Min-Max Theorems for Packing and Covering Odd (u, v)-trails[*]

Sharat Ibrahimpur[†]   Chaitanya Swamy[†]

Abstract

We investigate the problem of packing and covering odd (u, v)-trails in a graph. A (u, v)-trail is a (u, v)-walk that is allowed to have repeated vertices but no repeated edges. We call a trail *odd* if the number of edges in the trail is odd. Let $\nu(u, v)$ denote the maximum number of edge-disjoint odd (u, v)-trails, and $\tau(u, v)$ denote the minimum size of an edge-set that intersects every odd (u, v)-trail.

We prove that $\tau(u, v) \le 2\nu(u, v) + 1$. Our result is *tight*—there are examples showing that $\tau(u, v) = 2\nu(u, v) + 1$—and substantially improves upon the bound of 8 obtained in [4] for $\tau(u, v)/\nu(u, v)$. Our proof also yields a polynomial-time algorithm for finding a cover and a collection of trails satisfying the above bounds.

Our proof is simple and has two main ingredients. We show that (loosely speaking) the problem can be reduced to the problem of packing and covering odd $(\{u, v\}, \{u, v\})$-trails losing a factor of 2 (either in the number of trails found, or the size of the cover). Complementing this, we show that the odd-$(\{u, v\}, \{u, v\})$-trail packing and covering problems can be tackled by exploiting a powerful min-max result of [2] for packing vertex-disjoint nonzero A-paths in group-labeled graphs.

1   Introduction

Min-max theorems are a classical and central theme in combinatorics and combinatorial optimization, with many such results arising from the study of packing and covering problems. For instance, *Menger's theorem* [10] gives a *tight* min-max relationship for packing and covering edge-disjoint (or internally vertex-disjoint) (u, v)-paths: the maximum number of edge-disjoint (or internally vertex-disjoint) (u, v)-paths (i.e., *packing number*) is equal to the minimum number of edges (or vertices) needed to cover all u-v paths (i.e., *covering number*); the celebrated *max-flow min-cut theorem* generalizes this result to arbitrary edge-capacitated graphs. Another well-known example is the *Lucchesi-Younger theorem* [8], which shows that the maximum number of edge-disjoint directed cuts equals the minimum-size of an arc-set that intersects every directed cut.

Motivated by Menger's theorem, it is natural to ask whether similar (tight or approximate) min-max theorems hold for other variants of path-packing and path-covering problems. Questions of this flavor have attracted a great deal of attention. Perhaps the most prominent results known of this type are Mader's min-max theorems for packing vertex-disjoint $\mathcal{S}$-paths [9, 14], which generalize both the *Tutte-Berge formula* and Menger's theorem, and further far-reaching generalizations of this due to Chudnovsky et al. [2] and Pap [11] regarding packing vertex-disjoint non-zero paths and non-returning A-paths in group-labeled graphs and permutation-labeled graphs respectively.

We consider a different variant of the (u, v)-path packing and covering problems, wherein we impose *parity constraints* on the paths. Such constraints naturally arise in the study of multicommodity-flow problem, which can be phrased in terms of packing odd circuits in a signed graph, and consequently, such

---

odd-circuit packing and covering problems have been widely investigated (see, e.g., [15], Chapter 75). Focusing on (u, v)-paths, a natural variant that arises involves packing and covering *odd* (u, v)-paths, where a (u, v)-path is odd if it contains an odd number of edges. However, there are simple examples [4] showing an unbounded gap between the packing and covering numbers in this setting.

In light of this, following [4], we investigate the min-max relationship for packing and covering odd (u, v)-trails. An *odd* (u, v)-*trail* is a (u, v)-walk with *no repeated edges* and an odd number of edges. Churchley et al. [4] seem to have been the first to consider this problem. They showed that the (worst-case) ratio between the covering and packing numbers for odd (u, v)-trails is at most 8—which is in stark contrast with the setting of odd (u, v) paths, where the ratio is unbounded—and at least 2, so there is no tight min-max theorem like Menger's theorem. They motivate the study of odd (u, v)-trails from the perspective of studying *totally-odd immersions*. In particular, determining if a graph G has k edge-disjoint odd (u, v)-trails is equivalent to deciding if the 2-vertex graph with k parallel edges has a totally-odd immersion into G.

Our results. We prove a *tight bound* on the ratio of the covering and packing numbers for odd (u, v)-trails, which also substantially improves the bound of 8 shown in [4] for this covering-vs-packing ratio.[1] Let $\nu(u, v)$ and $\tau(u, v)$ denote respectively the packing and covering numbers for odd (u, v)-trails. Our main result (Theorem 3.1) establishes that $\tau(u, v) \leq 2\nu(u, v) + 1$. Furthermore, we obtain in polynomial time a certificate establishing that $\tau(u, v) \leq 2\nu(u, v) + 1$. This is because we show that, for any integer $k \geq 0$, we can compute in polynomial time, a collection of k edge-disjoint odd (u, v)-trails, or an odd-(u, v)-trail cover of size at most $2k - 1$. As mentioned earlier, there are examples showing $\tau(u, v) = 2\nu(u, v) + 1$ (see Fig. 1), so our result *settles* the question of obtaining worst-case bounds for the $\tau(u, v)/\nu(u, v)$ ratio.

Notably, our proof is also simple, and noticeably simpler than (and different from) the one in [4]. We remark that the proof in [4] constructs covers of a certain form; in Appendix C, we prove a *lower bound* showing that such covers cannot yield a bound better than 3 on the covering-vs-packing ratio.

Our techniques. We focus on showing that for any k, we can obtain either k edge-disjoint odd (u, v)-trails or a cover of size at most $2k - 1$. This follows from two other auxiliary results which are potentially of independent interest.

Our key insight is that one can *decouple* the requirements of parity and u-v connectivity when constructing odd (u, v)-trails. More precisely, we show that if we have a collection of k edge-disjoint odd $(\{u, v\}, \{u, v\})$-trails, that is, odd trails that start and end at a vertex of $\{u, v\}$, and the u-v edge connectivity, denoted $\lambda(u, v)$, is at least $2k$, then we can obtain k edge-disjoint odd (u, v)-trails (Theorem 3.3). Notice that if $\lambda(u, v) < 2k$, then a min u-v cut yields a cover of the desired size. So the upshot of Theorem 3.3 is that it reduces our task to the *relaxed* problem of finding k edge-disjoint odd $(\{u, v\}, \{u, v\})$-trails. The proof of Theorem 3.3 relies on elementary arguments (see Section 4). We show that given a fixed collection of $2k$ edge-disjoint (u, v)-paths, we can always modify our collection of edge-disjoint trails so as to make progress by decreasing the number of contacts that the paths make with the trails and/or by increasing the number of odd (u, v) trails in the collection. Repeating this process a small number of times thus yields the k edge-disjoint odd (u, v)-trails.

Complementing Theorem 3.3 we prove that we can either obtain k edge-disjoint $(\{u, v\}, \{u, v\})$-trails, or find an odd-$(\{u, v\}, \{u, v\})$-trail cover (which is also an odd-(u, v)-trail cover) of size at most $2k - 2$ (Theorem 3.2). This proof relies on a powerful result of [2] about *packing and covering nonzero A-paths in group-labeled graphs* (see Section 5, which defines these concepts precisely). The idea here is that [2] show that one can obtain either k vertex-disjoint nonzero A-paths or a set of at most $2k - 2$ vertices intersecting all nonzero A-paths, and this can be done in polytime [1, 5] (see also [12]). This is the same type of result that we seek, except that we care about edge-disjoint trails, as opposed to vertex-disjoint paths. However, by moving to a suitable gadget graph (essentially the line graph) where we replace each vertex by a clique,

---

[1]This bound was later improved to 5 [3, 6]. We build upon some of the ideas in [6].

we can encode trails as paths, and edge-disjointness is captured by vertex-disjointness. Applying the result in [2] then yields Theorem 3.2.

Related work. Churchley et al. [4] initiated the study of min-max theorems for packing and covering odd (u, v)-trails. They cite the question of totally-odd immersions as motivation for their work. We say that a graph H has an *immersion* [13] into another graph G, if one can map V_H bijectively to some $U \subseteq V(G)$, and E_H to edge-disjoint trails connecting the corresponding vertices in U. (As noted by [4], trails are more natural objects than paths in the context of reversing an edge-splitting-off operation, as this, in general, creates trails.) An immersion is called *totally odd* if all trails are of odd length. As noted earlier, the question of deciding if a graph G has k edge-disjoint odd (u, v)-trails can be restated as determining if the 2-vertex graph with k parallel edges has a totally-odd immersion into G.

In an interesting contrast to the unbounded gap between the covering and packing numbers for odd (u, v)-paths, [16] showed that the covering number is at most twice the *fractional packing number* (which is the optimal value of the natural odd-(u, v)-path-packing LP).

The notions of odd paths and trails can be generalized and abstracted in two ways. The first involves *signed graphs* [17], and there are various results on packing odd *circuits* in signed graphs, which are closely related to multicommodity flows (see [15], Chapter 75). The second involves *group-labeled graphs*, for which [2, 1] present strong min-max theorems and algorithms for packing and covering vertex-disjoint nonzero A-paths. Pap further generalized the latter results to the setting of packing vertex-disjoint non-returning A-paths in *permutation-labeled graphs*. He obtained both a min-max theorem for the packing problem [11] (which is analogous to the min-max theorem in [2]), and devised an algorithm for computing a maximum-cardinality packing [12].

2 Preliminaries and notation

Let $G = (V, E)$ be an undirected graph. For $X \subseteq V$, we use $E(X)$ to denote the set of edges having both endpoints in X and $\delta(X)$ to denote set of edges with exactly one endpoint in X. For disjoint $X, Y \subseteq V$, we use $E(X, Y)$ to denote the set of edges with one end in X and one end in Y.

A (p, q)-*walk* is a sequence $(x_0, e_1, x_1, e_2, x_2, \ldots, e_r, x_r)$, where $x_0, \ldots, x_r \in V$ with $x_0 = p$, $x_r = q$, and e_i is an edge with ends x_{i-1}, x_i for all $i = 1, \ldots, r$. The vertices $x_1, \ldots, x_{r-1}$ are called the *internal vertices* of this walk. We say that such a (p, q)-walk is a:

- (p, q)-*path*, if either $r > 0$ and all the x_is are distinct (so $p \neq q$), or $r = 0$, which we call a *trivial path*;
- (p, q)-*trail* if all the e_is are distinct (we could have $p = q$).

