Electronic Colloquium on Computational Complexity, Report No. 59 (2017)

The Matching Problem in General Graphs is in Quasi-NC

Ola Svensson* Jakub Tarnawski*

April 6, 2017

Abstract

We show that the perfect matching problem in general graphs is in quasi-NC. That is, we give a deterministic parallel algorithm which runs in $O(\log^3 n)$ time on $n^{O(\log^2 n)}$ processors. The result is obtained by a derandomization of the Isolation Lemma for perfect matchings, which was introduced in the classic paper by Mulmuley, Vazirani and Vazirani [1987] to obtain a Randomized NC algorithm.

Our proof extends the framework of Fenner, Gurjar and Thierauf [2016], who proved the analogous result in the special case of bipartite graphs. Compared to that setting, several new ingredients are needed due to the significantly more complex structure of perfect matchings in general graphs. In particular, our proof heavily relies on the laminar structure of the faces of the perfect matching polytope.

*School of Computer and Communication Sciences, EPFL.
Email: {ola.svensson,jakub.tarnawski}@epfl.ch.
Supported by ERC Starting Grant 335288-OptApprox.

ISSN 1433-8092

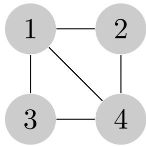

$$T(G) = \begin{pmatrix} 0 & X_{12} & X_{13} & X_{14} \\ -X_{12} & 0 & 0 & X_{24} \\ -X_{13} & 0 & 0 & X_{34} \\ -X_{14} & -X_{24} & -X_{34} & 0 \end{pmatrix}$$

Figure 1: Example of a Tutte matrix of an undirected graph.

1 Introduction

The perfect matching problem is a fundamental question in graph theory. Work on this problem has contributed to the development of many core concepts of modern computer science, including linear-algebraic, probabilistic and parallel algorithms. Edmonds [Edm65b] was the first to give a polynomial-time algorithm for it. However, half a century later, we still do not have full understanding of the deterministic parallel complexity of the perfect matching problem. In this paper we make progress in this direction.

We consider a problem to be efficiently solvable in parallel if it has an algorithm which uses polylogarithmic time and polynomially many processors. More formally, a problem is in the class NC if it has uniform circuits of polynomial size and polylogarithmic depth. The class RNC is obtained if we also allow randomization.

We study the decision version of the problem: given an undirected simple graph, determine whether it has a perfect matching – and the search version: find and return a perfect matching if one exists. The decision version was first shown to be in RNC by Lovász [Lov79]. The search version has proved to be more difficult, as a graph may contain exponentially many perfect matchings and it is necessary to coordinate the processors so that they identify and output the same one; our vocabulary for this is that we want to *isolate* one matching. This version was found to be in RNC by Karp, Upfal and Wigderson [KUW86] and Mulmuley, Vazirani and Vazirani [MVV87]. All these algorithms are randomized, and it remains a major open problem to determine whether randomness is required, i.e., whether either version is in NC.

A very successful approach to the perfect matching problem is the linear-algebraic one. It involves the Tutte matrix associated with a graph $G = (V, E)$, which is a $|V| \times |V|$ matrix defined as follows (see Figure 1 for an example):

$$T(G)_{u,v} = \begin{cases} X_{(u,v)} & \text{if } (u,v) \in E \text{ and } u < v, \\ -X_{(v,u)} & \text{if } (u,v) \in E \text{ and } u > v, \\ 0 & \text{if } (u,v) \notin E, \end{cases}$$

where $X_{(u,v)}$ for $(u,v) \in E$ are variables. Tutte's theorem [Tut47] says that $\det T(G) \neq 0$ if and only if G has a perfect matching. This is great news for parallelization, as computing determinants is in NC [Csa76, Ber84]. However, the matrix is defined over a ring of indeterminates, so randomness is normally used in order to test if the determinant is nonzero. One approach is to replace each indeterminate by a random value from a large field; this leads, among others, to the fastest known (single-processor) running times for dense graphs [MS04, Har09].

A second approach, adopted by Mulmuley, Vazirani and Vazirani [MVV87] for the search version, is to replace the indeterminates by randomly chosen powers of two. Namely, for each edge (u,v), a random weight $w(u,v) \in \{1, 2, ..., 2|E|\}$ is selected, and we substitute $X_{(u,v)} := 2^{w(u,v)}$. Now, let us make the crucial assumption that one perfect matching M is **isolated**, in the sense that it is the *unique minimum-weight perfect matching* (minimizing $w(M)$). Then $\det T(G)$ remains nonzero after the substitution: one can show that M contributes a term $\pm 2^{2w(M)}$ to $\det T(G)$, whereas all other terms are multiples of $2^{2w(M)+1}$ and thus they cannot cancel $2^{2w(M)}$ out. The determinant can still be computed in NC as all entries $2^{w(u,v)}$ of the matrix are of polynomial bit-length, and so we have a parallel algorithm for the decision version. An algorithm for the search version also follows: for every edge in parallel, test whether removing it causes this least-significant digit $2^{2w(M)}$ in the determinant to disappear; output those edges for which it does.

