is not the case, we obtain a solution by solving each connected components individually and isolated vertices can be ignored. To familiarize ourselves with the problem, let us model an instance $\mathcal{I} = (G, k)$ of τ -BDED as a 2GT-instance; we denote this linear-time transformation by $T(\mathcal{I})$:

- The vertex set of the network D is $V(D) := V(G) \cup E(G)$.
- For each edge $e = \{u, v\} \in E(G)$, add the arcs (e, u) and (e, v) with capacity 1 to D .
- For each $e \in E(G)$, set $B_e := \{0, 1\}$, that is, a flow of either 0 or 1 is allowed to leave e .
- For each $v \in V(G)$, set $B_v := [-\tau, 0]$, that is, any flow up to τ is allowed to arrive in v .

Note that in the above construction a vertex $v \in V(G)$ is assigned an infinitely large set $[-\tau, 0]$ of allowed excess values. We henceforth assume that a reasonable encoding for the sets B_v is used (here, e.g. via the numbers $-\tau$ and 0 encoding the interval).

It is straightforward to verify the correctness of this reduction:

► **Lemma 4.1.** *Let $\mathcal{I} = (G, k)$ be a τ -BDED instance and let $T(\mathcal{I}) = (D, \text{cap}, \mathcal{B})$ be the constructed 2GT instance. Then there is a \mathcal{B} -transshipment f in D with value $\text{val}(f) \geq m_G - k$ if and only if there is an edge set $F \subseteq E$, $|F| \leq k$, with $\rho^*(G - F) \leq \tau$.*

Proof.

“ \Rightarrow ”: Suppose that f is a \mathcal{B} -transshipment for $(D, \text{cap}, \mathcal{B})$ with $\text{val}(f) \geq m_G - k$. We define $F \subseteq E(G)$ to be the set of edges where no flow is sent from the corresponding vertices in D , that is, $F := \{e \in E(G) \mid f_\Sigma(e) = 0\}$. Since $f_\Sigma(e) \in B_e = \{0, 1\}$ for all $e \in E(G)$ and $\text{val}(f) \geq m_G - k$, it follows that $|F| \leq k$.

Using the flow f , we now construct a fractional orientation ϕ for $G - F$ with $\Delta_\phi^- \leq \tau$. Lemma 2.1 then yields the claim. For any edge $e = \{u, v\} \in E(G) \setminus F$, consider the arcs (e, u) and (e, v) in $A(D)$. Set $\phi(e)^u := f(e, u)$ and $\phi(e)^v := f(e, v)$. It holds that $\phi(e)^u + \phi(e)^v = 1$, since $f_\Sigma(e) = 1$. For any vertex $v \in V(G)$, it holds that $\deg_\phi^-(v) \leq \tau$, because $f_\Sigma(v) \geq -\tau$ by the excess constraint $B_v = [-\tau, 0]$.

“ \Leftarrow ”: Now, let $F \subseteq E(G)$, $|F| \leq k$, with $\rho^*(G - F) \leq \tau$. By Lemma 2.2, there exists a fractional orientation ϕ on $G - F$ with $\Delta_{\phi}^- \leq \tau$. We construct a \mathcal{B} -transshipment f with $\text{val}(f) \geq m_G - k$ from ϕ : For each $e = \{u, v\} \in E(G) \setminus F$, set the flows $f(e, u) := \phi(e)^u$ and $f(e, v) := \phi(e)^v$, and for each $e = \{u, v\} \in F$, set $f(e, u) = f(e, v) := 0$. Consider the vertices in $V(D)$ representing the edges of $E(G)$. The above assignment ensures for any vertex $e \in E(G) \setminus F$, that $f_{\Sigma}(e) = 1$ and for any vertex $e \in F$ that $f_{\Sigma}(e) = 0$. Moreover, $f_{\Sigma}(v) \in [-\tau, 0]$ is guaranteed since $\Delta_{\phi}^- \leq \tau$. Finally, we have $\text{val}(f) = \sum_{e \in E \setminus F} 1 = m_G - |F| \geq m_G - k$. \blacktriangleleft

4.2 Solving 2-Layer General Transshipment via General Factors

We now consider how to solve 2GT using GENERAL FACTORS. One important step will be to limit the excess constraints of vertices in a 2GT instance to integral values. For this, we first consider conditions under which 2GT yields integral solutions, meaning solutions in which the flow assigned to any edge is integral.

► **Lemma 4.2.** *Let $(D, \text{cap}: A(D) \rightarrow \mathbb{N}, \mathcal{B})$ be a 2GT instance with $A(D) \subseteq U \times W$. Furthermore, for each $v \in V(D)$, let $B_v \subseteq \mathbb{Z}$ or $B_v = [c_v, d_v]$ for some $c_v, d_v \in \mathbb{Z}$ with $c_v \leq d_v$. If $(D, \text{cap}, \mathcal{B})$ has a \mathcal{B} -transshipment, then it attains an integral maximum \mathcal{B} -transshipment.*

Proof. The idea is to construct a flow instance with lower and upper bounds for the flow on each arc, then use a known integrality property and finally transfer the flow back to our 2GT instance. Let f be a maximum \mathcal{B} -transshipment for $(D, \text{cap}, \mathcal{B})$ (such f must exist if there is a feasible transshipment) and consider the flow instance \mathcal{I}' , where we take the network D together with a source vertex s , a sink vertex t and the following arcs with lower and upper bounds (denoted by intervals) on the flow assigned to them:

- For each arc $a \in A(D)$, we set the permitted flow to $[0, \text{cap}(a)]$.
- For every $u \in U$ with $f_{\Sigma}(u) > 0$, we add the arc (s, u) and distinguish between two cases. If $B_u = [c_u, d_u]$, then we set the permitted flow to $[c_u, d_u]$. Otherwise, we set the permitted flow to $[f_{\Sigma}(u), f_{\Sigma}(u)]$.
- For every $w \in W$ with $f_{\Sigma}(w) < 0$, we add the arc (w, t) and distinguish between two cases. If $B_w = [c_w, d_w]$, then we set the permitted flow to $[-d_w, -c_w]$ and otherwise to $[-f_{\Sigma}(w), -f_{\Sigma}(w)]$.

We now show that there is a feasible flow f' for the constructed instance, which then implies the existence of an integral maximum flow for \mathcal{I}' (see [21, Corollary 11.2e]). Let a be some arc in our constructed instance. We define f' by

$$f'(a) := \begin{cases} f_{\Sigma}(u) & \text{if } a = (s, u) \text{ for some } u \in U \\ f(a) & \text{if } a \in A(D) \\ -f_{\Sigma}(w) & \text{if } a = (w, t) \text{ for some } w \in W. \end{cases}$$

Note that $\text{val}(f) = \text{val}(f')$ since both flows send the same amount of flow from U to W . The feasibility of f' follows directly from our construction and the fact that f is a \mathcal{B} -transshipment. Now, let f'' be an integral maximum flow for \mathcal{I}' . We define an integral \mathcal{B} -transshipment f^* by $f^*(u, w) := f''(u, w)$ for all $(u, w) \in A(D)$.