Thus, a (p, q)-trail is a (p, q)-walk that is allowed to have repeated vertices but *no repeated edges*. Given vertex-sets $A, B \subseteq V$, we say that a trail is an (A, B)-trail to denote that it is a (p, q)-trail for some $p \in A, q \in B$. A (p, q)-trail is called *odd* (respectively, *even*) if it has an odd (respectively, even) number of edges.

Definition 2.1. Let $G = (V, E)$ be a graph, and $u, v \in V$ (we could have $u = v$).
  (a) The *packing number for odd (u, v)-trails*, denoted $\nu(u, v; G)$, is the maximum number of edge-disjoint odd (u, v)-trails in G.
  (b) We call a subset of edges C an *odd (u, v)-trail cover* of G if it intersects every odd (u, v)-trail in G. The *covering number for odd (u, v)-trails*, denoted $\tau(u, v; G)$, is the minimum size of an odd (u, v)-trail cover of G.
We drop the argument G when it is clear from the context.

For any two distinct vertices x, y of G, we denote the size of a minimum (x, y)-cut in G by $\lambda(x, y; G)$, and drop G when it is clear from the context. By the max-flow min-cut (or Menger's) theorem, $\lambda(x, y; G)$

is also the maximum number of edge-disjoint (x, y)-paths in G.

3 Main results and proof overview

Our main result is the following *tight* approximate min-max theorem relating the packing and covering numbers for odd (u, v) trails.

Theorem 3.1. *Let $G = (V, E)$ be an undirected graph, and $u, v \in V$. For any nonnegative integer k, we can obtain in polynomial time, either:*

1. *k edge-disjoint odd (u, v)-trails in G, or*

2. *an odd (u, v)-trail cover of G of size at most $2k - 1$.*

Hence, we have $\tau(u, v; G) \leq 2 \cdot \nu(u, v; G) + 1$.

Theorem 3.1 is tight, as can be seen from Fig. 1; we show in Appendix A that $\nu(u, v; G) = k$ and $\tau(u, v; G) = 2k + 1$ for this instance. (The fact that Theorem 3.1 is tight was communicated to us by [3], who provided a different tight example.)
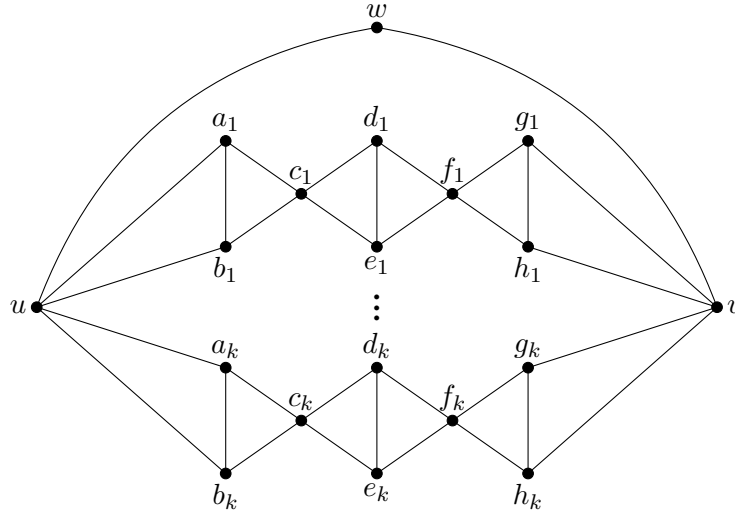

Figure 1: Graph with $\nu(u, v) = k$, $\tau(u, v) = 2k + 1$.

Theorem 3.2. *Let $G = (V, E)$ be an undirected graph and $s \in V$. For any nonnegative integer k, we can obtain in polynomial time:*

1. *k edge-disjoint odd (s, s)-trails in G, or*

2. *an odd (s, s)-trail cover of G of size at most $2k - 2$.*

Theorem 3.1 follows readily from the following two results.

Theorem 3.3. *Let $G = (V, E)$ be an undirected graph, and $u, v \in V$ with $u \neq v$. Let $\widehat{\mathcal{T}}$ be a collection of edge-disjoint odd $(\{u, v\}, \{u, v\})$-trails in G. If $\lambda(u, v) \geq 2 \cdot |\widehat{\mathcal{T}}|$, then we can obtain in polytime $|\widehat{\mathcal{T}}|$ edge-disjoint odd (u, v)-trails in G.*

Proof of Theorem 3.1. If $u = v$, then Theorem 3.2 yields the desired statement. So suppose $u \neq v$. We may assume that $\lambda(u, v) \geq 2k$, since otherwise a minimum (u, v)-cut in G is an odd (u, v)-trail cover of

4

the required size. Let E_{uv} be the uv edge(s) in G (which could be $\emptyset$). Let $\widehat{G}$ be obtained from $G - E_{uv}$ by identifying u and v into a new vertex s. (Note that $\widehat{G}$ has no loops.) Any odd (u,v)-trail in $G - E_{uv}$ maps to an odd (s,s)-trail in $\widehat{G}$. We apply Theorem 3.2 to $\widehat{G}, s, k' = k - |E_{uv}|$. If this returns an odd-(s,s)-trail cover C of size at most $2k' - 2$, then $C \cup E_{uv}$ is an odd-(u,v)-trail cover for G of size at most $2k - 2$. If we obtain a collection of k' edge-disjoint odd (s,s)-trails in $\widehat{G}$, then these together with E_{uv} yield k edge-disjoint odd $(\{u,v\},\{u,v\})$-trails in G. Theorem 3.3 then yields the required k edge-disjoint odd (u,v)-trails. Polytime computability follows from the polytime computability in Theorems 3.2 and 3.3. ∎

Theorem 3.3 is our chief technical insight, which facilitates the decoupling of the parity and u-v connectivity requirements of odd (u,v)-trails, thereby driving the entire proof. (It can be seen as a refinement of Theorem 5.1 in [6].) While Theorem 3.2 returns $(\{u,v\},\{u,v\})$-trails with the right *parity*, Theorem 3.3 supplies the missing ingredient needed to convert these into (u,v)-trails (of the same parity). We give an overview of the proofs of Theorems 3.2 and 3.3 below before delving into the details in the subsequent sections. We remark that both Theorem 3.2 and Theorem 3.3 are *tight* as well, as we show in Appendices D and B respectively.

The proof of Theorem 3.3 relies on elementary arguments and proceeds as follows (see Section 4). Let $\mathcal{P}$ be a collection of $2 \cdot |\widehat{\mathcal{T}}|$ edge-disjoint (u,v)-paths. We provide a simple, efficient procedure to iteratively modify $\widehat{\mathcal{T}}$ (whilst maintaining $|\widehat{\mathcal{T}}|$ edge-disjoint odd $(\{u,v\},\{u,v\})$-trails) and eventually obtain $|\widehat{\mathcal{T}}|$ odd (u,v)-trails. Let $\mathcal{P}_0 \subseteq \mathcal{P}$ be the collection of paths of $\mathcal{P}$ that are edge-disjoint from trails in $\widehat{\mathcal{T}}$. First, we identify the trivial case where $|\mathcal{P}_0|$ is sufficiently large. If so, these paths and $\widehat{\mathcal{T}}$ directly yield odd (u,v)-trails as follows: odd-length paths in $\mathcal{P}_0$ are already odd (u,v)-trails, and even-length paths in $\mathcal{P}_0$ can be combined with odd (u,u)- and odd (v,v)- trails to obtain odd (u,v)-trails.

The paths in $\mathcal{P} \setminus \mathcal{P}_0$, all share at least one edge with some trail in $\widehat{\mathcal{T}}$. Each path is a sequence of edges from u to v. If the first edge that a path $P \in \mathcal{P}$ shares with a trail in $\widehat{\mathcal{T}}$ lies on a (v,v)-trail T, then it is easy to use parts of P and T to obtain an odd (u,v)-trail that is edge-disjoint from all other trails in $\widehat{\mathcal{T}}$, and thereby make progress by increasing the number of odd (u,v)-trails in the collection. A similar conclusion holds if the last edge that a path shares with a trail in $\widehat{\mathcal{T}}$ lies on a (u,u)-trail. If neither of the above cases apply, then the paths in $\mathcal{P} \setminus \mathcal{P}_0$ are in a sense *highly tangled* (which we formalize later) with trails in $\widehat{\mathcal{T}}$. We then infer that $\mathcal{P} \setminus \mathcal{P}_0$ and $\widehat{\mathcal{T}}$ must satisfy some simple structural properties, and leverage this to carefully modify the collection $\widehat{\mathcal{T}}$ (while preserving edge-disjointness) so that the new set of trails are "less tangled" with $\mathcal{P}$ than $\widehat{\mathcal{T}}$, and thereby make progress. Continuing this procedure a polynomial number of times yields the desired collection of $|\widehat{\mathcal{T}}|$ edge-disjoint odd (u,v)-trails.

The proof of Theorem 3.2 relies on the key observation that we can cast our problem as the problem of packing and covering nonzero A-paths in a group-labeled graph (H, Γ) [2] for a suitable choice of A, H, and Γ (see Section 5). In the latter problem, (1) H denotes an oriented graph whose arcs are labeled with elements of a group Γ; (2) an A-path is a path in the undirected version of H whose ends lie in A; and (3) the (Γ-) length of an A-path P is the sum of $\pm\gamma_e$s (suitably defined) for arcs in P, and a nonzero A path is one whose length is not zero (where zero is the identity element of Γ). Chudnovsky et al. [2] show that either there are k *vertex-disjoint* non-zero A-paths, or there is a vertex-set of size at most $2k - 2$ intersecting every non-zero A-path (Theorem 1.1 in [2]). We show that applying their result to a suitable "gadget graph" H (essentially the line graph of G), yields Theorem 3.2 (see Section 5). Polytime computability follows because a subsequent paper [1] gave a polytime algorithm for finding a maximum-size collection of vertex-disjoint non-zero A-paths, and it is implicit in their proof that this also yields a suitable vertex-covering of non-zero A-paths [5].

We remark that while the use of the packing-covering result in [2] yields quite a compact proof of Theorem 3.2, it also makes the resulting proof somewhat opaque since we apply the result in [2] to the gadget graph. However, it is possible to translate the min-max theorem for packing vertex-disjoint nonzero

A-paths proved in [2] to our setting and obtain the following more-accessible min-max theorem for packing edge-disjoint odd (s, s)-trails (stated in terms of G and not the gadget graph). In Section 6, we prove that

$$\nu(s, s; G) = \min\left(|E(S) - F| + \sum_{C \in \text{comp}(G-S)} \left\lfloor \frac{|E(S, C)|}{2} \right\rfloor\right) \tag{1}$$

where the minimum is taken over all bipartite subgraphs (S, F) of G such that $s \in S$.