The fundamental claim in [MVV87] is that assigning random weights to edges does indeed isolate one matching with high probability. This is known as the Isolation Lemma and turns out to be true in the much more general setting of arbitrary set families:

Lemma 1.1 (Isolation Lemma). *Let $\mathcal{M} \subseteq 2^E$ be any nonempty family of subsets of a universe $E = \{1, 2, ..., |E|\}$. Suppose we define a weight function $w : E \to \{1, 2, ..., 2|E|\}$ by selecting each $w(e)$ for $e \in E$ independently and uniformly at random. Then with probability at least $1/2$, there is a* unique *set $M \in \mathcal{M}$ which minimizes the weight $w(M) = \sum_{e \in M} w(e)$.*

We call such a weight function w *isolating*. We take $\mathcal{M}$ in Lemma 1.1 to be the set of all perfect matchings.

Since Lemma 1.1 is the only randomized ingredient of the RNC algorithm, the natural approach to showing that the perfect matching problem is in NC is the derandomization of the Isolation Lemma. That is, we would like a set of polynomially many weight functions (with polynomially bounded values) which would be guaranteed to contain an isolating one. To get an NC algorithm, we should be able to generate this set efficiently in parallel; then we can try all weight functions simultaneously.

However, derandomizing the Isolation Lemma turns out to be a challenging open question. It has been done for certain classes of graphs: strongly chordal [DK98], planar bipartite [DKR10, TV12], or graphs with a small number of perfect matchings [GK87, AHT07].[1] The general set-family setting of the Isolation Lemma is also related to circuit lower bounds and polynomial identity testing [AM08].

Recently, in a major development, Fenner, Gurjar and Thierauf [FGT16] have almost derandomized the Isolation Lemma for bipartite graphs. Namely, they define a family of weight functions which can be computed obliviously (only using the number of vertices n) and prove that for any bipartite graph, one of these functions is isolating. Because their family has quasi-polynomial size and the weights are quasi-polynomially large, this has placed the perfect bipartite matching problem in the class $\mathsf{quasi\text{-}NC}$.

Nevertheless, the general-graph setting of the derandomization question (either using the Isolation Lemma or not) remained wide open. Even in the planar case, with NC algorithms for bipartite planar and small-genus graphs having been known for a long time [MN89, MV00], we knew no $\mathsf{quasi\text{-}NC}$ algorithm for non-bipartite graphs.[2] In general, the best known upper bound on the size of uniform circuits with polylogarithmic depth was exponential.

We are able to nearly bridge this gap in understanding. The main result of our paper is the following:

Theorem 1.2. *For any number n, we can in $\mathsf{quasi\text{-}NC}$ construct $n^{O(\log^2 n)}$ weight functions on $\{1, 2, ..., \binom{n}{2}\}$ with weights bounded by $n^{O(\log^2 n)}$ such that for any graph on n vertices, one of these weight functions isolates a perfect matching (if one exists).*

[1] More generally, there has been much interest in obtaining NC algorithms for the perfect matching problem on restricted graph classes (not necessarily using the Isolation Lemma), e.g.: regular bipartite [LPV81], P_4-tidy [Par98], dense [DHK93], convex bipartite [DS84], claw-free [CNN89], incomparability graphs [KVV85].

[2] Curiously, an NC algorithm to count the number of perfect matchings in a planar graph is known [Kas67, Vaz89], which implies an algorithm for the decision version; however, it is open to give an NC algorithm for the search version.

The results of [MVV87] and Theorem 1.2 together imply that the perfect matching problem (both the decision and the search variant) in general graphs is in quasi-NC; see Section 2.1 for more details on this. The implied algorithm is very simple; the complexity lies in the analysis, i.e., proving that one of the weight functions is isolating (Theorem 4.11).

In what follows, we first give an overview of the framework in [FGT16] for bipartite graphs. We then explain how we extend the framework to general graphs. Due to the more complex structure of perfect matchings in general graphs, we need several new ideas. In particular, we exploit structural properties of the perfect matching polytope.

1.1 Isolation in bipartite graphs

In this section we shortly discuss the elegant framework introduced by Fenner, Gurjar and Thierauf [FGT16], which we extend to obtain our result.

If a weight function w is not isolating, then there exist two minimum-weight perfect matchings, and their symmetric difference consists of alternating cycles. In each such cycle, the total weight of edges from the first matching must be equal to the total weight of edges from the second matching (as otherwise we could obtain another matching of lower weight). The difference between these two total weights is called the *circulation* of the cycle. By the above, if all cycles have nonzero circulation, then w is isolating. It is known how to obtain weight functions which satisfy a polynomial number of such non-equalities (see Lemma 3.4); unfortunately, a graph may have an exponential number of cycles.

The crucial idea of [FGT16] is to build the weight function in $\log n$ rounds. In the first round, we find a weight function with the property that each cycle *of length 4* has nonzero circulation; this is possible since there are at most n^4 such cycles. We apply this function and from now on consider only those edges which belong to a minimum-weight perfect matching. Crucially, it turns out that in the subgraph obtained this way, all cycles of length 4 have disappeared – this follows from the simple structure of the bipartite perfect matching polytope (a face is simply the bipartite matching polytope of a subgraph) and fails to hold for general graphs. In the second round, we start from this subgraph and apply another weight function which ensures that all even cycles of length up to 8 have nonzero circulation (one proves that there are again n^4 many since the graph contains no 4-cycles). Again, these cycles disappear from the next subgraph, and so on. After $\log n$ rounds, the current subgraph has no cycles, i.e., it is a perfect matching. The final weight function is obtained by combining the $\log n$ polynomial-sized weight functions. To get a parallel algorithm, we need to simultaneously try each such possible combination, of which there are quasi-polynomially many.