First, we show that f^* is indeed a \mathcal{B} -transshipment. For this, note that for each $u \in U$, we have $f_{\Sigma}^*(u) = f''(s, u)$. We distinguish between two cases. If $B_u \subseteq \mathbb{Z}$, then the permitted flow on (s, u) is exactly $f_{\Sigma}(u)$ and therefore $f''(s, u) = f_{\Sigma}(u) \in B_u$. Otherwise, the permitted flow on (s, u) is $[c_u, d_u]$ which directly implies $f''(s, u) \in [c_u, d_u] = B_u$. For all vertices $w \in W$, we have $f_{\Sigma}^*(w) = -f''(w, t)$ and by analogous arguments as above, it follows $-f''(w, t) \in B_w$.

It remains to show that f^* is indeed a maximum flow. Recall that f is a maximum flow for $(D, \text{cap}, \mathcal{B})$ and f' with $\text{val}(f') = \text{val}(f)$ is a feasible solution for \mathcal{I}' . Further, f'' is a maximum integral flow for \mathcal{I}' , that is, $\text{val}(f') \leq \text{val}(f'')$. Finally, since $\text{val}(f^*) = \text{val}(f'')$, we have $\text{val}(f) = \text{val}(f') \leq \text{val}(f'') = \text{val}(f^*)$. \blacktriangleleft

Using the above lemma we now observe that any 2GT instance with an integral maximum transshipment can be restricted to integral capacities and integral excess values.

► Observation 4.3. *If a 2GT instance $(D, \text{cap}, \mathcal{B})$ has an integral maximum \mathcal{B} -transshipment, then it is equivalent to $(D, \text{cap}', \mathcal{B}')$ where $\mathcal{B}' := \{B_v \cap \mathbb{Z} \mid B_v \in \mathcal{B}\}$ and $\text{cap}'(a) := \lfloor \text{cap}(a) \rfloor$ for each arc $a \in A(D)$.*

With this observation we can now describe a reduction of integral 2GT instances to GENERAL FACTORS. This will later allow us to use known algorithms for GENERAL FACTORS to solve 2GT instances.

► Lemma 4.4. *Let $(D, \text{cap}: A(D) \rightarrow \mathbb{N}, \mathcal{B} = \{B_v \subseteq \mathbb{Z} \mid v \in V(D)\})$ be an integral 2GT instance, where all capacities are bounded polynomially in the size of D . Then, one can construct in polynomial time, a GENERAL FACTORS instance (G, \mathcal{B}') with the following properties:*

- (i) $(D, \text{cap}, \mathcal{B})$ has an integral \mathcal{B} -transshipment of value p if and only if (G, \mathcal{B}') has a factor of size $2p$.
- (ii) Given a tree decomposition for the underlying undirected graph of D of width tw , we can compute a tree decomposition of width $\max\{2, \text{tw}\}$ for G in polynomial time.
- (iii) The max-gap of \mathcal{B}' is $\max\{1, \gamma\}$, where γ is the max-gap of \mathcal{B} .

Proof. Let $(D, \text{cap}, \mathcal{B} = \{B_v \subseteq \mathbb{Z} \mid v \in V(D)\})$ be a given 2GT instance with $A(D) \subseteq U \times W$. Now, consider the multigraph, which is obtained by replacing every arc (u, w) with $\text{cap}(u, w)$ many undirected parallel edges between u and w . Moreover, set $\mathcal{B}' := \{\text{abs}(B_v) \mid B_v \in \mathcal{B}\}$, where $\text{abs}(B_v) := \{|x| \mid x \in B_v\}$. We can transform the multigraph into a simple graph by subdividing every edge and adding the degree constraint $\{0, 2\}$ to \mathcal{B}' for every newly created vertex. We refer to the set of new vertices by H and if necessary, we call the vertex, which was introduced for the i^{th} parallel edge between u and w , by h_{uw}^i . We call the obtained graph G and the constructed GENERAL FACTORS instance is (G, \mathcal{B}') . We next show property (i).

“ \Rightarrow ”: Let f be an integral \mathcal{B} -transshipment for $(D, \text{cap}, \mathcal{B})$. We show that there is a general factor $F \subseteq E(G)$ for the instance (G, \mathcal{B}') of size $|F| = 2 \text{val}(f)$. Intuitively, if the flow value on the arc (u, w) is $f(u, w)$, then we add $2f(u, w)$ edges from G to F , such that they correspond to $f(u, w)$ copies of the edge $\{u, w\}$. Formally, we set

$$F := \bigcup_{(uw) \in A(D)} \bigcup_{i=1}^{f(u,w)} \{\{u, h_{uw}^i\}, \{h_{uw}^i, w\}\}.$$

Since every flow from U to W counts towards $\text{val}(f)$, it directly follows that $|F| = 2 \text{val}(f)$. It remains to show that F satisfies the degree constraints \mathcal{B}' , that is, $\deg_{(V(G), F)}(v) \in B_v$ for every $v \in V(G)$. Recall, that for $h \in H$, we have $B_h = \{0, 2\}$. By construction, we either add zero or two edges incident to h to F and therefore we have $\deg_{(V(G), F)}(h) \in B_h$. For $v \in U \cup W$, note that the degree in $(V(G), F)$ coincides exactly with the absolute value of $f_\Sigma(v)$. This together with $f_\Sigma(v) \in B_v$ implies that $\deg_{(V(G), F)}(v) \in \text{abs}(B_v) \in \mathcal{B}'$.

“ \Leftarrow ”: Now suppose $F \subseteq E(G)$ is a \mathcal{B}' -factor. We show that there exists a \mathcal{B} -transshipment f with $\text{val}(f) = \lfloor \frac{|F|}{2} \rfloor$. Let (u, w) be an arc in D . By construction, if $\{u, h_{uw}^i\} \in F$ for some $h_{uw}^i \in H$, then also $\{h_{uw}^i, w\} \in F$, because $B_{h_{uw}^i} = \{0, 2\}$. Now, we define a flow f by setting $f(u, w)$ to the number of vertices from H to which both u and w are connected in $(V(G), F)$. This is an integral feasible flow since for every arc $(u, w) \in A(D)$ both u and w have by construction at most $\text{cap}(u, w)$ common neighbors in G . Moreover, for each $u \in U$, we have $f_\Sigma(u) = \deg_{(V(G), F)}(u) \in \text{abs}(B_u) = B_u$ and for each $w \in W$, we have $f_\Sigma(w) = -\deg_{(V(G), F)}(w) \in \{-x \mid x \in \text{abs}(B_w)\} = B_w$. Since exactly $\lfloor \frac{|F|}{2} \rfloor$ edges are between U and H and also between H and W , it follows $\text{val}(f) = \lfloor \frac{|F|}{2} \rfloor$.

It remains to show the two properties (ii) and (iii).