Notice that Theorem 3.2 follows easily from this min-max formula: if (S^*, F^*) is a bipartite subgraph of G with $s \in S^*$ that attains the minimum above, then the edges in $E(S^*) - F$ combined with $\max\{0, |E(S^*, C)| - 1\}$ edges from $E(S^*, C)$ for every component C of $G - S^*$ yields a cover of size at most twice the RHS of (1).

4 Proof of Theorem 3.3: converting edge-disjoint odd $(\{u, v\}, \{u, v\})$-trails to edge-disjoint odd (u, v)-trails

Recall that $\widehat{\mathcal{T}}$ is a collection of edge-disjoint odd $(\{u, v\}, \{u, v\})$-trails in G. We denote the subset of odd (u, u)-trails, odd (v, v)-trails, and odd (u, v)-trails in $\widehat{\mathcal{T}}$ by $\widehat{\mathcal{T}}_{uu}$, $\widehat{\mathcal{T}}_{vv}$, and $\widehat{\mathcal{T}}_{uv}$, respectively. Let $k_{uu}(\widehat{\mathcal{T}}) = |\widehat{\mathcal{T}}_{uu}|$, $k_{vv}(\widehat{\mathcal{T}}) = |\widehat{\mathcal{T}}_{vv}|$, and $k_{uv}(\widehat{\mathcal{T}}) = |\widehat{\mathcal{T}}_{uv}|$. To keep notation simple, we will drop the argument $\widehat{\mathcal{T}}$ when its clear from the context. Since we are given that $\lambda(u, v) \geq 2 \cdot |\widehat{\mathcal{T}}|$, we can obtain a collection $\mathcal{P}$ of $2 \cdot |\widehat{\mathcal{T}}|$ edge-disjoint (u, v)-paths in G. In the sequel, while we will modify our collection of odd $(\{u, v\}, \{u, v\})$-trails, $\mathcal{P}$ stays fixed.

We now introduce the key notion of a *contact* between a trail T and a (u, v)-path P. Suppose that $P = (x_0, e_1, x_1, \ldots, e_r, x_r)$ for some $r \geq 1$.

Definition 4.1. A *contact* between P and T is a *maximal subpath S of P containing at least one edge* such that S is also a subtrail of T i.e., for $0 \leq i < j \leq r$, we say that $(x_i, e_{i+1}, x_{i+1}, \ldots, e_j, x_j)$ is a contact between P and T if $(x_i, e_{i+1}, x_{i+1}, \ldots, e_j, x_j)$ is a subtrail of T, but neither $(x_{i-1}, e_i, x_i, \ldots e_j, x_j)$ (if $i > 0$) nor $(x_i, e_{i+1}, x_{i+1}, \ldots, x_j, e_{j+1}, x_{j+1})$ (if $j < r$) is a subtrail of T.

Define $\quad \mathcal{C}(P, T) = \left| \left\{ (i, j) : 0 \leq i < j \leq r, \quad (x_i, e_{i+1}, x_{i+1}, \ldots, e_j, x_j) \text{is a contact between } P \text{ and } T \right\} \right|$

By definition, contacts between P and T are edge disjoint. For an edge-disjoint collection $\mathcal{T}$ of trails, we use $\mathcal{C}(P, \mathcal{T})$ to denote $\sum_{T \in \mathcal{T}} \mathcal{C}(P, T)$. So if $\mathcal{C}(P, \mathcal{T}) = 0$, then P is edge-disjoint from every trail in $\mathcal{T}$. Otherwise, we use the term *first contact* of P to refer to the contact arising from the first edge that P shares with some trail in $\mathcal{T}$ (note that P is a (u, v)-walk so is a sequence from u to v). Similarly, the *last contact* of P is the contact arising from the last edge that P shares with some trail in $\mathcal{T}$. If $\mathcal{C}(P, \mathcal{T}) = 1$, then the first and last contacts of P are the same. We further overload notation and use $\mathcal{C}(\mathcal{P}, \mathcal{T})$ to denote $\sum_{P \in \mathcal{P}} \mathcal{C}(P, \mathcal{T}) = \sum_{P \in \mathcal{P}, T \in \mathcal{T}} \mathcal{C}(P, T)$. We use $\mathcal{C}(\mathcal{P}, \mathcal{T})$ as a measure of how "tangled" $\mathcal{T}$ is with $\mathcal{P}$. The following lemma classifies five different cases that arise for any pair of edge-disjoint collections of odd $(\{u, v\}, \{u, v\})$-trails and (u, v)-paths.

Lemma 4.2. *Let $\mathcal{T}$ be a collection of edge-disjoint odd $(\{u, v\}, \{u, v\})$-trails in G. If $|\mathcal{P}| \geq 2 \cdot |\mathcal{T}|$, then one of the following conditions holds.*

(a) *There are at least $k_{uu}(\mathcal{T}) + k_{vv}(\mathcal{T})$ paths in $\mathcal{P}$ that make no contact with any trail in $\mathcal{T}$.*

(b) *There exists a path $P \in \mathcal{P}$ that makes its first contact with a trail $T \in \mathcal{T}_{vv}$.*

(c) *There exists a path $P \in \mathcal{P}$ that makes its last contact with a trail $T \in \mathcal{T}_{uu}$.*

(d) *There exist three distinct paths $P_1, P_2, P_3 \in \mathcal{P}$ that make their first contact with a trail $T \in \mathcal{T}_{uu} \cup \mathcal{T}_{uv}$.*

6

(e) There exist three distinct paths $P_1, P_2, P_3 \in \mathcal{P}$ that make their last contact with a trail $T \in \mathcal{T}_{uv} \cup \mathcal{T}_{vv}$.

Proof. To keep notation simple, we drop the argument $\mathcal{T}$ in the proof. Suppose that conclusion (a) does not hold. Then there are at at least $2 \cdot |\mathcal{T}| - (k_{uu} + k_{vv} - 1) = 2k_{uv} + k_{uu} + k_{vv} + 1$ paths in $\mathcal{P}$ that make at least one contact with some trail in $\mathcal{T}$. Let $\mathcal{P}' \subseteq \mathcal{P}$ be this collection of paths. If either conclusions (b) or (c) hold (for some $P \in \mathcal{P}'$), then we are done, so assume that this is not the case. Then, every path $P \in \mathcal{P}'$ makes its first contact with a trail in $\mathcal{T}_{uu} \cup \mathcal{T}_{uv}$ and its last contact with a trail in $\mathcal{T}_{uv} \cup \mathcal{T}_{vv}$. Note that the number of first and last contacts are both at least $2k_{uv} + k_{uu} + k_{vv} + 1 > 2 \cdot \min(k_{uv} + k_{uu}, k_{uv} + k_{vv})$. So if $k_{uu} \le k_{vv}$, then by the Pigeonhole principle, there are at least 3 paths that make their first contact with some $T \in \mathcal{T}_{uu} \cup \mathcal{T}_{uv}$, i.e., conclusion (d) holds. Similarly, if $k_{vv} \le k_{uu}$, then conclusion (e) holds. ∎

We now leverage the above classification and show that in each of the above five cases, we can make progress by "untangling" the trails (i.e., decreasing $C(\mathcal{P}, \mathcal{T})$) and/or increasing the number of odd (u, v)-trails in our collection.

Lemma 4.3. *Let $\mathcal{T}$ be a collection of edge-disjoint odd $(\{u, v\}, \{u, v\})$-trails. If $|\mathcal{P}| \ge 2 \cdot |\mathcal{T}|$, we can obtain another collection $\mathcal{T}'$ of edge-disjoint odd $(\{u, v\}, \{u, v\})$-trails such that at least one of the following holds.*

(i) $k_{uv}(\mathcal{T}') = |\mathcal{T}|$.

(ii) $C(\mathcal{P}, \mathcal{T}') \le C(\mathcal{P}, \mathcal{T})$ and $k_{uv}(\mathcal{T}') = k_{uv}(\mathcal{T}) + 1$.

(iii) $C(\mathcal{P}, \mathcal{T}') \le C(\mathcal{P}, \mathcal{T}) - 1$ and $k_{uv}(\mathcal{T}') \ge k_{uv}(\mathcal{T}) - 1$.

Proof. If $k_{uv}(\mathcal{T}) = |\mathcal{T}|$, then (i) holds trivially by taking $\mathcal{T}' = \mathcal{T}$. So we may assume that $\mathcal{T}$ contains some odd (u, u)- or odd (v, v)-trail. Observe that $\mathcal{T}$ and $\mathcal{P}$ satisfy the conditions of Lemma 4.2, so at least one of the five conclusions of Lemma 4.2 applies. We handle each case separately.

(a) At least $k_{uu}(\mathcal{T}) + k_{vv}(\mathcal{T})$ paths in $\mathcal{P}$ have zero contacts with $\mathcal{T}$. Let $\mathcal{P}_0 = \{P \in \mathcal{P} : C(P, \mathcal{T}) = 0\}$. Consider some $P \in \mathcal{P}_0$. If P is odd, we can replace an odd (u, u)- or odd (v, v)- trail in $\mathcal{T}$ with P. If P is even, then P can be combined with an odd (u, u)- or odd (v, v)- trail to obtain an odd (u, v)-trail. Since $|\mathcal{P}_0| \ge k_{uu}(\mathcal{T}) + k_{vv}(\mathcal{T})$, we create $k_{uu}(\mathcal{T}) + k_{uv}(\mathcal{T})$ odd (u, v)-trails this way, and this new collection $\mathcal{T}'$ satisfies (i).

(b) Some $P \in \mathcal{P}$ makes its first contact with an odd (v, v)-trail $T \in \mathcal{T}$. Let the first vertex in the first contact between P and T be x. Observe that x partitions the trail T into two subtrails S_1 and S_2. Since T is an odd trail, exactly one of S_1 and S_2 is odd. We can now obtain an odd (u, v)-trail T' by traversing P from u to x, and then traversing S_1 or S_2, whichever yields odd parity (see Fig. 2). Since P already made a contact with T, we have $C(P, T') \le C(P, T)$, and $C(Q, T') \le C(Q, T)$ for any other path $Q \in \mathcal{P}$. Thus, taking $\mathcal{T}' = (\mathcal{T} \cup \{T'\}) \setminus \{T\}$, we have $C(\mathcal{P}, \mathcal{T}') \le C(\mathcal{P}, \mathcal{T})$, and (ii) holds.

$$T_1 = u \xrightarrow[P]{} x \xrightarrow[S_1]{} v$$

$$T_2 = u \xrightarrow[P]{} x \xrightarrow[S_2]{} v$$

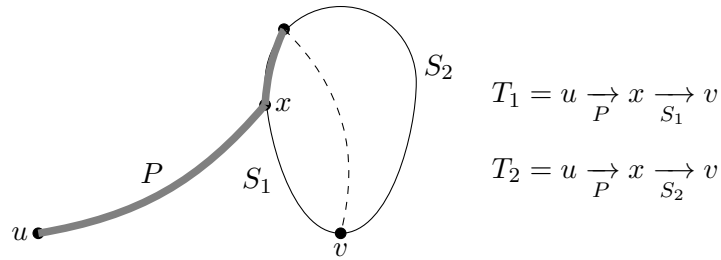Figure 2: Path P makes its first contact with an odd (v, v)-trail.