This result has later been generalized by Gurjar and Thierauf [GT16] to the linear matroid intersection problem – a natural extension of bipartite matching. From the work of Narayanan, Saran and Vazirani [NSV94], who gave an RNC algorithm for that problem (also based on computing a determinant), it again follows that derandomizing the Isolation Lemma implies a quasi-NC algorithm.

1.2 Challenges of non-bipartite graphs

We find it useful to look at the method explained in the previous section from a polyhedral perspective (also used by [GT16]). We begin from the set of all perfect matchings, of which we take the convex hull: the perfect matching polytope. After applying the first weight function, we want to consider only those perfect matchings which minimize the weight; this is exactly the definition of a face of the polytope. In the bipartite case, any face was characterized by just taking a subset of edges (i.e., making certain constraints $x_e \geq 0$ tight), so we could simply think about recursing on a smaller subgraph. This was used to show that any cycle whose circulation

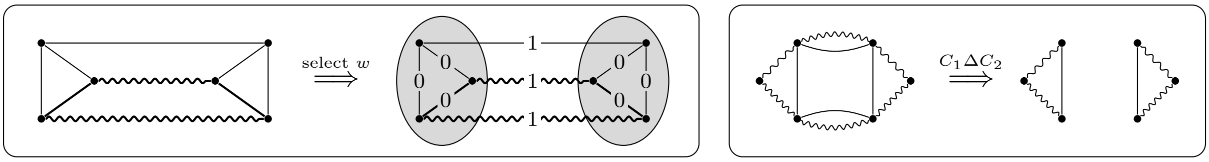

Main difficulty   Minor difficulty

Figure 2: An illustration of the difficulties of derandomizing the Isolation Lemma for general graphs as compared to bipartite graphs.
On the left: in trying to remove the bold cycle, we select a weight function w such that the circulation of the cycle is $1-0+1-0 \neq 0$. By minimizing over w we obtain a new, smaller subface – the convex hull of perfect matchings of weight 1 – but every edge of the cycle is still present in one of these matchings. The cycle has only been eliminated in the following sense: it can no longer be obtained in the symmetric difference of two matchings in the new face (since none of them select both swirly edges). The vertex sets drawn in gray represent the new tight odd-set constraints that describe the new face (indeed: for a matching to have weight 1, it must take only one edge from the boundary of a gray set). We will say that the cycle does not *respect* the gray vertex sets (see Section 3).
On the right: two even cycles whose symmetric difference contains no even cycle.

has been made nonzero will not retain all of its edges in the next subgraph. The progress we made in the bipartite case could be measured by the minimum length of a cycle in the current subgraph, which doubled as we moved from face to subface.

Unfortunately, in the non-bipartite case, the description of the perfect matching polytope is more involved (see Section 2.2). Namely, moving to a new subface may also cause new tight *odd-set constraints* to appear; these require that, for an odd set $S \subseteq V$ of vertices, exactly one edge of a matching should cross S. This complicates our task; see the left part of Figure 2 for an example (the same one as given by [FGT16] to demonstrate the difficulty of the general-graph case). Now a face is described by not only a subset of edges, but also a family of tight odd-set constraints. Thus we can no longer guarantee that any cycle whose circulation has been made nonzero will disappear from the support. Our idea of what it means to remove a cycle needs to be refined (see Section 3), as well as the measure of progress we use to prove that a single matching is isolated after $\log n$ rounds (see Section 4).

Another difficulty, of a more technical nature, concerns the counting argument used to prove that a graph with no cycles of length at most λ contains only polynomially many cycles of length at most 2λ. In the bipartite case, the symmetric difference of two cycles (which are even) contains a simple cycle (which is also even, short, and thus should not exist); this enables a simple checkpointing argument. In the general case, we are still only interested in removing even cycles, but the symmetric difference of two even cycles may not contain an even simple cycle (see the right part of Figure 2). This forces us to remove not only even simple cycles, but all even walks, which may contain repeated edges (we call these *alternating circuits* – see Definition 3.1), and to rework the counting scheme, obtaining a bound of n^{17} rather than n^4 (see Lemma 5.4). Moreover, instead of simple graphs, we work on node-weighted multigraphs, which arise by contracting certain tight odd-sets.

1.3 Our approach

This section is a high-level, idealized explanation of how to deal with the main difficulty (see the left part of Figure 2); we ignore the more technical one in this description.

Removing cycles which do not cross a tight odd-set. As discussed in Section 1.2, when moving from face to subface we cannot guarantee that, for each even cycle whose circulation we make nonzero, one of its edges will be absent in the new face. However, this will at least be true for cycles that do not cross any odd-set tight for the new face.[3] This implies that if we apply $\log n$ weight functions in succession, then the resulting face will not contain in its support any even cycle that crosses no tight odd-set (by the same argument as in Section 1.1). This is less than we need, but it is a good first step. If, at this point, there were no tight sets, then we would be done, as we would have removed all cycles. However, in general there will still be cycles crossing tight sets, which make our task more difficult.