- (ii): Let $(\mathcal{T}, (X_i)_{i \in V(\mathcal{T})})$ be a tree-decomposition of width tw for the underlying undirected graph of D . For each arc $(u, w) \in A(D)$, there exists a bag X_i with $\{u, w\} \subseteq X_i$. We obtain a tree decomposition for G of width $\max\{2, \text{tw}\}$ by appending one bag $\{u, w, h_{uw}^j\}$ to X_i for each $j \in [\text{cap}(u, w)]$.
- (iii): Note that, for any vertex $h \in H$, the set B_h has a gap of exactly one. For $v \in U \cup W$, the collection \mathcal{B}' contains the set $\text{abs}(B_v)$ which has the same gap as B_v since either $B_v \subseteq \mathbb{Q}_0^+$ or $B_v \subseteq \mathbb{Q}_0^-$. \blacktriangleleft

To sum up, we have seen that a 2GT instance allowing flows in an interval with integer endpoints can be restricted to an integral instance and that integral 2GT instances can be translated into equivalent GENERAL FACTORS instances. For non-integral 2GT instances, the reduction to GENERAL FACTORS does not work. A common trick we will use later, is to scale capacities and excess constraints by some factor in order to make them integral. This scaling obviously influences the overall flow value.

► **Lemma 4.5.** *Let $(D, \text{cap}, \mathcal{B})$ be a 2GT instance, $q \in \mathbb{Q}_0^+$ and $(D, \text{cap}', \mathcal{B}')$ be the instance, which is obtained by multiplying all excess constraints and capacities from $(D, \text{cap}, \mathcal{B})$ with q . Then, $(D, \text{cap}, \mathcal{B})$ has a maximum \mathcal{B} -transshipment value of p if and only if $(D, \text{cap}', \mathcal{B}')$ has a maximum \mathcal{B}' -transshipment value of pq .*

Proof. Clearly, since all capacities and excess values are multiplied by q , it follows that f is a maximum \mathcal{B} -transshipment if and only if f' with $f'(u, w) := qf(u, w)$ for all $(u, w) \in A(D)$ is a maximum \mathcal{B}' -transshipment. \blacktriangleleft

We write $q\mathcal{B}$ for \mathcal{B}' to indicate that \mathcal{B}' was obtained from \mathcal{B} by scaling with factor q .

4.3 Algorithmic Consequences

2-Layer General Transshipment. Having discussed how to translate a 2GT instance into a GENERAL FACTORS instance, we now consider the algorithmic implications for the 2GT problem. We obtain three algorithmic results as stated in the following theorem.

- **Theorem 4.6.** *Let $(D, \text{cap}: A(D) \rightarrow \mathbb{N}, \mathcal{B})$ be an instance of 2GT.*
 - (i) *If $B_v = [a_v, b_v] \cap \mathbb{Z}$ for each vertex v with $a_v \leq 0 \leq b_v$ and the excess constraints and capacities are polynomially bounded in the size of D , then one can compute in randomized $O(m^{1+o(1)})$ time a maximum \mathcal{B} -transshipment.*
 - (ii) *If $B_v \subseteq \mathbb{Z}$ has gap at most one for each vertex v and capacities are bounded polynomially, then one can compute in polynomial time a maximum \mathcal{B} -transshipment or output that there exist no feasible flow.*

- (iii) Suppose we are given a tree decomposition for the underlying undirected graph of D of width tw . Moreover, suppose that $B_v \subseteq \mathbb{Z}$ for each vertex v and that capacities are bounded polynomially. Let $M := \max_{v \in V(D)} \max\{|x| \mid x \in B_v\}$. Then, one can compute in $(M + 1)^{\text{tw} + 1} n^{O(1)}$ time a maximum \mathcal{B} -transshipment or output that there exists no feasible flow.

Proof. We prove each of the three statements separately.

- (i): We provide a linear-time reduction to MAXIMUM FLOW. If the edge capacities are integers and polynomially bounded, then an (integral) maximum flow for this instance can be computed in randomized $m^{1+o(1)}$ time [7]. Now consider the following MAXIMUM FLOW instance, where we take the network D , add a source vertex s , a sink vertex t and the following arcs:

- For every $u \in U$, we add the arc (s, u) with capacity b_u .
- For every $w \in W$, we add the arc (w, t) with capacity $-a_w$.

This network can be constructed in $O(m_D + n_D) = O(m_D)$ time. We now compute a maximum flow f for this network and obtain a \mathcal{B} -transshipment by setting $f^*(u, w) := f(u, w)$ for all $(u, w) \in A(D)$. Note that f^* is feasible since $f_\Sigma^*(u) = f_\Sigma^+(u) \in [0, b_u]$ for all $u \in U$ and $f_\Sigma^*(w) = -f_\Sigma^-(w) \in [a_w, 0]$ for all $w \in W$. Furthermore, f^* is a maximum flow since $\text{val}(f^*) = \text{val}(f)$ and any flow in the transshipment instance can be directly translated to an s - t -flow f' of the same value as follows:

$$f'(a) := \begin{cases} f_\Sigma^*(u) & \text{if } a = (s, u) \text{ for some } u \in U \\ f^*(a) & \text{if } a \in A(D) \\ -f_\Sigma^*(w) & \text{if } a = (w, t) \text{ for some } w \in W. \end{cases}$$

- (ii): We apply Lemma 4.4 and obtain a GENERAL FACTORS instance (G, \mathcal{B}') which has gap exactly one by Lemma 4.4 (iii). Using an algorithm from Dudycz and Paluch [10] we can now find a maximum \mathcal{B}' -factor F of G in polynomial time. In particular, Lemma 4.4 gives us a \mathcal{B} -transshipment of value $\frac{|F|}{2}$ and implies that this is indeed the maximum. On the other hand, if there is no solution to (G, \mathcal{B}') , then there is also no feasible flow for $(D, \text{cap}, \mathcal{B})$.
- (iii): We again apply Lemma 4.4 and obtain a GENERAL FACTORS instance (G, \mathcal{B}') . By Lemma 4.4 (ii), we get a tree decomposition for G of width at most $\text{tw} + 1$ and, by construction of the GENERAL FACTORS instance, we know that $\max\{\max B' \mid B' \in \mathcal{B}'\} \leq M$ (note that we can assume $M \geq 2$ since otherwise we are in case (ii) which is polynomial-time solvable). This allows us to apply a result from Marx et al. [17], which states that a maximum \mathcal{B}' -factor of G can be computed in $(M + 1)^{\text{tw} + 1} n^{O(1)}$ time. If there is no solution to (G, \mathcal{B}') , then there is also no feasible flow for $(D, \text{cap}, \mathcal{B})$. ◀

τ -Bounded-Density Edge Deletion. In Section 4.1 we described how to reduce τ -BDED to 2GT. Therefore, the results from Theorem 4.6 for 2GT translate directly to τ -BDED and imply the algorithmic results in Theorems 1.1 and 1.2:

► **Theorem 4.7.** Let $\mathcal{I} = (G, k)$ be a τ -BOUNDED-DENSITY EDGE DELETION instance.

- (i) If $\tau \in \mathbb{N}$, then one can decide \mathcal{I} in randomized $O(m^{1+o(1)})$ time.
- (ii) If $2\tau \in \mathbb{N}$, then one can decide \mathcal{I} in polynomial time.
- (iii) If $\tau = \frac{a}{b}$ for some fixed $a, b \in \mathbb{N}$ and a tree decomposition for G of width tw is given, then one can decide \mathcal{I} in $(\max\{a, b\} + 1)^{\text{tw} + 1} n^{O(1)}$ time.