(c) Some $P \in \mathcal{P}$ makes its last contact with an odd (u, u)-trail $T \in \mathcal{T}$. This is completely symmetric to (b), so a similar strategy works and we satisfy (ii).

(d) Paths $P_1, P_2, P_3 \in \mathcal{P}$ that make their first contact with an odd $(u, \{u, v\})$-trail $T \in \mathcal{T}$. Note that all contacts between paths in $\mathcal{P}$ and trails in $\mathcal{T}$ are edge disjoint, since the paths in $\mathcal{P}$ are edge disjoint and the trails in $\mathcal{T}$ are edge disjoint. For $i = 1, 2, 3$, let the first vertex in the first contact of P_i (with T) be x_i. Let Q_i denote the subpath of P_i between u and x_i. Note that T is a sequence of edges from u to some vertex in $\{u, v\}$. Without loss of generality, assume that in T, the first contact of P_1 appears before the first contact of P_2, which appears before the first contact of P_3. The vertices x_1, x_2, x_3 partition the trail T into four subtrails S_0, S_1, S_2, S_3 (see Fig. 3). For a trail X, we denote the reverse sequence of X by $\overline{X}$. Now consider the following trails (where $+$ denotes concatenation):

$$T_1 = S_0 + \overline{Q}_1, \quad T_2 = Q_1 + S_1 + \overline{Q}_2, \quad T_3 = Q_2 + S_2 + \overline{Q}_3, \quad T_4 = Q_3 + S_3.$$

Observe that the disjoint union of edges in T_1, T_2, T_3, and T_4 has the same parity as that of T, and hence at least one of the T_is is an odd trail; call this trail T'. Let $\mathcal{T}' = \mathcal{T} \cup \{T'\} \setminus \{T\}$. By construction, *every* T_i avoids at least one of the (first) contacts made by P_1, P_2, or P_3 (with T). Also, for any other path $Q \in \mathcal{P} \setminus \{P_1, P_2, P_3\}$, we have $\mathcal{C}(Q, T') \leq \mathcal{C}(Q, T)$. Therefore, $\mathcal{C}(\mathcal{P}, \mathcal{T}') \leq \mathcal{C}(\mathcal{P}, \mathcal{T}) - 1$. It could be that T was an odd (u, v)-trail, which is now replaced by an odd (u, u)-trail, so $k_{uv}(\mathcal{T}') \geq k_{uv}(\mathcal{T}) - 1$. So we satisfy (iii).

$$T_1 = u \xrightarrow[S_0]{} x_1 \xrightarrow[\overline{P}_1]{} u$$

$$T_2 = u \xrightarrow[P_1]{} x_1 \xrightarrow[S_1]{} x_2 \xrightarrow[\overline{P}_2]{} u$$

$$T_3 = u \xrightarrow[P_2]{} x_2 \xrightarrow[S_2]{} x_3 \xrightarrow[\overline{P}_3]{} u$$
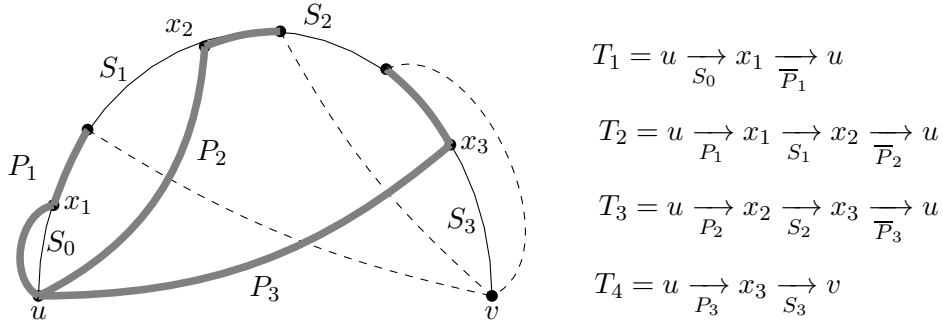
$$T_4 = u \xrightarrow[P_3]{} x_3 \xrightarrow[S_3]{} v$$

Figure 3: Paths P_1, P_2, P_3 make their first contact with an odd (u, v)-trail.

(e) Paths $P_1, P_2, P_3 \in \mathcal{P}$ make their last contact with an odd $(\{u, v\}, v)$-trail in $\mathcal{T}$. This is completely symmetric to (d) and the same approach works, so we again satisfy (iii). ∎

Theorem 3.3 now follows by simply applying Lemma 4.3 starting with the initial collection $\mathcal{T}^0 := \widehat{\mathcal{T}}$ until conclusion (i) of Lemma 4.3 applies. The $\mathcal{T}'$ returned by this final application of Lemma 4.3 then satisfies the theorem statement.

We now argue that this process terminates in at most $2 \cdot |E(G)| + |\widehat{\mathcal{T}}|$ steps, which will conclude the proof. Let $k = |\widehat{\mathcal{T}}|$. Consider the following potential function defined on a collection $\mathcal{T}$ of k edge-disjoint odd $(\{u, v\}, \{u, v\})$-trails: $\phi(\mathcal{T}) := 2 \cdot \mathcal{C}(\mathcal{P}, \mathcal{T}) - k_{uv}(\mathcal{T})$. Consider any iteration where we invoke Lemma 4.3 and move from a collection $\mathcal{T}$ to another collection $\mathcal{T}'$ with $k_{uv}(\mathcal{T}') < k$. Then, either conclusion (ii) or (iii) of Lemma 4.3 applies, and it is easy to see that $\Phi(\mathcal{T}') \leq \phi(\mathcal{T}) - 1$. Finally, we have $-k \leq \Phi(\mathcal{T}) \leq 2 \cdot |E(G)|$ for all $\mathcal{T}$ since $0 \leq \mathcal{C}(\mathcal{P}, \mathcal{T}) \leq |E(G)|$ as the contacts between paths in $\mathcal{P}$ and trails in $\mathcal{T}$ are edge-disjoint, so the process terminates in at most $2|E(G)| + k$ steps.

5 Proof of Theorem 3.2

Our proof relies on two reductions both involving non-zero A-paths in a group-labeled graph, which we now formally define. A *group-labeled graph* is a pair (H, Γ), where Γ is a group, and $H = (N, E')$ is an oriented graph (i.e., for any $u, v \in N$, if $(u, v) \in E'$ then $(v, u) \notin E'$) whose arcs are labeled with elements of Γ. All addition (and subtraction) operations below are always with respect to the group Γ. A path P in H is a sequence $(x_0, e_1, x_1, \ldots, e_r, x_r)$, where the x_is are distinct, and each e_i has ends x_i, x_{i+1} but could be oriented either way (i.e., as (x_i, x_{i+1}) or (x_{i+1}, x_i)). (So upon removing arc directions, P yields a path in the undirected version of H.) We say that P traverses e_i in the direction (x_i, x_{i+1}). The Γ-*length* (or simply length) of P, denoted $\gamma(P)$, is the sum of $\pm \gamma_e$s for arcs in P, where we count $+\gamma_e$ for e if P's traversal of e matches e's orientation and $-\gamma_e$ otherwise. Given $A \subseteq N$, an A-path is a path $(x_0, e_1, \ldots, e_r, x_r)$ where $r \geq 1$, and $x_0, x_r \in A$; finally, call an A-path P a nonzero A-path if $\gamma(P) \neq 0$ (where 0 denotes the identity element for Γ).

Chudnovsky et al. [2] proved the following theorem as a consequence of a min-max formula they obtain for the maximum number of vertex-disjoint nonzero A-paths. Subsequently, [1] devised a polytime algorithm to compute the maximum number of vertex-disjoint non-zero A-paths. Their algorithm also implicitly computes the quantities needed in (the minimization portion of) their min-max formula to show the optimality of the collection of A-paths they return [5]; this in turn easily yields the vertex-set mentioned in Theorem 5.1.

Theorem 5.1 ([2, 1]). *Let $\big(H = (N, E'), \Gamma\big)$ be a group-labeled graph, and $A \subseteq V$. Then, for any integer k, one can obtain in polynomial time, either:*

1. *k vertex-disjoint nonzero A-paths, or*

2. *a set of at most $2k - 2$ vertices that intersects every nonzero A-path.*

Recall that G is the undirected graph in the theorem statement, and $s \in V$. For a suitable choice of a group-labeled graph (H, Γ), and a vertex-set A, we show that: (a) vertex-disjoint nonzero A-paths in (H, Γ) yield edge-disjoint odd (s, s)-trails; and (b) a *vertex-set* covering all nonzero A-paths in (H, Γ) yields an odd (s, s)-trail cover of G. Combining this with Theorem 5.1 finishes the proof.

Since we are dealing with parity, it is natural to choose $\Gamma = \mathbb{Z}_2$ (so the orientation of edges in H will not matter). To translate vertex-disjointness (and vertex-cover) to edge-disjointness (and edge-cover), we essentially work with the line graph of G, but slightly modify it to incorporate edge labels. We replace each vertex $x \in V$ with a clique of size $\deg_G(x)$, with each clique-node corresponding to a distinct edge of G incident to x; we use $[x]$ to denote this clique, both its set of nodes and edges; the meaning will be clear from the context. (Note that if $\deg_G(x) = 0$, then there are no nodes and edges corresponding to x in H; this is fine since isolated nodes in G can be deleted without affecting anything.) For every edge $e = xy \in E$, we create an edge between the clique nodes of $[x]$ and $[y]$ corresponding to e. We arbitrarily orient the edges to obtain H. We give each clique edge a label of 0, and give every other edge a label of 1. Finally, we let $A = [s]$.