Decomposition into two subinstances. To deal with the tight odd-sets, we will make use of two crucial properties. The first property is easy to see: once we fix the single edge e in the matching which crosses a tight set S, the instance breaks up into two *independent* subinstances – that is, every perfect matching in the graph which contains e is the union of: the edge e, a perfect matching on the vertex set S (ignoring the S-endpoint of e), and a perfect matching on the vertex set $V \setminus S$ (ignoring the other endpoint of e).

This will allow us to employ a divide-and-conquer strategy: to isolate a matching in the entire graph, we will take care of both subinstances and of the cut separating them. We formulate the task of dealing with such a subinstance (a subgraph induced on an odd-cardinality vertex set) as follows: we want that, once the (only) edge of a matching which lies on the boundary of the tight odd-set is fixed, the entire matching inside the set is uniquely determined (we will then call this set *contractible* – see Definition 4.1). This can be seen as a generalization of our isolation objective to subgraphs with an odd number of vertices. If we can get that for the tight set and for its complement, then each edge from the cut separating them induces a unique perfect matching in the graph; therefore there are at most n^2 perfect matchings left in the current face. Now, in order to isolate the entire graph, we only need a weight function w which assigns different weights to all these matchings. This can be written as n^4 linear non-equalities on w, and we can generate a weight function w satisfying all of them (see Lemma 3.4).[4] While it is not clear how to continue this scheme beyond the first level or why we could hope to have a low depth of recursion, we will soon explain how we utilize this basic strategy in the sequel.

Laminarity. The second crucial property that we utilize is that the family of odd-set constraints tight for a face exhibits good structural properties; it is known that a *laminar* family of odd sets is enough to describe any face (see Section 2.2). This enables a scheme where we use this family to make progress in a bottom-up fashion. This is still challenging as the family does not stay fixed as we move from face to face. The good news is that it can only increase: whenever a new tight odd-set constraint appears which is not spanned by the previous ones, we add that odd-set to our laminar family.

[3] This is because if there are no tight odd-set constraints, then our faces behave as in the bipartite case. Now intuitively, if we only consider those cycles which do not cross any tight set, then we can remove them using the same arguments as in that case.

[4] This can also be seen as assigning nonzero circulation to n^4 cycles.

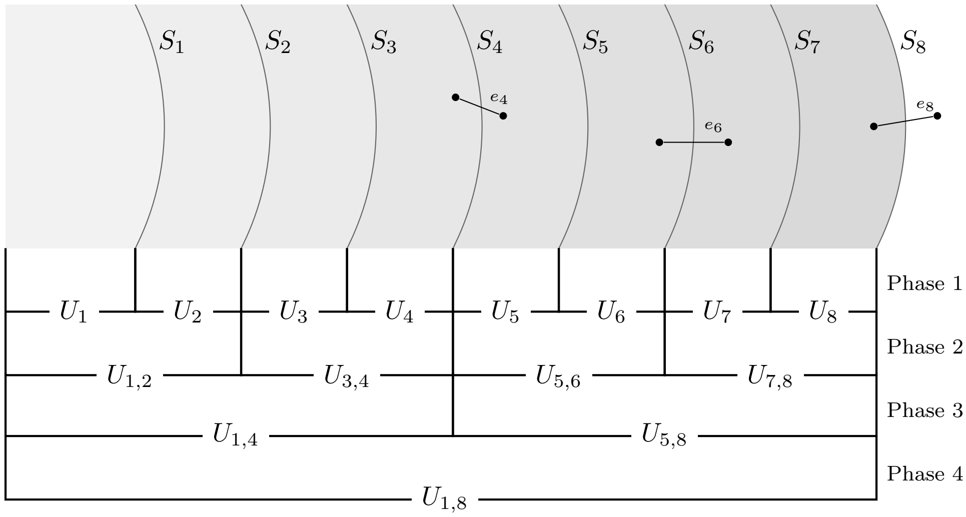

Figure 3: Example of a chain consisting of 8 tight sets, and our divide-and-conquer argument.

Chain case. To get started, let us first discuss the special case where the family of tight constraints is a chain, i.e., an increasing sequence of odd-sets $S_1 \subsetneq S_2 \subsetneq \ldots \subsetneq S_\ell$. For this introduction, assume $\ell = 8$; see Figure 3. (This will be an informal and simplified version of the proof of Lemma 5.6.) Denote by $U_1, \ldots, U_8$ the *layers* of this chain, i.e., $U_1 = S_1$ and $U_p = S_p \setminus S_{p-1}$ for $p = 2, 3, \ldots, 8$. Suppose this chain describes the face that was obtained by applying $\log n$ weight functions as described above; then there is no cycle that lies inside a single layer U_p.