Proof. Recall the construction of a 2GT instance $T(\mathcal{I}) = (D, \text{cap}, \mathcal{B})$ from Section 4.1, which has maximum \mathcal{B} -transshipment of value at least $m_G - k$ if and only if (G, k, τ) is a yes-instance (see Lemma 4.1). The capacities on the arcs in D are one and the excess constraints are of the form $\{0, 1\}$ and $[-\tau, 0]$. Since τ is bounded polynomially in n_D , clearly so are all the capacities and excess constraints. This also holds when scaling the instance with a constant, as described in Lemma 4.5. Using the construction, we can derive the three statements.

- (i): First, note that $T(\mathcal{I})$ can be constructed in $O(m_G + n_G) = O(m_G)$ time. The excess constraints are of the form $\{0, 1\}$ and $[-\tau, 0]$. Since $\tau \in \mathbb{N}$ and a feasible solution exists (e.g. sending a flow of 0 along all arcs), we can apply Lemma 4.2 which guarantees us the existence of an integral maximum \mathcal{B} -transshipment. Subsequently, Observation 4.3 tells us that we can restrict all the excess constraints to their intersection with \mathbb{Z} . This allows us to apply Theorem 4.6 (i), which concludes the algorithm.
- (ii): By Lemma 4.5, we can scale the 2GT instance $T(\mathcal{I})$ by the factor two and the obtained excess constraints are of the form $\{0, 2\}$ and $[-2\tau, 0]$. Since 2τ is integer and a feasible solution exists, we can again apply Lemma 4.2 which gives us the existence of an integral maximum flow. Again, Observation 4.3 tells us that we can restrict all the excess constraints to their intersection with \mathbb{Z} . Note that the excess constraints are now of the form $\{0, 2\}$ and $[-2\tau, 0] \cap \mathbb{Z}$. In particular, they have gap at most one and therefore the obtained transshipment instance satisfies the conditions of Theorem 4.6 (ii). Since we scaled $T(\mathcal{I})$ by two, the maximum $2\mathcal{B}$ -transshipment value is at least $2(m_G - k)$ if and only if (G, k, τ) is a yes-instance.
- (iii): By Lemma 4.5, we can scale the 2GT instance $T(\mathcal{I})$ by the factor b and the obtained excess constraints are of the form $\{0, b\}$ and $[-b\tau, 0] = [-a, 0]$. This scaled version has a maximum $b\mathcal{B}$ -transshipment of value at least $b(m_G - k)$ if and only if the original instance has a maximum flow of at least $m_G - k$. By the same arguments as above, we can restrict the excess constraints to their intersections with \mathbb{Z} . Furthermore, we easily obtain a tree decomposition for the underlying undirected graph of D of width tw since the graph can be obtained by subdividing all edges of G [9]. To finish the proof, note that we can now apply Theorem 4.6 (iii) to compute a maximum transshipment in $(\max\{a, b\} + 1)^{tw+1} n^{O(1)}$ time. We then return yes if and only if the value is at least $b(m_G - k)$. \blacktriangleleft

5 NP-hardness Results

In this section, we complement our algorithmic findings from Section 4 by showing NP-hardness for every constant target density $\tau \in \{q \in \mathbb{Q} \mid q \geq 2/3 \wedge q \notin \mathbb{H}\}$, thereby completing the proof of Theorem 1.1. We prove the result in two steps. First, we show that any fixed target density $\tau \in [2/3, 1)$ yields NP-hardness. Afterwards, we show that τ -BDED is NP-hard for any constant target density $\tau \in \{q \in \mathbb{Q} \mid q > 1 \wedge q \notin \mathbb{H}\}$.

Case $\tau < 1$. Note that for $\tau < 1$ any solution graph (the graph obtained by removing at most k edges) cannot contain a cycle as a cycle has density 1. Hence, any solution graph is a forest. The maximum density of a forest is determined by the size of a largest tree in it as a tree with n nodes has $n - 1$ edges and thus density $1 - 1/n$, which is monotonically increasing for increasing values of n . Thus, any target density $\tau < 1$ can be replaced by the largest value $\tau' \leq \tau$ with $\tau' = \ell/\ell+1$ for some $\ell \in \mathbb{N}$. For $\ell = 0$, we have to remove all edges and thus τ -BDED becomes trivial. For $\ell = 1$, the problem is equivalent to finding a smallest set of




Figure 2 A schematic illustration of the reduction in Theorem 5.1. Given an instance of $X\ell C$ with universe $X = \{x_1, x_2, \dots, x_{\ell q}\}$ and set collection $C = \{C_1, C_2, \dots, C_t\}$, we add a vertex y_i for each element $x_i \in X$ and a vertex c_j for each set $C_j \in C$. We add an edge $\{y_i, c_j\}$ whenever $x_i \in C_j$. We further add dummy stars with $\ell - 1$ leaves and fully connect their centers to the vertices representing C . Any τ -BDED solution then leaves t stars with ℓ leaves and there is a one-to-one mapping between (green) star centers and sets selected in an $X\ell C$ solution.

edges that only leaves a matching. This can be solved in polynomial time by computing a maximum matching and returning all non-matching edges. For $\ell \in \{2, 3\}$, Bazgan et al. [3] showed NP-hardness. We extend their hardness result for $\ell = 3$ to arbitrary ℓ . The proof consists of a reduction from EXACT COVER BY ℓ -SETS ($X\ell C$) and is sketched in Figure 2.

► Theorem 5.1. τ -BDED is NP-hard for any $\tau = \ell/\ell+1$ with $\ell \geq 3$.

Proof. We reduce from EXACT COVER BY ℓ -SETS, which is NP-hard for $\ell \geq 3$ [13]. It is defined as follows.

EXACT COVER BY ℓ -SETS ($X\ell C$)

Input: A set X with $|X| = \ell q$ and a collection C of ℓ -element subsets of X .

Question: Is there a subcollection C' of C where every element of X occurs in exactly one set in C' ?

The proof is a straightforward adaption of the hardness proof for $\tau = 3/4$ [3]; see Figure 2 for a sketch. Let (X, C) be an instance of $X\ell C$ with $X = \{x_1, x_2, \dots, x_{\ell q}\}$ and $C = \{C_1, C_2, \dots, C_t\}$. If $t < q$, then we trivially have a no-instance as not all elements appear in a set in C . So we assume $t \geq q$ and construct a graph G as follows: We start with one vertex y_i for each element $x_i \in X$ and one vertex c_j for each set $C_j \in C$. We add an edge $\{y_i, c_j\}$ whenever $x_i \in C_j$. We then add a set of $t - q$ stars S_1, S_2, \dots, S_{t-q} each with $\ell - 1$ leaves to G . Formally, for each $i \in [t - q]$, we add a vertex u_i and for each $i \in [t - q]$ and each $j \in [\ell - 1]$, we add a vertex z_i^j and the edge $\{u_i, z_i^j\}$. Next, we add all edges $\{c_j, u_i\}$ with $j \in [t]$ and $i \in [t - q]$. This concludes the construction of G . Let m be the number of edges in it. To conclude the entire construction, we set $k := m - \ell t$.