Lemma 5.2. *The following properties hold.*

(a) *Every A-path P in H maps to an (s, s)-trail $T = \pi(P)$ in G such that $\gamma(P) = 1$ iff T is an odd trail.*

(b) *If two A-paths P, Q are vertex disjoint then the (s, s)-trails $\pi(P)$ and $\pi(Q)$ are edge disjoint.*

(c) *Every (s, s)-trail T in G with at least one edge maps to an A-path $P = \sigma(T)$ in G such that: T is an odd trail iff $\gamma(P) = 1$, and P contains a vertex x iff T contains the corresponding edge of G.*

Proof. The proof is straightforward. By definition, the 1-labeled edges of H are in bijective correspondence with the edges of G.

For part (a), let $P = (x_0, e_1, \ldots, e_r, x_r)$ be an A-path in H. Let $P' = (e_{i_1}, \ldots, e_{i_q})$ be the subsequence of P consisting of the 1-labeled edges of P. If $P' = ()$, then define $\pi(P)$ to be the trivial (s, s)-trail (s), which satisfies the required property. Let f_j be the edge of G corresponding to e_{i_j}. Then, for $j = 2, \ldots, q - 1$, we have $f_j = u_{j-1} u_j$, where u_{j-1} is such that $[u_{j-1}]$ contains one end of both e_{i_j} and $e_{i_{j-1}}$, and u_j is such that $[u_j]$ contains one end of both e_{i_j} and $e_{i_{j+1}}$. The end of e_{i_1} not in $[u_1]$ must lie in $[s]$ (as all edges of P occurring before e_{i_1} have label 0, and so must be edges of $[s]$). Similarly, the end of e_{i_q} that does not lie in $[u_{q-1}]$ must lie in $[s]$ (as all edges of P occurring after e_{i_q} must be edges of $[s]$). Therefore, the sequence $f_1, \ldots, f_q$ yields an (s, s)-trail $\pi(P)$ in G, and we have $|\pi(P)| = |P'| = \gamma(P)$. Observe that an edge f of G lies in $\pi(P)$ iff the corresponding 1-labeled edge of H lies in P.

For part (b), let $T = \pi(P)$, $T' = \pi(Q)$ be the corresponding (s, s)-trails. We may assume that T and T' contain at least one edge, otherwise the statement is vacuously true. Since P and Q are vertex disjoint, their subsequences of 1-labeled edges, which map to the edges of T and T', do not contain any common edges, so T and T' are edge disjoint.

For part (c), consider an (s, s)-trail $(u_0 = s, f_1, \ldots, f_r, u_r = s)$, where $r \geq 1$. Each f_i maps to a distinct 1-labeled edge e_i of H. The edges e_i and e_{i+1} both have one end incident to a distinct node in $[u_i]$, and there is a clique edge joining these ends. Since the f_is are all distinct and no two 1-labeled edges of H share an endpoint, the ends of all the e_is are distinct. Let x_0 be the end of e_1 in $[u_0] = [s]$, and x_r be the end of e_r in $[u_r] = [s]$. Note that $x_0 \neq x_r$ since $e_1 \neq e_r$. So the edges e_i interspersed with the clique edges from $[u_i]$ yield an (x_0, x_r)-path $\sigma(T)$ in H, which is therefore an A-path. Also, $\gamma(P) = r \bmod 2$, so P is a nonzero A-path iff T is an odd trail. This proves part (ii). By construction, we ensure that a node x lies in P iff the corresponding edge f of G lies in T. ∎

To complete the proof of Theorem 3.2, we apply Theorem 5.1 to the nonzero A-path instance $(H, [s], \gamma, \mathbb{Z}_2)$ constructed above. If we obtain k vertex-disjoint nonzero A-paths in H, then parts (a) and (b) of Lemma 5.2 imply that we can map these to k edge-disjoint odd (s, s)-trails. Alternatively, if we obtain a set C of at most $2k - 2$ vertices of H that intersect every nonzero A-path, then we obtain a cover F for odd (s, s)-trails in G by taking the set of edges in G corresponding to the vertices in C. To see why F is a cover, suppose that the graph $G - F$ has an odd (s, s)-trail. This then maps to a nonzero A-path P in H such that $P \cap C = \emptyset$ by part (c) of Lemma 5.2, which yields a contradiction.

6 Min-max theorem for packing odd edge-disjoint (s, s)-trails

Let G be an undirected graph. For a node $s \in V(G)$, recall that $\nu(s, s; G)$ denotes the maximum number of edge-disjoint odd (s, s)-trails in G. Let $\mathrm{comp}(G)$ denote the set of components of G. Given disjoint vertex sets S, T, we use $E_G(S, T)$ to denote the set of edges in G with one end in S and one end in T. Let $E_G(S)$ be the set of edges with both ends in S. We drop the subscript G above when it is clear from the context. We prove the following min-max theorem.

Theorem 6.1. *Let G be an undirected graph, and s be a node in $V(G)$. Then,*

$$\nu(s, s; G) = \min \left\{ |E(S) - F| + \sum_{C \in \mathrm{comp}(G-S)} \left\lfloor \frac{|E(S, C)|}{2} \right\rfloor \; : \; s \in S, \, (S, F) \text{ is a bipartite subgraph of } G \right\}.$$

(2)

In the sequel, we fix G and s to be the graph and the node mentioned in the statement of Theorem 6.1. It is easy to see that for any bipartite subgraph (S, F) with $s \in S$, the expression in the RHS of (2) gives an upper bound on $\nu(s, s; G)$: an odd (s, s)-trail either lies completely within S and must therefore use some edge of $E(S) \setminus F$, or must include a vertex of some component $K \in \mathrm{comp}(G - S)$; there are at most

10

$E(S) \setminus F$ edge-disjoint trails of the first type, and at most $\left\lfloor \frac{|E(S,K)|}{2} \right\rfloor$ edge-disjoint trails that may use a vertex of a component K. Our goal is to show that for some choice of (S, F), the expression in the RHS of (2) equals $\nu(s, s; G)$.

Let $(H, [s], \gamma : E(H) \mapsto \{0, 1\}, \mathbb{Z}_2)$ be the nonzero A-path packing instance obtained from (G, s) as described in Section 5. Let $\nu(H, [s], \gamma)$ denote the maximum number of vertex-disjoint non-zero $[s]$-paths in H under the labeling γ, so $\nu(s, s; G) = \nu(H, [s], \gamma)$. As mentioned earlier, we obtain our min-max theorem by applying the min-max formula of Chudnovsky et al. [2] for packing vertex-disjoint nonzero A-paths in group-labeled graphs to $(H, [s], \mathbb{Z}_2)$, and simplifying the resluting expression by leveraging the structure underlying the gadget graph H. Our starting point therefore is the following result obtained by specializing Theorem 1.2 in [2] to the A-path instance $(H, [s], \gamma, \mathbb{Z}_2)$. We need the following notation from [2]. Given $y \in V(H)$, *switching* the labeling γ at y means that we flip the labels of the edges incident to y; that is, we obtain a new labeling γ' where $\gamma'_e = 1 - \gamma_e$ if e is incident to y and $\gamma'_e = \gamma_e$ otherwise. Note that if $y \notin [s]$, then $\nu(H, [s], \gamma) = \nu(H, [s], \gamma')$. For $A' \subseteq V(H)$ and a labeling γ', let $E(A', \gamma')$ denote the set of all edges $e \in E(H)$ with both ends in A' and having $\gamma'_e = 0$.

Theorem 6.2 (Corollary of Theorem 1.2 in [2]). *There exists a labeling* $\gamma' : E(H) \mapsto \{0, 1\}$ *obtained by switching* γ *at some (suitably chosen) vertices in* $V(H) - [s]$, *and vertex-sets* $X, A' \subseteq V(H)$ *with* $[s] - X \subseteq A' \subseteq V(H) - X$ *such that*

$$\nu(H, [s], \gamma) = |X| + \sum_{K \in \text{comp}(H - X - E(A', \gamma'))} \left\lfloor \frac{|A' \cap V(K)|}{2} \right\rfloor. \tag{3}$$

We may assume that the labeling γ' given above is obtained by switching γ at some subset of $A' - [s]$, since we can always switch γ' at vertices not in A' without affecting the RHS of (3). It will be convenient to restate Theorem 6.2 by explicitly referring to the subset of A' at which γ was switched, as follows. Let $Y, B_0, B_1 \subseteq V(H)$ be disjoint sets such that $[s] - Y \subseteq B_0$ and $B_0 \cup B_1 \subseteq V(H) - Y$; we call (Y, B_0, B_1) a valid triple. Let $\gamma(B_1)$ denote the labeling obtained from γ by switching at every vertex in B_1. Let $H(Y, B_0, B_1) := H - Y - E(B_0 \cup B_1, \gamma(B_1))$. Define

$$p(Y, B_0, B_1) := |Y| + \sum_{K \in \text{comp}(H(Y, B_0, B_1))} \left\lfloor \frac{|(B_0 \cup B_1) \cap V(K)|}{2} \right\rfloor.$$

As noted in [2], it is not hard to see that $\nu(H, [s], \gamma) \le p(Y, B_0, B_1)$ for any valid triple (Y, B_0, B_1). We call a valid triple (Y, B_0, B_1) *tight* if $\nu(H, [s], \gamma) = p(Y, B_0, B_1)$. Let the labeling γ' given by Theorem 6.2 be obtained by switching γ at $A'_1 \subseteq A'$. Let $A'_0 := A' - A'_1$. Theorem 6.2 then tells us that (X, A'_0, A'_1) is a tight triple. Exploiting the structure of H, we show that there exists a tight triple of a convenient form, which will then lead to our min-max formula.

Lemma 6.3. *There exists a tight triple* $(\emptyset, B_0, B_1)$ *such that, for every* $x \in V(G)$, *we have*

$$[x] \subseteq B_0 \quad or \quad [x] \subseteq B_1 \quad or \quad [x] \cap (B_0 \cup B_1) = \emptyset. \tag{4}$$

Before proving Lemma 6.3, we show how this yields Theorem 6.1.

Proof of Theorem 6.1. Let $(\emptyset, B_0, B_1)$ be the tight triple given by Lemma 6.3. Let $\widetilde{\gamma} = \gamma(B_1)$. Take $S = S_0 \cup S_1$, where $S_0 := \{x \in V(G) : [x] \subseteq B_0\}$ and $S_1 := \{x \in V(G) : [x] \subseteq B_1\}$. Note that $s \in S_0$. Let $F := E_G(S_0, S_1)$. Let $C_1, \ldots, C_\ell$ be the components of $G - S$. We show that $p(\emptyset, B_0, B_1) = |E(S) - F| + \sum_{i=1}^{\ell} \left\lfloor \frac{|E(S, C_i)|}{2} \right\rfloor$, which together with the fact that $(\emptyset, B_0, B_1)$ is a tight triple, completes the proof.