Notice that every layer U_p is of even size and it touches two boundaries of tight odd-sets: S_{p-1} and S_p (that is, $\delta(U_p) \subseteq \delta(S_{p-1}) \cup \delta(S_p)$). Any perfect matching (in the current face) will have one edge from $\delta(S_{p-1})$ and one edge from $\delta(S_p)$ (possibly the same edge), therefore U_p will have two (or zero) boundary edges in the matching. (An exception is U_1, which is odd, only touches S_1 and will have one boundary edge in the matching.) This motivates us to generalize our isolation objective to layers as follows: we say that a layer U_p is *contractible* if choosing an edge from $\delta(S_{p-1})$ and an edge from $\delta(S_p)$ induces a unique matching inside U_p. This definition naturally extends to layers of the form $S_r \setminus S_{p-1} = U_p \cup U_{p+1} \cup \ldots \cup U_r$, which we will denote by $U_{p,r}$.

Recall that we have ensured that there is no cycle that lies inside a single layer $U_p = U_{p,p}$. It follows that these layers are contractible.[5] Let us say that this was the first phase of our approach (see Figure 3). In the second phase, we want to ensure contractibility for double layers: $U_{1,2}$, $U_{3,4}$, $U_{5,6}$ and $U_{7,8}$. In general, we double our progress in each phase: in the third one we deal with the quadruple layers $U_{1,4}$ and $U_{5,8}$, and in the fourth phase we deal with the octuple layer $U_{1,8}$.

Let us now describe a single phase. Take e.g. the layer $U_{5,8}$ and two boundary edges $e_4 \in \delta(S_4)$ and $e_8 \in \delta(S_8)$ (see Figure 3); we want to have only a unique matching in $U_{5,8}$ including these edges. Now we will realize our divide-and-conquer approach. Note that the layers $U_{5,6}$ and $U_{7,8}$ have already been dealt with (made contractible) in the previous phase. Therefore, for each choice of boundary edge $e_6 \in \delta(S_6)$ for the matching, there is a unique matching inside both of these layers. Just like previously, this implies that there are only n^2

[5] This is because two different matchings (but with the same boundary edges) in the current face would induce an alternating cycle in their symmetric difference.

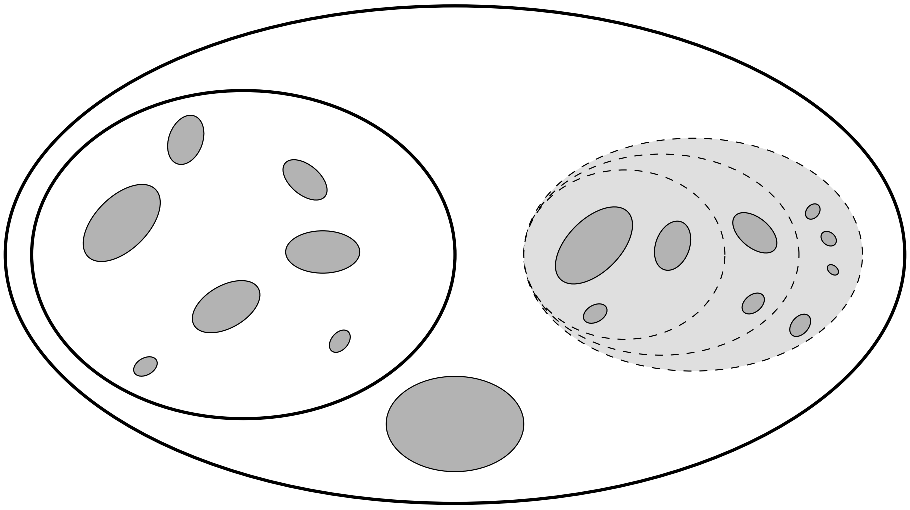

Figure 4: Example of a general laminar family.
Dark-gray sets are of size at most λ and thus contractible.
Dashed sets are of size more than λ but at most 2λ; they must form chains (due to the cardinality constraints). We make them contractible in the first step. Then we contract them (so now all light-gray and dark-gray sets are contracted).
Thick sets are of size more than 2λ. For the second step, we erase the edges on their boundaries. Then we remove cycles of length up to 2λ from the resulting instance (the *contraction*), which has no tight odd-sets (and no cycles of length up to λ).

matchings using e_4 and e_8 in the layer $U_{5,8}$, and we can select a weight function that isolates one of them.[6]

By generalizing this strategy in the natural way, we can deal with any chain in $\log \ell \leq \log n$ phases, even if it consists of $\Omega(n)$ tight sets.[7] (By applying one more weight function, we can isolate a unique perfect matching in the entire graph using the same arguments as above.)

General case. Of course, there is no reason to expect that the laminar family of tight cuts we obtain after applying the initial $\log n$ weight functions will be a chain. It also does not seem easy to directly generalize our inductive scheme from a chain to an arbitrary family. Therefore we put forth a different progress measure, which allows us to make headway even in the absence of such a favorable odd-set structure.

Since a laminar family can be represented as a tree, we might think about a bottom-up strategy based on it; however, we cannot deal with its nodes level-by-level, since it may have height $\Omega(n)$ and we can only afford $\text{poly}(\log n)$ many phases. Instead, we will first deal with all tight odd-sets of size up to 4, then up to 8, then up to 16 and so on (by making them contractible). At the same time, we will also remove all even cycles of length up to 4, then up to 8 and so on.[8] These two components of our progress measure, which we call λ-goodness, are mutually beneficial, as we will see below.