Since the construction takes polynomial time, it only remains to show correctness. To this end, first assume that there is a solution C' for the input instance (X, C) of $X\ell C$. Note that $|C'| = q$. We construct a set B of ℓt edges such that deleting all edges except for those in B is a solution for the constructed instance of τ -BDED. For each $C_j \in C'$, include the ℓ edges joining c_j to vertices y_i with $x_i \in C_j$. Note that there are exactly ℓq such edges. Moreover, for each $i \in [t - q]$ and each $j \in [\ell - 1]$, we add the edge $\{u_i, z_i^j\}$ to B . Finally, for each $C_j \notin C'$, we pick a distinct vertex u_i and include $\{u_i, c_j\}$ in B . Note that there are exactly $t - q$ sets in $C \setminus C'$ and also $t - q$ vertices u_i . Therefore, B contains a perfect matching

between these two sets of vertices. Overall, B has size $\ell q + (\ell - 1)(t - q) + (t - q) = \ell t$. Note that each vertex c_j belongs to a different connected component in the graph induced by B . Moreover, each such vertex belongs to a tree with ℓ additional vertices (either a star S_j or ℓ y -vertices). Thus, the maximum density is $\ell/\ell+1$, concluding the first part of the proof.

For the other direction, assume that there is a set F of at most k edges whose removal yields a graph with maximum density at most $\tau = \ell/\ell+1$. Note that the constructed graph G has $(\ell + 1)t$ vertices. Moreover, $G - F$ has at least $m - |F| \geq m - k = \ell t$ edges. Since any cycle has density 1, $G - F$ is a forest and since any tree with $\ell + 2$ vertices has $\ell + 1$ edges and therefore density $\ell+1/\ell+2 > \ell/\ell+1$, it holds that each tree in $G - F$ contains at most $\ell + 1$ vertices. If $G - F$ contains exactly t trees, each with $\ell + 1$ vertices and ℓ edges, then $G - F$ contains ℓt edges and F contains $m - \ell t = k$ edges. Since G is connected (assuming each element appears in at least one set in C), increasing the number of trees in $G - F$ or decreasing the size of any tree in $G - F$ increases the size of F to above k . Thus, $G - F$ contains exactly t trees and each tree contains $\ell + 1$ vertices and ℓ edges. This implies that no edge $\{u_i, z_i^j\}$ for any $i \in [t - q]$ and any $j \in [\ell - 1]$ is contained in F . Moreover, each vertex u_j is adjacent to exactly ℓ vertices in $G - F$ and thus exactly one edge $\{u_i, c_j\}$ is not contained in F for each $i \in [t - q]$ (and the index j is unique for each i). Consider the remaining q vertices c_j that are not connected to a star S_i . We claim that the corresponding sets C_j are a solution for the original instance of X ℓ C. Consider an arbitrary vertex y_i (corresponding to an element $x_i \in X$). By construction, y_i is only adjacent to vertices c_j with $x_i \in C_j$. Since each tree in $G - F$ contains exactly $\ell + 1$ vertices and only q vertices c_j and ℓq vertices y_i are not yet distributed, each such undistributed vertex c_j must be adjacent to ℓ vertices y_i each and each vertex y_i is only adjacent to a single vertex c_j in $G - F$. Note that this exactly corresponds to an exact cover, concluding the proof. \blacktriangleleft

Case $\tau > 1$. Our goal is to show that τ -BDED is NP-hard for any constant $\tau = p/q > 1$ with $\tau \notin \mathbb{H}$. The proof will be a reduction from VERTEX COVER on q -regular graphs which mostly consists of two types of gadgets: *vertex gadgets* and *edge gadgets*. On a high level, the vertex gadgets are balanced graphs with density τ . The edge gadgets have density slightly below τ , but there are also two edges connecting each edge gadget to the two corresponding vertex gadgets. Combined, an edge gadget and the two corresponding vertex gadgets have density slightly larger than τ (in such a way that $(m^{-1/q})/n = \tau$). The idea is that deleting one edge in a vertex gadget is sufficient to “fix” all adjacent edge gadgets.

The vertex gadget will be based on the following lemma. We state it in a more general form which can be applied whenever a solution graph contains a cycle. It states that in a cycle, we can always assume that at least one edge in a fractional orientation is fully assigned to one of its endpoints. The idea is to simply shift the fractional orientation around the cycle until this is no longer possible in which case one of the edges fully assigns its fractional orientation to one endpoint.

► Lemma 5.2. *Let G be a graph containing a cycle C and let ϕ be a fractional orientation of G . Then, there exists a fractional orientation ϕ' with $\deg_{\phi'}^-(v) = \deg_{\phi}^-(v)$ for all $v \in V(G)$ and an edge $e = \{u, v\}$ in C with $\phi'(e)^u = 1$, that is, e is completely assigned to one of its endpoints by the fractional orientation ϕ' . Given G , C , and ϕ , then ϕ' can be found in linear time with respect to the size of G .*

Proof. Let $C = (v_1, \dots, v_\ell, v_{\ell+1} = v_1)$. If there already exists an edge $\{v_i, v_{i+1}\}$ in C with $\phi(v_i v_{i+1})^{v_i} = 1$ or $\phi(v_i v_{i+1})^{v_{i+1}} = 1$, then we are done. Otherwise, let $\varepsilon_i := \phi(v_i v_{i+1})^{v_i}$ for each $i \in [\ell]$, $\varepsilon := \min_{i \in [\ell]} \varepsilon_i$ and let j be an index such that $\varepsilon = \varepsilon_j$. We obtain the desired

fractional orientation ϕ' from ϕ by setting $\phi'(e) = \phi(e)$ for all edges in G that are not $\{v_i, v_{i+1}\}$ for some $i \in [\ell]$, $\phi'(v_i v_{i+1})^{v_{i+1}} := \phi(v_i v_{i+1})^{v_{i+1}} + \varepsilon$ and $\phi'(v_i v_{i+1})^{v_i} = \phi(v_i v_{i+1})^{v_i} - \varepsilon$ for all $i \in [\ell]$. Note that this is possible by the definition of ε and $\phi'(v_j v_{j+1})^{v_{j+1}} = 1$. Moreover, $\deg_{\phi'}^-(v) = \deg_{\phi}^-(v)$ for all $v \in V(G)$. Given G , C and ϕ , clearly ϕ' can be determined in linear time with respect to the size of G considering edges in the cycle at most twice and all other edges at most once. \blacktriangleleft

For the edge gadget, the basic idea is to start with a path adding additional edges in regular intervals to construct a near balanced graph with two orientations such that all vertices except for the two endpoints of the path have indegree τ and one endpoint has indegree $\tau - 1$ and the other has indegree $\tau - 1 + 1/q$. Connecting each endpoint of the path via a single edge with the corresponding vertex gadget then ensures that at least one edge in one of the three gadgets has to be deleted. Recall that we want to encode that a vertex is part of a vertex cover if an edge is removed in the corresponding vertex gadget. The excess value of $1/q$ will be small enough to distribute the available capacity generated by a single edge deletion to all incident edge gadgets.

► **Lemma 5.3.** *For every constant rational number $\tau = \frac{p}{q} > 1$, where p and q are coprime, a graph W with two special vertices u and w can be constructed in constant time, such that there exist two fractional orientations ϕ_1 and ϕ_2 with the following properties:*


- (i) $\deg_{\phi_1}^-(v) = \deg_{\phi_2}^-(v) = \tau$ for all $v \in V(W) \setminus \{u, w\}$,
- (ii) $\deg_{\phi_1}^-(u) = \deg_{\phi_2}^-(w) = \tau - 1 + \frac{1}{q}$ and
- (iii) $\deg_{\phi_1}^-(w) = \deg_{\phi_2}^-(u) = \tau - 1$.