11

Observe that for every vertex $x \in S$, all the clique edges of H in $[x]$ have $\widetilde{\gamma}$-label 0 (since γ is switched at either both endpoints or no endpoint) and thus belong to $E(B_0 \cup B_1, \gamma(B_1))$. Therefore such edges do not appear in $\widetilde{H} := H(\emptyset, B_0, B_1)$. Also, for every edge $e = xy \in F$, there is a corresponding edge e in H between some vertex w in B_0 and some vertex z in B_1, and $\widetilde{\gamma}_e = 0$; so $e \notin E_{\widetilde{H}}$. Next, for every $e = xy \in E(S) - F = E(S_0) \cup E(S_1)$, there is a corresponding edge in $\widetilde{H}$ between some vertex $w \in [x]$ and some vertex $z \in [y]$, and this edge has $\widetilde{\gamma}$-label 1. Thus, $\{w, z\}$ is a connected component of $\widetilde{H}$ of size 2, and contributes 1 to $p(\emptyset, B_0, B_1)$.

We now argue that components $\{C_i\}_{i=1}^{\ell}$ are in bijection with the remaining components of $\widetilde{H}$, and if C_i corresponds to component K_i of $\widetilde{H}$, then K_i contributes $\left\lfloor \frac{|E(S,C_i)|}{2} \right\rfloor$ to $p(\emptyset, B_0, B_1)$. Consider any component C_i. Let $E(S, C_i) = \{x_1 y_1, \ldots, x_r y_r\}$, where $x_j \in S$ and $y_j \in C_i$ for all $j = 1, \ldots, r$. For $j = 1, \ldots, r$, let $w_j \in [x_j]$ be the end in $B_0 \cup B_1$ of the unique edge in $\widetilde{H}$ corresponding to edge $x_j y_j$. For every y_j, all vertices of H in $[y_j]$ belong to the same component of $\widetilde{H}$. So since $y_1, \ldots, y_r$ are connected in $G - S$, all vertices in $\bigcup_{j=1}^{r}[y_j]$ belong to the same component of $\widetilde{H}$, say K_i. It follows that $K_i \supseteq \{w_1, \ldots, w_r\} \cup \left(\bigcup_{j=1}^{r}[y_j]\right)$. Thus we have,

$$\left\lfloor \frac{|E(S, C_i)|}{2} \right\rfloor = \left\lfloor \frac{r}{2} \right\rfloor = \left\lfloor \frac{|V(K_i) \cap (B_0 \cup B_1)|}{2} \right\rfloor .$$

Moreover, since $G - S = \bigcup_{i=1}^{\ell} C_i$, it follows that $\widetilde{H}$ has no components other than $K_1, \ldots, K_\ell$. So to summarize, we have shown that $p(\emptyset, B_0, B_1) = |E(S) - F| + \sum_{i=1}^{\ell} \left\lfloor \frac{|E(S,C_i)|}{2} \right\rfloor$, completing the proof. ∎

We now prove Lemma 6.3. We will need the following claim.

Claim 6.4. *For any integer $q \geq 2$ and nonnegative integers $r_1, \ldots, r_q$ we have*

$$\sum_{i=1}^{q} \left\lfloor \frac{r_i}{2} \right\rfloor \geq \left\lfloor \frac{\sum_{i=1}^{q} r_i - (q-1)}{2} \right\rfloor . \tag{5}$$

Proof. For any integer x, we have $\lfloor \frac{x}{2} \rfloor \geq \frac{x-1}{2}$, and equal to $\frac{x}{2}$ if x is even. So if any of the r_i's are even, the statement follows. Otherwise, $\sum_{i=1}^{q} r_i - q$ is even, and both the LHS and RHS of (5) are equal to $\frac{\sum_{i=1}^{q} r_i - q}{2}$. ∎

Proof of Lemma 6.3. Let (X, A_0', A_1') be the tight triple given by Theorem 6.2. First, we show how to modify (X, A_0', A_1') to obtain a tight triple $(X, \widetilde{A}_0, \widetilde{A}_1)$ that satisfies

$$([y] - X) \subseteq \widetilde{A}_0 \quad \text{or} \quad ([y] - X) \subseteq \widetilde{A}_1 \quad \text{or} \quad ([y] - X) \subseteq V(H(X, \widetilde{A}_0, \widetilde{A}_1)) - (\widetilde{A}_0 \cup \widetilde{A}_1) \tag{6}$$

for every $y \in V(G)$. Then, we show how vertices in X can be appropriately moved out of X to obtain the required tight triple $(\emptyset, B_0, B_1)$.

Step 1: satisfying (6). Let A_*' denote $V(H(X, A_0', A_1')) - A_0' - A_1' = V(H) - X - A_0' - A_1'$. Suppose there is a vertex $y \in V(G)$ that violates (6), so $([y] - X) \not\subseteq T$ for $T \in \{A_0', A_1', A_*'\}$. Note that $y \neq s$ since $[s] \subseteq A_0'$ (as (X, A_0', A_1') is a valid triple). Define $D_0 := A_0' - [y]$, $D_1 := A_1' - [y]$ and $D_* := V(H) - X - D_0 - D_1$. Clearly, (X, D_0, D_1) is a valid triple, and $[y] - X \subseteq D_*$. We show that (X, D_0, D_1) is a tight triple; so by repeatedly applying the above modification for every vertex that violates (6), we obtain a tight triple $(X, \widetilde{A}_0, \widetilde{A}_1)$ satisfying (6).

12

Let $H_1 := H(X, A_0', A_1')$, and $H_2 := H(X, D_0, D_1)$. Let $\gamma_1 = \gamma(A_1')$, and $\gamma_2 = \gamma(D_1)$. To show that (X, D_0, D_1) is tight, it suffices to show that $p(X, D_0, D_1) \leq p(X, A_0', A_1')$. This boils down to showing that

$$\sum_{K \in \mathrm{comp}(H_2)} \left\lfloor \frac{|V(K) \cap (D_0 \cup D_1)|}{2} \right\rfloor \leq \sum_{K \in \mathrm{comp}(H_1)} \left\lfloor \frac{|V(K) \cap (A_0' \cup A_1')|}{2} \right\rfloor. \tag{7}$$

Observe that H_1 and H_2 have the same vertex-set $V(H) - X$, and H_1 is a subgraph of H_2. Every component of H_2 is either also a component of H_1, or is obtained by merging some components of H_1. Moreover, there is at most one component of H_2 that corresponds to the latter case; call this component K_2, and let $\mathcal{K}$ be the components of H_1 that are merged to form K_2. Clearly, every component of H_2 that is also a component of H_1 contributes the same to both the LHS and RHS of (7). So we focus on component K_2 and show that the contribution of K_2 to the LHS of (7) is at most the total contribution of components in $\mathcal{K}$ to the RHS of (7).

Note that all vertices in $[y] - X$ belong to the same component of H_1: since y violates (6), it is not hard to see that there is a spanning subgraph of $[y] - X$ that is not contained in $E(A_0' \cup A_1', \gamma_1)$. Every edge $e = (w, z) \in E(H_2)$ that merges two components in $\mathcal{K}$ is therefore such that $\{w, z\} \not\subseteq [y]$, and lies in $E(A_0' \cup A_1', \gamma_1) - E(D_0 \cup D_1, \gamma_2)$. So we have $\gamma_{1,e} = 0$, $w, z \in A_0' \cup A_1'$, and $\{w, z\} \not\subseteq D_0 \cup D_1$. It follows that e has exactly one end in $([y] - X) \cap (A_0' \cup A_1')$, and the other end not in $[y]$. Therefore, $|\mathcal{K}| \leq 1 + |([y] - X) \cap (A_0' \cup A_1')|$. It follows that

$$|V(K_2) \cap (D_0 \cup D_1)| = \sum_{K \in \mathcal{K}} |V(K) \cap (A_0' \cup A_1')| - |([y] - X) \cap (A_0' \cup A_1')|$$

$$\leq \sum_{K \in \mathcal{K}} |V(K) \cap (A_0' \cup A_1')| - (|\mathcal{K}| - 1).$$

Combining with Claim 6.4, we obtain that

$$\sum_{K \in \mathcal{K}} \left\lfloor \frac{|V(K) \cap (A_0' \cup A_1')|}{2} \right\rfloor \geq \left\lfloor \frac{\sum_{k \in \mathcal{K}} |V(K) \cap (A_0' \cap A_1') - (|\mathcal{K}| - 1)}{2} \right\rfloor \geq \left\lfloor \frac{|V(K_2) \cap (D_0 \cup D_1)|}{2} \right\rfloor.$$

This completes the proof that (X, D_0, D_1) is a tight triple.

Step 2: removing vertices in X. Recall that Step 1 yields a tight triple $(X, \widetilde{A}_0, \widetilde{A}_1)$ satisfying (6) for every vertex $y \in V(G)$. We now show how to convert $(X, \widetilde{A}_0, \widetilde{A}_1)$ to a tight triple $(\emptyset, B_0, B_1)$ satisfying (4) for every $x \in V(G)$. Let $\widetilde{A}_*$ denote $V(H(X, \widetilde{A}_0, \widetilde{A}_1)) - \widetilde{A}_0 - \widetilde{A}_1 = V(H) - X - \widetilde{A}_0 - \widetilde{A}_1$. Let $y \in V(G)$ be such that $[y] \cap X \neq \emptyset$. Choose some vertex $z \in [y] \cap X$. We know that (6) holds for y. Let $Y := X - \{z\}$. If $\widetilde{A}_0 \cap [y] \neq \emptyset$, set $J_0 = \widetilde{A}_0 \cup \{z\}$, else set $J_0 = \widetilde{A}_0$. If $\widetilde{A}_1 \cap [y] \neq \emptyset$, set $J_1 = \widetilde{A}_1 \cup \{z\}$, else set $J_1 = \widetilde{A}_1$. Notice that for every vertex $x \in V(G)$, all vertices in $[x] - Y$ are either in J_0, or in J_1, or in $J^* := V(H) - Y - J_0 - J_1$. We show that (Y, J_0, J_1) is a tight triple; so by repeating this process, we will eventually obtain a tight triple $(\emptyset, B_0, B_1)$ satisfying (4).