[6] We actually select only one function per phase, which works simultaneously for all layers $U_{p,r}$ in this phase (here: $U_{1,4}$ and $U_{5,8}$) and all pairs of boundary edges e_{p-1} and e_r.

[7] In the general proof, we do not quite have a binary tree structure like in the example. Instead, in the t-th phase, we deal with all layers $U_{p,r}$ having $1 \leq p \leq r \leq \ell$ with $r - p \leq 2^{t-1} - 1$. This makes our proof simpler if ℓ is not a power of two.

[8] As discussed in Section 1.2 and Figure 2, the meaning of the term *remove* needs to be refined, as we cannot hope to always delete an edge of the cycle from the support of the current face.

Making odd-sets contractible enables us not only to achieve progress, but also to simplify our setting. A contractible tight set can be, for our purposes, thought of as a single vertex – much like a blossom in Edmonds' algorithm. This is because such a set has exactly one boundary edge in a perfect matching (as does a vertex), and choosing that edge determines the matching in the interior. We will contract such sets (hence the name).

Suppose that our current face is already λ-good. Roughly, this means that we have made odd-sets of size up to λ (which we will call small) contractible and removed cycles of length up to λ. Now we want to obtain a face which is 2λ-good.

The first step is to make odd-sets of size up to 2λ contractible. Let us zoom in on one such odd-set (specifically, a maximal set of size at most 2λ – see the largest dashed set in Figure 4). Once we have contracted all the small sets into single vertices, all interesting sets are now of size more than λ but at most 2λ, and any laminar family consisting of such sets must be a chain, since a set of such size cannot have two disjoint subsets of such size (see Figure 4). But this is the chain case that we have already solved!

Having made odd-sets of size up to 2λ contractible, we can contract them. The second step is now to remove cycles of length up to 2λ. However, here we do not need to care about those cycles which cross an odd-set S of size larger than 2λ – the reason being, intuitively, that in our technical arguments we define the length of a cycle based on the sizes of sets that it crosses, and thus such a cycle actually becomes longer than 2λ. In other words, we can think about removing cycles of length up to 2λ from a version of the input graph where all small odd-sets have been contracted and all larger ones have had their boundaries erased (see Figure 4). We call this version the *contraction* (see Definition 4.5). Our λ-goodness progress measure (see Definition 4.7) is actually defined in terms of cycles in the contraction.

Now the second step is easy: we just need to remove all cycles of length up to 2λ from the contraction, which has no tight odd-sets and no cycles of length up to λ – a simple scenario, already known from the bipartite case. Applying one weight function is enough to do this.

Finally, what does it mean for us to remove a cycle? When we make a cycle's circulation nonzero, it is then eliminated from the new face in the following sense: either one of its edges disappears from the support of the face (recall that this is what always happened in the bipartite case), or a new tight odd-set appears, with the following property: the cycle crosses the set with fewer (or more) even-indexed edges than odd-indexed edges (see the example in Figure 2). In short, we say that the cycle does not *respect* the new face (see Section 3). This notion of removal makes sense when viewed in tandem with the contraction, because once a cycle crosses a set in the laminar family, there are two possibilities in each phase: either this set is large – then its boundary is not present in the contraction, which cancels the cycle, or it is small – then it is contracted and the cycle also disappears (for somewhat more technical reasons).

To reiterate, our strategy is to simultaneously remove cycles up to a given length and make odd-sets up to a given size contractible. We can do this in $\log n$ phases. In each such phase we need to apply a sequence of $\log n$ weight functions in order to deal with a chain of tight odd-sets (as outlined above). In all, we are able to isolate a perfect matching in the entire graph using a sequence $O(\log^2 n)$ weight functions with polynomially bounded weights.

1.4 Future work

The most immediate open problem left by our work is to get down from quasi-NC to NC for the perfect matching problem. Even for the bipartite case, this will require new insights or methods, as it is not clear how we could e.g. reduce the number of weight functions from $\log n$ to only a constant.

Proving that the search version of the perfect matching problem in planar graphs is in NC is also open[2]. While the $\mathsf{quasi\text{-}NC}$ result of [FGT16] gives rise to a new NC algorithm for bipartite planar graphs, which proceeds by verifying at each step whether the chosen weight function has removed the wanted cycles (it computes the girth of the support of the current face in NC), our λ-goodness progress measure seems to be difficult to verify in NC.

A related problem which has resisted derandomization so far is exact matching [PY82]. Here we are given a graph whose some edges are colored red and an integer k; the question is to find a perfect matching containing exactly k red edges. The problem is in RNC [MVV87], but not known to even be in P.

Finally, our polyhedral approach motivates the question of what other zero-one polytopes admit such a derandomization of the Isolation Lemma. One class that comes to mind are polyhedra with totally unimodular constraint matrices.

1.5 Outline

The rest of the paper is organized as follows. In Section 2 we introduce notation and define basic notions related to the perfect matching polytope and to the weight functions that we use. In Section 3 we define alternating circuits (our generalization of alternating cycles), discuss what it means for such a circuit to respect a face, and develop our tools for circuit removal. In Section 4 we introduce our measure of progress (λ-goodness), contractible sets and the contraction multigraph. We also state Theorem 4.11, which implies our main result. Finally, in Section 5 we prove our key technical theorem: that applying $\log_2 n+1$ weight functions allows us to make progress from λ-good to 2λ-good.