Proof. Let $p = cq + r$ where c and r are integers and $0 \leq r < q$. Note that $r \geq 1$ since p and q are coprime and $c \geq 1$ since $\tau > 1$. Let $p' := p - (c-1)q = q + r$, and $\tau' := p'/q$. Note that $1 < \tau' < 2$ and that p' and q are coprime. We construct the graph W in two stages. First, we construct a graph W' for target density τ' and then we use this construction to construct the graph W for target density τ .

Note that if the graph W' has n vertices and fulfills the requirements for the fractional orientations (for the target density τ'), then the number m of edges in it fulfills $m = \tau'n - 2 + \frac{1}{q}$. This number is not necessarily an integer, but since p' and q are coprime, we can choose n such that $(p' \cdot n) \bmod q = q - 1$ (see for example Burton [4] for more details). Then, $m = \tau'n - 2 + \frac{1}{q} = \frac{p'n+1-2q}{q}$ is an integer. Moreover, we can increase n by aq for any positive integer a and this increases m by ap' and thus m remains an integer. Since $2q > p' > q$ as $1 < \tau' < 2$, this allows us to pick² a positive integer $n \geq 4$ such that m is also a positive integer with $n + 1 \leq m \leq 2n - 3$.

To construct W' with n vertices and m edges, we start with a path $(u = v_1, \dots, v_n = w)$ and iteratively add edges to the construction while simultaneously ensuring the existence of the desired fractional orientation ϕ'_1 . We will later show how to modify ϕ'_1 to get ϕ'_2 . First, we orient the edge $\{u, v_2\}$ with a value of $\tau' - 1 + 1/q$ to u and with a value of $1 - (\tau' - 1 + 1/q) = 2 - \frac{p'-1}{q}$ to v_2 . This is possible since $1 \leq \tau' = \frac{p'}{q} < \frac{p'+1}{q} \leq 2$. Now starting from $i = 2$ and for increasing values, we orient the edge $\{v_i, v_{i+1}\}$ in such a

² Use the extended Euclidean algorithm to first get an integer t , such that $1 \equiv tp' \pmod{q}$. Setting $n' = ((q-1)t) \bmod q$ satisfies $(n'p') \bmod q = ((q-1)tp') \bmod q = q - 1$. Let $m' = \frac{p'}{q}n' - 2 + \frac{1}{q}$. If $n' + 1 \leq m'$ (implying $n' \geq 4$), set $n = n'$, $m = m'$, otherwise iteratively increment n' by q until the condition holds. Since every increase of n' by q increases m' by p' with $p' < q$ and since initially $n' < q$, this procedure terminates after at most q steps.



■ **Figure 3** Illustration of the edge gadget described in Lemma 5.3 with fractional orientation ϕ_1 on the top and ϕ_2 on the bottom. Vertices $\{u = v_1, \dots, v_j, \dots, v_n = w\}$ form a path to which edges are added to ensure an indegree of τ' w.r.t. ϕ_1 for all vertices except u and w (which have an indegree of $\tau' - 1 + \frac{1}{q}$ and $\tau' - 1$ respectively). Vertex v_j is the first vertex which requires an additional edge to reach τ' and x is the amount of fractional orientation it receives from v_{j-1} . To create ϕ_2 from ϕ_1 , fractional orientation is “pushed” from u to w , flipping their indegrees.

way, that v_i gets indegree exactly τ' from the already oriented edges until we reach a point where even assigning a value of 1 is not enough. In the latter case, we add the edge $\{v_i, w\}$ to the graph and orient it completely to v_i . We then continue with the edge $\{v_i, v_{i+1}\}$ as before but now, whenever a value of 1 is not enough, we add the edge $\{u, v_i\}$ and completely orient it towards v_i . See Figure 3 for an illustration.

Note that since $m \geq n + 1$ (we will next show that the constructed graph has exactly m edges) and $n - 1$ edges belong to the initial path P , we add at least two edges in that way and hence the edges we try to add are not part of the path already (and we never try to add two edges to the same vertex as $\tau' < 2$). Moreover, since $m \leq 2n - 3$ and $n - 1$ edges belong to the initial path, we only add at most $n - 2$ edges and hence there are enough vertices to add sufficiently many edges. We next show that after doing the above for all $2 \leq i \leq n - 1$, the graph contains exactly m edges and ϕ'_1 fulfills all requirements from the lemma statement (for target density τ'). By construction, the indegree of u with respect to ϕ'_1 is $\tau' - 1 + \frac{1}{q}$ and the indegree of each vertex v_i with $2 \leq i \leq n - 1$ is τ' . Added over all vertices, the sum of in-degrees is $(n - 1)\tau' - 1 + \frac{1}{q}$. The indegree of w is some value t between 0 and 1 (only the edge $\{v_{n-1}, w\}$ assigns some positive value to w). Since the sum of indegrees equals the number of edges, it holds that $(n - 1)\tau' - 1 + \frac{1}{q} + t \in \mathbb{N}$. Note that this is the case for $t = \tau' - 1$ since then, $(n - 1)\tau' - 1 + \frac{1}{q} + t = n\tau' - 2 + \frac{1}{q} = \frac{\tau' n + 1}{q} - 2$, which is an integer by the choice of n . Moreover, $t = \tau' - 1$ is the only value between 0 and 1 for which this is the case as increasing/decreasing the number of edges by one also increases or decreases t by 1. Hence, the number of edges in the constructed graph is $n\tau' - 2 + \frac{1}{q} = m$, and $t = \tau' - 1$. Thus, the orientation ϕ'_1 fulfills all requirements of the lemma statement.

It remains to show the existence of the fractional orientation ϕ'_2 and to adapt the construction for target densities larger than 2. Towards the former, let j be the index where we added the edge $\{v_j, w\}$. Note that this index exists as $m \geq n$. Intuitively, we modify ϕ'_1

by sending an orientation value of $1/q$ from u to w along the path (u, v_2, \dots, v_j, w) . Note that by construction all orientations in ϕ'_1 are multiples of $1/q$. To construct ϕ'_2 from ϕ'_1 more formally, we set (for all $i \in [2, j-1]$)

$$\begin{aligned} \phi'_2(uv_2)^u &:= \phi'_1(uv_2)^u - \frac{1}{q}, & \phi'_2(uv_2)^{v_2} &:= \phi'_1(uv_2)^{v_2} + \frac{1}{q}, \\ \phi'_2(v_i v_{i+1})^{v_i} &:= \phi'_1(v_i v_{i+1})^{v_i} - \frac{1}{q}, & \phi'_2(v_i v_{i+1})^{v_{i+1}} &:= \phi'_1(v_i v_{i+1})^{v_{i+1}} + \frac{1}{q}, \\ \phi'_2(v_j w)^{v_j} &:= \phi'_1(v_j w)^{v_j} - \frac{1}{q}, \text{ and} & \phi'_2(v_j w)^w &:= \phi'_1(v_j w)^w + \frac{1}{q}. \end{aligned}$$

Note that if none of the above edges assign a value of 1 to the “later” vertex (the one with higher index), then the above modification is possible and ϕ'_2 also fulfills all requirements from the lemma statement. Assume towards a contradiction that some edge $\{v_i, v_{i'}\}$ assigns a value of 1 to the later vertex. Let $i < i'$ without loss of generality and note that by construction $i' \neq n$ as the edge $\{v_j, w\}$ is fully assigned to v_j and $i \neq 1$ as $\tau' > 1$ (and hence u is assigned a positive value $\tau' - 1 + 1/q > 1/q$). Hence, $i' = i + 1$ and $i < j$. This implies that v_i is assigned a value of at most $1 < \tau'$ by ϕ'_1 as it is only incident to $\{v_{i-1}, v_i\}$ and $\{v_i, v_{i+1}\}$ as $i < j$ and j is the first index where we added an additional edge. Thus, the assumption that $\{v_i, v_{i'}\}$ is fully oriented towards $v_{i'}$ leads to a contradiction. This shows that the above modification is possible and ϕ'_2 exists as claimed.