Let $H_1 := H(X, \widetilde{A}_0, \widetilde{A}_1)$ and $H_2 := H(Y, J_0, J_1)$, so H_1 is a subgraph of H_2, and $V(H_2) = V(H_1) \cup \{z\}$. Again, it suffices to show that $p(Y, J_0, J_1) \leq p(X, \widetilde{A}_0, \widetilde{A}_1)$. This amounts to showing that

$$\sum_{K \in \mathrm{comp}(H_2)} \left\lfloor \frac{|V(K) \cap (J_0 \cup J_1)|}{2} \right\rfloor \leq 1 + \sum_{K \in \mathrm{comp}(H_1)} \left\lfloor \frac{|V(K) \cap (\widetilde{A}_0 \cup \widetilde{A}_1)|}{2} \right\rfloor. \tag{8}$$

Let $\gamma_2 := \gamma(J_1)$. Let w be the other end of the non-clique edge of H incident to z. Let K_1 be the component of H_1 containing w. Let K_2 be the component of H_2 containing z. Observe that since $(X, \widetilde{A}_0, \widetilde{A}_1)$ is a tight triple, the inequality in (8) cannot be strict (otherwise, we would have $\nu(H, [s], \gamma) > p(Y, J_0, J_1)$). This implies that z cannot be an isolated vertex of H_2.

If $z \in J_0 \cup J_1$, then all edges of H_2 belonging to the clique $[y]$ lie in $E(J_0 \cup J_1, \gamma_2)$. So we must have that $V(K_2) \supseteq \{w, z\}$ (as z cannot be an isolated vertex), and $V(K_2) = V(K_1) \cup \{z\}$. Clearly, $\mathrm{comp}(H_2) - \{K_2\} = \mathrm{comp}(H_1) - \{K_1\}$, and all components in this set contribute equally to the LHS and RHS of (8), so (8) follows in this case by noting that $|V(K_2) \cap (J_0 \cup J_1)| = 1 + |V(K_1) \cap (\widetilde{A}_0 \cup \widetilde{A}_1)|$.

If $z \in J_*$, then $w \in V(K_2)$, so $V(K_1) \subseteq V(K_2)$. All vertices of $[y] - X$ belong to the same component of H_1; call this component K_1'. (One can argue that $K_1 \neq K_1'$, but we do not need this below.) Clearly, $V(K_1') \subseteq V(K_2)$, so we have $V(K_2) = V(K_1) \cup V(K_1') \cup \{z\}$. We have $\mathrm{comp}(H_2) - \{K_2\} = \mathrm{comp}(H_1) - \{K_1, K_1'\}$, and all these components contribute equally to the LHS and RHS of (8). Finally,

$$
\left\lfloor \frac{|V(K_2) \cap (J_0 \cup J_1)|}{2} \right\rfloor = \left\lfloor \frac{|V(K_1) \cap (\widetilde{A}_0 \cup \widetilde{A}_1)|}{2} + \frac{|V(K_1') \cap (\widetilde{A}_0 \cup \widetilde{A}_1)|}{2} \right\rfloor
$$
$$
\leq 1 + \left\lfloor \frac{|V(K_1) \cap (\widetilde{A}_0 \cup \widetilde{A}_1)|}{2} \right\rfloor + \left\lfloor \frac{|V(K_1') \cap (\widetilde{A}_0 \cup \widetilde{A}_1)|}{2} \right\rfloor
$$

where the inequality above follows from Claim 6.4. Thus, (8) holds in this case as well. ∎

7 Extensions

Odd trails in signed graphs. A *signed graph* is a tuple $\big(G = (V, E), \Sigma\big)$, where G is undirected and $\Sigma \subseteq E$. A set F of edges is now called odd if $|F \cap \Sigma|$ is odd. Our results extend to the more-general setting of packing and covering odd (u, v)-trails in a signed graph. In particular, Theorems 3.1, 3.2 and 3.3 *hold without any changes*. Theorem 3.2 follows simply because it utilizes Theorem 5.1, which applies to the even more-general setting of group-labeled graphs. Theorem 3.3 holds because it uses basic parity arguments: if we simply replace parity with parity with respect to Σ (i.e., instead of parity of F, we now consider parity of $|F \cap \Sigma|$), then everything goes through. Finally, as before, combining the above two results yields (the extension of) Theorem 3.1.

Odd (C, D)-trails. This is the generalization of the odd (u, v)-trails setting, where we have disjoint sets $C, D \subseteq V$. Our results yield a factor-2 gap between the the minimum number of edges needed to cover all odd (C, D)-trails and the maximum number of edge-disjoint odd (C, D)-trails.

We achieve this as follows. First, we prove a generalization of Theorem 3.2 showing that for any integer $k \geq 0$, we can either obtain k edge-disjoint odd $(C \cup D, C \cup D)$-trails, or an odd-$(C \cup D, C \cup D)$-trail cover of size at most $2k - 2$. This follows by again utilizing Theorem 5.1: we use the same gadget graph H and group $\Gamma = \mathbb{Z}_2$ as in Section 5, but now take A to be all the clique nodes corresponding to the nodes of C and D. It is not hard to see that the same translation from the group-labeled setting applies here, and yields the above generalization of Theorem 3.2. Next, we observe that Theorem 3.3 can still be applied in this more-general setting to show that if we have a collection $\widehat{\mathcal{T}}$ of k edge-disjoint odd $(C \cup D, C \cup D)$-trails, and (at least) $2k$ edge-disjoint (C, D)-paths, then we can obtain k edge-disjoint odd (C, D)-trails. A simple way of seeing this is that we may simply contract C and D to form supernodes u and v, and then utilize the earlier proof; the resulting graph may have loops, but this can be avoided by reworking the proof of Theorem 3.3 to work directly with $\widehat{\mathcal{T}}$.

References

[1] M. Chudnovsky, W. H. Cunningham, and J. Geelen. An algorithm for packing non-zero A-paths in group-labelled graphs. *Combinatorica*, 28(2): 145–161, 2008.

[2] M. Chudnovsky, J. Geelen, B. Gerards, L. A. Goddyn, M. Lohman, and P. D. Seymour. Packing non-zero A-paths in group-labelled graphs. *Combinatorica*, 26(5):521–532, 2006.

[3] R. Churchley. Personal communication. 2016.

[4] R. Churchley, B. Mohar, and H. Wu. Weak duality for packing edge-disjoint odd (u, v)-trails. In *Proceedings of the 27th SODA*, 2086–2094, 2016.

[5] J. Geelen. Personal communication. 2016.

[6] S. Ibrahimpur. *Packing and covering odd (u, v)-trails in a graph*. Master's thesis, University of Waterloo. https://uwspace.uwaterloo.ca/handle/10012/10939.

[7] S. Ibrahimpur and C. Swamy. Min-max theorems for packing and covering Odd (u, v)-trails. In *Proceedings of IPCO*, pages 279–291, 2017.

[8] C. L. Lucchesi, and D. H. Younger. A minimax theorem for directed graphs. *J. London Math. Soc*, 17:369–374, 1978.

[9] W. Mader. Über die maximalzahl kreuzungsfreier H-Wege. *Archiv der Mathematik*, 31(1), 387–402, 1978.

[10] K. Menger. Zur allgemeinen kurventheorie. *Fund. Math.*, 10:96–115, 1927.

[11] G. Pap. Packing non-returning A-paths. *Combinatorica*, 27(2):247–251, 2007.

[12] G. Pap. Packing non-returning A-paths algorithmically. *Discrete Mathematics*, 308(8):1472–1488, 2008.

[13] N. Robertson and P. Seymour. Graph minors XXIII. Nash-Williams' immersion conjecture. *J. Combin. Theory Ser. B*, 100(2):181–205, 2010.

[14] A. Schrijver. A short proof of Mader's $\mathcal{S}$-paths theorem. *JCTB*, 82:319–321, 2001.

[15] A. Schrijver. Combinatorial Optimization: Polyhedra and Efficiency. *Algorithms and Combinatorics. Springer Berlin Heidelberg*, 2003.

[16] A. Schrijver, and P. D. Seymour. Packing odd paths. *JCTB*, 62:280–288, 1994.

[17] T. Zaslavsky. Glossary of signed and gain graphs and allied areas. *Electron. J. Combin.*, 5:Dynamic Surveys 9, 41 pp. (electronic), 1998.

A   Tight example for Theorem 3.1

We prove that for the graph $G = (V, E)$ shown in Fig. 1, we have $\nu(u, v) = k$, and $\tau(u, v) = 2k + 1$. For any $i = 1, \ldots, k$, let B_i denote the subgraph of G induced by the vertices $\{u, a_i, b_i, \ldots, g_i, h_i, v\}$.

The size of a minimum (u, v)-cut in G_k is $2k+1$, so $\tau(u, v) \leq 2k+1$. It is easy to see by inspection that any cover Z of G must contain at least 2 edges from each B_i, as otherwise $B_i - Z$ will have an odd (u, v)-path. Suppose, for a contradiction, that $|Z| \leq 2k$. Then $G - Z$ contains a (u, v)-path P, and Z contains exactly 2 edges from each B_i. We argue that $G - Z$ contains an odd (u, v)-trail, yielding a contradiction and thereby showing that $\tau(u, v) = 2k + 1$. Z must contain exactly one edge from each of the two edge-disjoint triangles $\{u, a_1, b_1\}$ and $\{v, g_1, h_1\}$ in B_1 containing u and v respectively, as otherwise, one can combine

P and such a triangle to obtain an odd (u, v)-trail in $G - Z$. But then $B_1 - Z$ is connected, and contains all edges incident to the nodes c_1, d_1, e_1, f_1, which implies that $B_1 - Z$ contains an odd (u, v)-path.

We now argue that $\nu(u, v) = k$. We have $\nu(u, v) \geq k$, as each B_i contains an odd (u, v)-path. Among all collections of $\nu(u, v)$ edge-disjoint odd (u, v)-trails, let $\mathcal{T}$ be one with the fewest number of edges. Let $P_1, \ldots, P_r$ be the odd (u, v)-paths in $\mathcal{T}$, and $T_1, \ldots, T_\ell$ be the odd (u, v)-trails in $\mathcal{T}$ that are not (u, v)-paths. Observe that each (u, v)-path $P \in \{P_1, \ldots, P_r\}$ must be contained in some subgraph B_i, and $B_i - P$ is disconnected. Next, observe that every odd (u, v)-trail $T \in \{T_1, \ldots, T_\ell\}$ consists of an even (u, v)-path concatenated with an odd (u, u)-trail or an odd (v, v)-trail, say Z. Moreover, Z must in fact be a triangle containing u or v due to the minimality of $\mathcal{T}$. (Suppose Z is an odd (u, u)-trail. Then Z must contain an odd cycle contained in some subgraph B_i and the edges in B_i incident to u; replacing Z in the trail T with the triangle in B_i incident to u yields an odd (u, v)-trail T' with $|E(T')| < |E(T)|$ that is edge-disjoint from $\mathcal{T} - \{T\}$, contradicting the minimality of $\mathcal{T}$. The case where Z is an odd (v, v)-trail is completely analogous.) Thus, every trail in $\{T_1, \ldots, T_\ell\}$ consists of an even (u, v)-path and a triangle incident to u or v. Without loss of generality, let $\lceil \frac{\ell}{2} \rceil$ of these trails contain triangles incident to u. Let $\mathcal{Z}_u$ be the collection of triangles contained in the trails $\{T_1, \ldots, T_\ell\}$ that are incident to u, so $|\mathcal{Z}_u| \geq \lceil \frac{\ell}{2} \rceil$.