2 Preliminaries

Throughout the paper we consider a fixed graph $G=(V,E)$ with n vertices. (The isolating weight functions whose existence we prove can be generated without knowledge of the graph.) For notational convenience, we assume that $\log_2(n)$ evaluates to an integer; otherwise simply replace $\log_2(n)$ by $\lceil\log_2(n)\rceil$.

We use the following notation. For a subset $S\subseteq V$ of the vertices, let $\delta(S)=\{e\in E: |e\cap S|=1\}$ denote the edges crossing the cut $(S,V\setminus S)$ and $E(S)=\{e\in E:|e\cap S|=2\}$ denote the edges inside S. We shorten $\delta(\{v\})$ to $\delta(v)$ for $v\in V$. For a vector $(x_e)_{e\in E}\in\mathbb{R}^{|E|}$, we define $x(\delta(S))=\sum_{e\in\delta(S)}x_e$, as well as $\mathrm{supp}(x)=\{e\in E: x_e>0\}$. For a subset $F\subseteq E$ we define $\mathbb{1}_F$ to be the vector with 1 on coordinates in F and 0 elsewhere. We again shorten $\mathbb{1}_{\{e\}}$ to $\mathbb{1}_e$ for $e\in E$. Sometimes we identify matchings M with their indicator vectors $\mathbb{1}_M$.

A matching is a set of edges $M\subseteq E$ such that no two edges share an endpoint. A matching M is perfect if $|M|=\frac{n}{2}$.

2.1 Parallel complexity

The complexity class $\mathsf{quasi\text{-}NC}$ is defined as $\mathsf{quasi\text{-}NC}=\bigcup_{k\geq 0}\mathsf{quasi\text{-}NC}^k$, where $\mathsf{quasi\text{-}NC}^k$ is the class of problems having uniform circuits of quasi-polynomial size $2^{\log^{O(1)} n}$ and polylogarithmic depth $O(\log^k n)$ [Bar92]. Here by "uniform" we mean that the circuit can be generated in quasi-polynomial time.

By the results of [MVV87], Theorem 1.2 implies that the perfect matching problem (both the decision and the search variant) in general graphs is in $\mathsf{quasi\text{-}NC}$. The same can be said about maximum cardinality matching, as well as minimum-cost perfect matching for small costs (given in unary); see Section 5 of [MVV87].

Some care is required to obtain our postulated running time, i.e., that the perfect matching problem has uniform circuits of size $n^{O(\log^2 n)}$ and depth $O(\log^3 n)$. We could get a $\mathsf{quasi\text{-}NC}^4$ algorithm by applying the results of [MV97, Section 6.1] to compute the determinant(s). To shave off one $\log n$ factor, we use the following Chinese remaindering method, pointed out to us by Rohit Gurjar (it will also appear in the full version of [FGT16]). We first compute determinants modulo small primes; since the determinant has $2^{O(\log^3 n)}$ bits, we need as many primes (each of $O(\log^3 n)$ bits). For one prime this can be done in NC^2 [Ber84]. Then we reconstruct the true value from the remainders. Doing this for an n-bits result would be in NC^1 [BCH86], and thus for a result with $2^{O(\log^3 n)}$ bits it is in $\mathsf{quasi\text{-}NC}^3$.

2.2 Perfect matching polytope

Edmonds [Edm65a] showed that the following set of equalities and inequalities on the variables $(x_e)_{e\in E}$ determines the perfect matching polytope (i.e., the convex hull of indicator vectors of all perfect matchings):

$$\begin{aligned} x(\delta(v)) &= 1 && \text{for } v \in V, \\ x(\delta(S)) &\geq 1 && \text{for } S \subseteq V \text{ with } |S| \text{ odd}, \\ x_e &\geq 0 && \text{for } e \in E. \end{aligned}$$

Note that the constraints imply that $x_e \leq 1$ for any $e \in E$. We refer to the perfect matching polytope of the graph $G = (V, E)$ by $\mathrm{PM}(V, E)$ or simply by PM. Our approach exploits the special structure of faces of the perfect matching polytope. Recall that a face of a polytope is obtained by setting a subset of the inequalities to equalities. We follow the definition of a face from the book of Schrijver [Sch03] – in particular, every face is nonempty.

Throughout the paper, we will only consider the perfect matching polytope and so the term "face" will always refer to a face of PM. When talking about faces, we use the following notation:

Definition 2.1. *For a face F we define*

$$E(F) = \{e \in E : (\exists x \in F)\ x_e > 0\} \quad \text{and} \quad \mathcal{S}(F) = \{S \subseteq V : |S| \text{ odd and } (\forall x \in F)\ x(\delta(S)) = 1\}\,.$$

In other words, $E(F)$ contains the edges that appear in a perfect matching in F and $\mathcal{S}(F)$ contains the tight cut constraints of F.

Notice that if a set is tight for a face, then it is also tight for any of its subfaces.

Standard uncrossing techniques imply that faces can be defined using laminar families of tight constraints. This is proved using Lemma 2.2 below, which is also useful in our approach.