Finally, we adapt the construction to increase the density from $\tau' = (q+r)/q$ by $c-1$ to $(q+r)/q + c - 1 = (cq+r)/q = p/q = \tau$. To this end, we add a balanced graph A with density τ (which always exists and Ruciński and Vince [20] show how to construct such a graph for each $\tau \geq 1$). Note that since any graph with n vertices has density at most $\binom{n}{2}/n = (n-1)/2 < n$, it holds that A contains at least $c-1$ vertices (but the size of A is constant for each constant τ). We select a set D of $c-1$ arbitrary vertices in A . To conclude the construction of the edge gadget, we add an edge between each vertex in W' and each vertex in D . To construct the orientations ϕ_1 and ϕ_2 for target density τ , we start with an orientation for A where each vertex is assigned exactly τ . Note that this exists as A is balanced and has density τ . For the edges in W' , we use the orientations ϕ'_1 for ϕ_1 and ϕ'_2 for ϕ_2 , respectively. Finally, we orient all edges between W' and D completely towards the vertices in W' . It is now easy to verify that all requirements of the lemma are satisfied by ϕ_1 and ϕ_2 . Clearly if p and q are constant, the construction takes constant time. ◀


We are now ready to prove the main result of this section. This will conclude the proof of Theorem 1.1, which then follows from Theorem 4.7 (for the algorithmic part) and Theorems 5.1 and 5.4 (for the NP-hardness).

► **Theorem 5.4.** τ -BDED is NP-hard for any $\tau > 1$ with $\tau \notin \mathbb{H}$.

Proof. Let $\tau = p/q > 1$ where p and q are coprime. Note that since $\tau \notin \mathbb{H}$, it holds that $q \geq 3$. We give a reduction from VERTEX COVER on q -regular graphs. It is folklore knowledge that this problem is NP-complete.³ The idea is to construct a vertex gadget for each vertex and an edge gadget for each edge in G using Lemmas 5.2 and 5.3, respectively.

Let $(G = (V, E), k)$ be an instance of VERTEX COVER where G is q -regular. We construct an equivalent instance (G', k') of τ -BDED as follows. For each vertex $v \in V$, we add a vertex gadget H_v with a special edge e_v and an incident vertex x_v to G' . The graph H_v is any

³ Amiri [2] showed that the problem cannot be solved in $2^{o(m)}$ time unless the ETH fails. The same reduction also shows NP-hardness, but this is not explicitly mentioned.



■ **Figure 4** Illustration of the hardness reduction for $\tau = p/q > 1$. Depicted is the construction for two vertices u, v in the VERTEX COVER instance connected by an edge $e = \{u, v\}$ with $\tau = \frac{7}{5}$. To the left and right are vertex gadgets H_u and H_v (highlighted in green) with special vertices x_u and x_v and the special edges e_u and e_v oriented fully towards x_u and x_v respectively (indicated with the thickening edge). The vertex gadgets are connected by the edge gadget W_e with special vertices w_u and w_v . Light red lines indicate edges from x_u and x_v to other edge gadgets. Note that for $\tau = \frac{7}{5}$, a balanced graph on 5 vertices with 7 edges can be created using the construction of Ruciński and Vince [20]. For $p > \binom{q}{2}$, observe that a graph on pq vertices and p^2 edges (of density $\tau = p/q$) can be created.

balanced graph with density τ (which always exists and Ruciński and Vince [20] show how to construct such a graph for each $\tau \geq 1$). Let ϕ_v be any fractional orientation for H_v where each vertex has indegree τ and let C be any cycle in H_v . Both can be found in polynomial time. Using Lemma 5.2, we construct another fractional orientation ϕ'_v where each vertex has indegree τ and one edge e fully assigns its orientation to one of its endpoints u . We set $e_v := e$ and $x_v := u$. This concludes the construction of the vertex gadget. For each edge $e = \{u, v\} \in E$, we add the edge gadget W_e , which is the graph obtained from the construction in Lemma 5.3 for density τ with special vertices w_u and w_v . We connect the gadgets by adding the edges $\{x_u, w_u\}$ and $\{x_v, w_v\}$ to our construction. Note that we can use Lemma 5.3 since we assumed that p and q are coprime. To conclude the construction, we set $k' := k$. A visualization of the reduction is shown in Figure 4.

Since the reduction can be computed in polynomial time as the graphs H_v and W_e have constant size, it only remains to show correctness. To this end, first assume that (G, k) is a yes-instance of VERTEX COVER and let C be a vertex cover of size at most k in G . For each $v \in C$, we add the edge e_v in the vertex gadget H_v to the solution set F . Note that these are $|C| \leq k = k'$ edges. By Lemma 2.1, it suffices to provide a fractional orientation ϕ of $G' - F$ with $\Delta_{\phi}^- \leq \tau$, which we construct as follows. We use the fractional orientation ϕ'_v for each vertex gadget (with potentially the edge e_v removed). Hence, the indegree of each vertex in a vertex gadget is τ unless it is the vertex x_v in a vertex gadget corresponding to a vertex $v \in C$. In this case, the indegree is $\tau - 1$. We next construct the fractional orientation for each edge gadget W_e for an edge $e = \{u, v\}$ (including the two edges connecting it to vertex gadgets). We assume without loss of generality that $u \in C$ (if both are contained, we pick one endpoint arbitrarily; the case where only v is part of the vertex cover is analogous). This means that the vertex x_u has indegree $\tau - 1$ in the current orientation. We use Lemma 5.3 to obtain a fractional orientation of the edges in W_e such that now all vertices in it except for w_u and w_v have indegree τ and w_u has indegree $\tau - 1 + 1/q$ and w_v has indegree $\tau - 1$. We orient the edge $\{w_v, x_v\}$ completely towards w_v and the edge $\{w_u, x_u\}$ with a value of $1 - 1/q$ to w_u and with a value of $1/q$ to x_u . By doing this for every edge e , we obtain a fractional orientation ϕ of $G' - F$ where each vertex has indegree at most τ since each vertex x_u is only connected to at most q edge gadgets that assign an orientation of $1/q$ each.