Consider the graph $G' = G - \bigcup_{i=1}^{r} P_i - \bigcup_{Z \in \mathcal{Z}_u} Z$. The (u, v) connectivity in G' is at most $2k + 1 - 2r - 2|\mathcal{Z}_u|$: deleting each P_i leaves some $B_{i'}$ disconnected, and so deleting either a P_i or a triangle in $\mathcal{Z}_u$ decreases the (u, v)-connectivity by 2. Also, G' contains ℓ (even) (u, v)-paths. So we have $\ell \leq 2k + 1 - 2r - 2 \lceil \frac{\ell}{2} \rceil$, which implies that $\ell + r \leq \frac{2k+1}{2}$, and hence that $\ell + r \leq k$. Thus, $\nu(u, v) = k$.

B Tight example for Theorem 3.3

We now provide an example showing that Theorem 3.3 is tight; in particular, the requirement that $\lambda(u, v) \geq 2|\widehat{\mathcal{T}}|$ cannot be weakened. Consider the graph G shown in Figure 4. $\mathcal{T} = \{ux_i y_i v\}_{i=1}^{k-1} \cup \{uz_1 z_2 u\} \cup \{vz_5 z_6 v\}$ is a collection of $k + 1$ edge-disjoint odd $(\{u, v\}, \{u, v\})$-trails in G. Observe that $\lambda(u, v) = 2k + 1 = 2|\mathcal{T}| - 1$. But G contains only k edge-disjoint odd (u, v)-trails: the subgraph induced by $\{u, v\}$, the z_is and w contains only one odd (u, v)-trail.
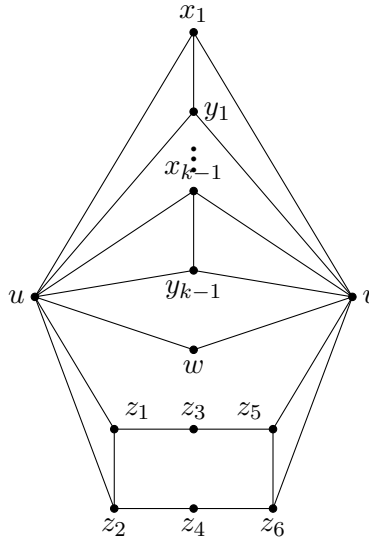

Figure 4: Example showing that Theorem 3.3 is tight.

C Lower bound for covers constructed in [4]

Churchley et al. [4] construct odd-(u,v)-trail covers of the form $\delta(X)$, where X is a u-v cut, or of the form $\delta(X) \cup (E(X) \setminus F)$, where $\{u,v\} \subseteq X$, and $(X, E(X) \cap F)$ is bipartite with u and v on the same side of the bipartition. The following is a family $\{H_k\}_{k \geq 1}$ of graphs where, for H_k we have $\nu(u,v) = k$, $\tau(u,v) = 2k$, but any cover of the above form has size at least $3k$. This shows that covers of the above form cannot yield a bound better than 3 on the covering-vs-packing ratio for odd (u,v) trails. (The graph H_k has parallel edges; if desired, the parallel edges can be eliminated by replacing each parallel edge with a length-3 path (to preserve trail parities).)
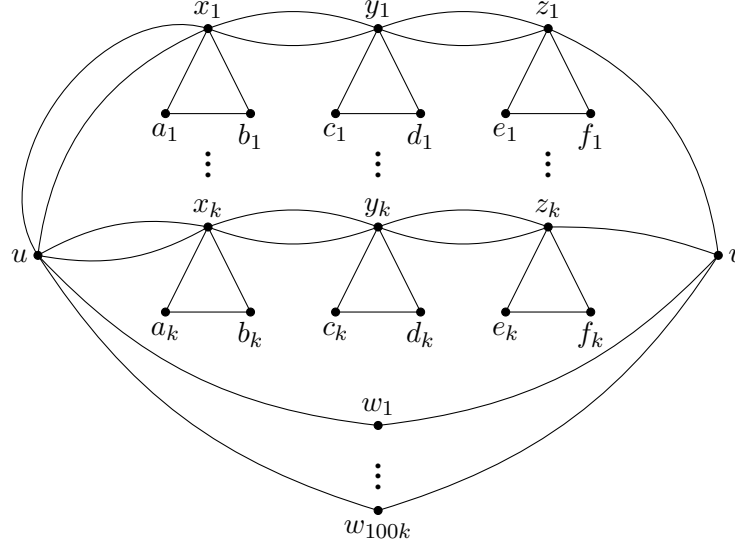

Figure 5: Graph H_k with $\nu(u,v) = k$, $\tau(u,v) = 2k$. We have $|\delta(X)| \cup |E(X) \setminus F| \geq 3k$ for any X, F such that $\{u,v\} \subseteq X$ and $(X, E(X) \cap F)$ is bipartite with u, v on the same side.

Let C_i denote the vertex-set $\{u, v, x_i, y_i, z_i, a_i, b_i, c_i, d_i, e_i, f_i\}$. Observe that every odd trail T in G uses an edge from at least one of the odd cycles in $\{x_i a_i b_i x_i\}_{i=1}^k \cup \{y_i c_i d_i y_i\}_{i=1}^k \cup \{z_i e_i f_i z_i\}_{i=1}^k$. For each $1 \leq i \leq k$, removing the two parallel edges between u and x_i along with the edge vz_i separates $C_i \setminus \{u,v\}$ from $\{u,v\}$, so at most one trail can use the nodes of $C_i \setminus \{u,v\}$. It follows that the packing number is at most k. It is trivial to find k edge-disjoint odd (u,v)-trails in G and hence $\nu(u,v) = k$. Next, the covering number is at most $2k$ since the edge-set consisting of all the edges (including parallel copies) between u and x_i for all $1 \leq i \leq k$ is an odd (u,v)-trail cover. Deleting at most one edge from the induced subgraph $G[C_i]$ does not destroy all the odd (u,v)-trails using at least one odd cycle from $G[C_i]$ hence we need at least $2k$ edges in any cover. Thus, $\tau(u,v) = 2k$.

Since $\lambda(u,v) \geq 100k$, every u-v cut X has $|\delta(X)| \geq 100k$. Finally, we verify that $|\delta(X)| \cup |E(X) \setminus F| \geq 3k$ for any X, F such that $\{u,v\} \subseteq X$ and $(X, E(X) \cap F)$ is bipartite with u, v on the same side. For this it is convenient to consider the graph shown in Fig. 6, which is the graph in Fig. 5, with u and v identified to form s.

Let $C_i' = \{s, x_i, y_i, z_i, a_i, b_i, c_i, d_i, e_i, f_i\}$. For any node-set X with $s \in X$, we have $|\delta(X)| \geq \sum_{i=1}^k |E(C_i' \cap X, C_i' \setminus X)|$. For F such that $(X, E(X) \cap F)$ is bipartite, we lower bound $|E(X) \setminus F|$ by counting the number of triangles (i.e., $\{x_i, a_i, b_i\}$, $\{y_i, c_i, d_i\}$, or $\{z_i, e_i, f_i\}$ for some i) in $E(X)$. This contribution is decoupled across the C_i's, so overall each C_i' contributes $|E(C_i' \cap X, C_i' \setminus X)| +$ (number of triangles in $C_i' \cap X$). We argue that this contribution is at least 3 for each C_i'. Each C_i' is 2-edge-connected, and the cuts of size 2 in C_i' are those that contain two edges of one of the C_i'-triangles. So
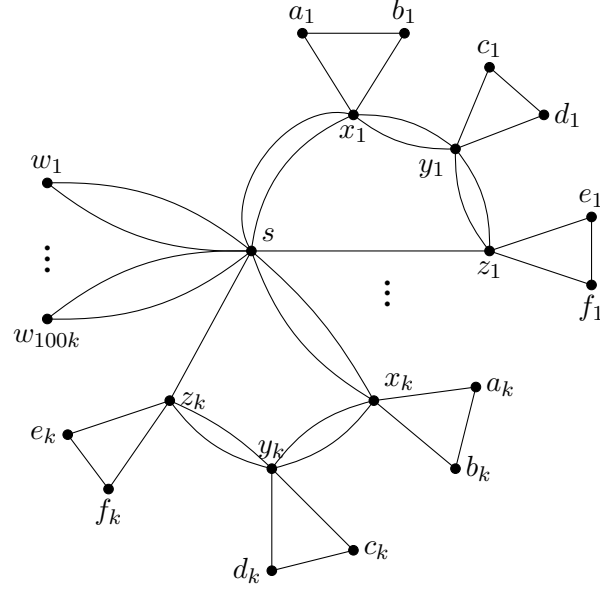
Figure 6: Graph from Figure 5 with vertices u and v identified to form s.

if $|E(C_i' \cap X, C_i' \setminus X)| \le 2$, then $C_i' \cap X$ contains at least one triangle, and the contribution is at least 3.

D   Tight example for Theorem 3.2

The example is the same as the graph G in Fig. 6, which, recall is obtained from the graph G' in Fig. 5 by identifying u and v. Since $\tau(u, v; G') = 2k$, we have $\tau(s, s; G) \ge 2k$. We argue that $\nu(s, s, G) \le k$, showing that the two conclusions of the theorem are tight as seen by the inputs k and $k + 1$.

The upper bound on $\nu(s, s, G)$ follows from the same reasoning as before. Recall that we have $C_i' = \{s, x_i, y_i, z_i, a_i, b_i, c_i, d_i, e_i, f_i\}$. Every odd trail T in G must use an edge in at least one of the triangles. For each $1 \le i \le k$, the two parallel edges between s and x_i and sz_i separate the triangle vertices of C_i' from s, so each C_i' contributes at most 1 to $\nu(s, s, G)$. Therefore, $\nu(s, s, G) \le k$.