Conversely, assume that the constructed instance of τ -BDED is a yes-instance and let F be a solution of size at most $k' = k$. We construct a vertex cover C of size at most k for G as follows. For all edges $e \in F$ which are contained in some vertex gadget H_v or connect a vertex in H_v with an edge gadget, we add v to the vertex cover C . If $e \in F$ is contained in an edge gadget $W_{e'}$, then we arbitrarily pick one of the two endpoints of e' and add it to C . Since we selected at most k vertices, it only remains to prove that C is a vertex cover. Assume towards a contradiction that this is not the case. Then, there exists some edge $e = \{u, v\}$ where $u, v \notin C$. Note that this implies that no edge incident to a vertex in the vertex gadgets corresponding to u and v and no edge in the edge gadget corresponding to e are contained in F . Let n_v be the number of vertices in a vertex gadget and let n_e be the number of vertices in an edge gadget. The union of the three aforementioned gadgets (plus the two edges between the gadgets) contains $2n_v + n_e$ vertices and $\tau(2n_v) + (\tau n_e - 2 + 1/q) + 2 = \tau(2n_v + n_e) + 1/q > \tau(2n_v + n_e)$ edges. Hence, the density of this subgraph is larger than τ , contradicting the fact that F is a solution. This concludes the proof. \blacktriangleleft

6 Conclusion

We provided a complexity dichotomy of τ -BDED with respect to the target density τ by showing that the problem is polynomial-time solvable if and only if τ is half-integral or smaller than $2/3$ (assuming $P \neq NP$). We also observed an encouraging result regarding the treewidth of the input graph. Studying τ -BDED with respect to further natural parameters like clique-width, or the diameter of the input graph and/or from the viewpoint of approximation algorithms are natural next steps. Another next step is to analyze cases in which τ is not constant but can depend on e.g. n . Our polynomial-time algorithms for (half-)integral target densities generalize to this setting. However, our hardness results and the FPT algorithm for treewidth do not. It is known that for some value $\tau = 1 - 1/o(n)$ depending on n the problem becomes $W[1]$ -hard with respect to the treewidth [3] and it is polynomial-time solvable for $\tau = 1 - 1/n$ since this is equivalent to computing a feedback edge set of the input graph. Another generalization our polynomial-time algorithms cover are weighted problem variants of τ -BDED as the weighted variants of MAXIMUM s - t -FLOW (FLOW) and GENERAL FACTORS (GF) are polynomial-time solvable [10, 21].

Concerning 2GT, an interesting question is whether the efficiently solvable cases discovered in this work (max-gap at most one, bounded treewidth) can be extended to cover larger classes of networks. For example, the NP-hardness for the case that only some arcs have an even flow restriction requires to “stack” several of them on paths: Each path from s to t in the constructed acyclic flow instance contains exactly three arcs requiring even flow (see full version of the paper). Hence, the question arises if the problem is polynomial-time solvable if each such restricted arc is incident to s or t .

References

- 1 Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. *Network flows – Theory, algorithms and applications*. Prentice Hall, 1993. 4, 5
- 2 Saeed Akhondian Amiri. A note on the fine-grained complexity of MIS on regular graphs. *Information Processing Letters*, 170:106123, 2021. doi:10.1016/J.IPL.2021.106123. 17
- 3 Cristina Bazgan, André Nichterlein, and Sofia Vazquez Alferez. Destroying densest subgraphs is hard. *Journal of Computer and System Sciences*, 151:103635, 2025. doi:10.1016/j.jcss.2025.103635. 2, 3, 13, 19

- 4 David M. Burton. *Elementary Number Theory*. Tata McGraw-Hill Publishing Company Limited, 2006. 15
- 5 Karthekeyan Chandrasekaran, Chandra Chekuri, and Shubhang Kulkarni. On deleting vertices to reduce density in graphs and supermodular functions. In *Proceedings of the 52nd International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 43:1–43:20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPIcs.ICALP.2025.43. 3, 6
- 6 Moses Charikar. Greedy approximation algorithms for finding dense components in a graph. In *Proceedings of the 3rd International Workshop on Approximation Algorithms for Combinatorial Optimization (APPROX)*, pages 84–95. Springer, 2000. doi:10.1007/3-540-44436-X_10. 4
- 7 Li Chen, Rasmus Kyng, Yang P. Liu, Richard Peng, Maximilian Probst Gutenberg, and Sushant Sachdeva. Almost-linear-time algorithms for maximum flow and minimum-cost flow. *Communications of the ACM*, 66(12):85–92, 2023. doi:10.1145/3610940. 2, 4, 5, 6, 11
- 8 Gérard Cornuéjols. General factors of graphs. *Journal of Combinatorial Theory*, 45(2):185–198, 1988. doi:10.1016/0095-8956(88)90068-8. 5
- 9 Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshantov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015. doi:10.1007/978-3-319-21275-3. 12
- 10 Szymon Dudycz and Katarzyna Paluch. Optimal general matchings. In *Proceedings of the 44th International Workshop on Graph-Theoretic Concepts in Computer Science (WG)*, pages 176–189. Springer, 2018. doi:10.1007/978-3-030-00256-5_15. 11, 19
- 11 Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979. 5
- 12 Andrew V. Goldberg. Finding a maximum density subgraph. Technical report, University of California at Berkeley, USA, 1984. 2
- 13 David G. Kirkpatrick and Pavol Hell. On the completeness of a generalized matching problem. In *Proceedings of the 10th Annual ACM Symposium on Theory of Computing (STOC)*, pages 240–245. ACM, 1978. doi:10.1145/800133.804353. 13
- 14 Tuukka Korhonen. A single-exponential time 2-approximation algorithm for treewidth. In *Proceedings of the 62nd IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 184–192, 2021. doi:10.1109/FOCS52979.2021.00026. 2
- 15 Tuukka Korhonen and Daniel Lokshantov. An improved parameterized algorithm for treewidth. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing (STOC)*, pages 528–541, 2023. doi:10.1145/3564246.3585245. 2
- 16 Tommaso Lanciano, Atsushi Miyachi, Adriano Fazzino, and Francesco Bonchi. A survey on the densest subgraph problem and its variants. *ACM Computing Surveys*, 56(8):208:1–208:40, 2024. doi:10.1145/3653298. 1, 2, 3
- 17 Dániel Marx, Govind S. Sankar, and Philipp Schepper. Degrees and gaps: Tight complexity results of general factor problems parameterized by treewidth and cutwidth. In *Proceedings of the 48th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 95:1–95:20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.ICALP.2021.95. 5, 11
- 18 Luke Mathieson and Stefan Szeider. Editing graphs to satisfy degree constraints: A parameterized approach. *Journal of Computer and System Sciences*, 78(1):179–191, 2012. doi:10.1016/J.JCSS.2011.02.001.
- 19 Jean-Claude Picard and Maurice Queyranne. A network flow solution to some nonlinear 0-1 programming problems, with applications to graph theory. *Networks*, 12(2):141–159, 1982. doi:10.1002/NET.3230120206. 2
- 20 Andrzej Ruciński and Andrew Vince. Strongly balanced graphs and random graphs. *Journal of Graph Theory*, 10(2):251–264, 1986. doi:10.1002/jgt.3190100214. 17, 18
- 21 Alexander Schrijver. *Combinatorial optimization: Polyhedra and efficiency*. Springer, 2003. 3, 4, 5, 7, 8, 19