Two subsets $S, T \subseteq V$ of vertices are said to be crossing if they intersect and none is contained in the other, i.e., $S \cap T, S \setminus T, T \setminus S \neq \emptyset$. A family $\mathcal{L}$ of subsets of vertices is *laminar* if no two sets $S, T \in \mathcal{L}$ are crossing. Furthermore, we say that $\mathcal{L}$ is a *maximal laminar subset* of a family $\mathcal{S}$ if no set in $\mathcal{S} \setminus \mathcal{L}$ can be added to $\mathcal{L}$ while maintaining laminarity.

Note that any single-vertex set is tight for any face, and therefore a maximal laminar family contains all these sets; by convention, in our arguments all laminar families will always contain all singleton sets.

The following lemma is known; for completeness, its proof is included in Appendix A.

Lemma 2.2. *Consider a face F. For any maximal laminar subset $\mathcal{L}$ of $\mathcal{S}(F)$ we have*

$$\mathrm{span}(\mathcal{L}) = \mathrm{span}(\mathcal{S}(F))\,,$$

where for a subset $\mathcal{T} \subseteq \mathcal{S}(F)$, $\mathrm{span}(\mathcal{T})$ denotes the linear subspace of $\mathbb{R}^E$ spanned by the boundaries of sets in $\mathcal{T}$, i.e., $\mathrm{span}(\mathcal{T}) = \mathrm{span}\{\mathbb{1}_{\delta(S)} : S \in \mathcal{T}\}$.

Intuitively, Lemma 2.2 implies that a maximal laminar family $\mathcal{L}$ of $\mathcal{S}(F)$ is enough to describe a face F (together with the edge set $E(F)$). Furthermore, given a subface $F' \subseteq F$, we can extend $\mathcal{L}$ to a larger laminar family $\mathcal{L}' \supseteq \mathcal{L}$ which describes F'.

It is also well-known that the perfect matching polytope PM is integral, i.e., all of its extreme points are integral. It follows that every face of PM is also integral.

2.3 Weight functions

For our derandomization of the Isolation Lemma we will use families of weight functions which are possible to generate obliviously (i.e., using only the number of vertices in G). We define them below.

Definition 2.3. *Given $t \geq 7$, we define the family of weight functions $\mathcal{W}(t)$ as follows. Number the edge set $E = \{e_1, ..., e_{|E|}\}$ arbitrarily. Let $w_k : E \to \mathbb{Z}$ be given by $w_k(e_j) = (4n^2 + 1)^j \bmod k$ for $j = 1, ..., |E|$ and $k = 2, ..., t$. We define $\mathcal{W}(t) = \{w_k : k = 2, ..., t\}$.*

For brevity, we write $\mathcal{W} := \mathcal{W}(n^{20})$.

In our argument we will obtain a decreasing sequence of faces. Each face arises from the previous by minimizing over a linear objective (given by a weight function).

Definition 2.4. *Let F be a face and w a weight function. The subface of F minimizing w will be called $F[w]$:*

$$F[w] := \operatorname{argmin}\{\langle w, x\rangle : x \in F\}.$$

Instead of minimizing over one weight function and then over another, we can *concatenate* them in such a way that minimizing over the concatenation yields the same subface. In particular, we will argue that one just needs to try all possible concatenations of $O(\log^2 n)$ weight functions from $\mathcal{W}$ in order to find one which isolates a unique perfect matching in G (i.e., it produces a single extreme point as the minimizing subface).

Definition 2.5. *For two weight functions w and w', where $w : E \to \mathbb{Z}$ and $w' \in \mathcal{W}$, we define their* concatenation $w \circ w' := n^{21}w + w'$, *i.e.,*

$$\left(w \circ w'\right)(e) := n^{21} \cdot w(e) + w'(e).$$

We also define $\mathcal{W}^k$ to be the set of all concatenations of k weight functions from $\mathcal{W}$, i.e.,[9]

$$\mathcal{W}^k := \{w_1 \circ w_2 \circ ... \circ w_k : w_1, w_2, ..., w_k \in \mathcal{W}\}.$$

Fact 2.6. *We have $F[w][w'] = F[w \circ w']$.*

Proof. Both faces are integral; therefore we only need to show that $F[w][w'] \cap \mathbb{Z}^E = F[w \circ w'] \cap \mathbb{Z}^E$. The first set consists of matchings in F minimizing w and, among such matchings, minimizing w'. The second set consists of matchings in F minimizing $w \circ w'$. These two sets are equal because for any M:

- $(w \circ w')(M) = n^{21} \cdot w(M) + w'(M)$,
- $w' \in \mathcal{W} = \mathcal{W}(n^{20})$ implies that $0 \leq w'(M) < n^{20} \cdot \frac{n}{2} < n^{21}$ for any matching M,
- $w(M) \in \mathbb{Z}$, so that for any two matchings M_1 and M_2, $w(M_1) > w(M_2)$ implies $(w \circ w')(M_1) - (w \circ w')(M_2) = n^{21}\left(w(M_1) - w(M_2)\right) + \left(w'(M_1) - w'(M_2)\right) > 0$,

so that the ordering given by $w \circ w'$ is the same as the lexicographic ordering given by (w, w'). □

[9] By $w_1 \circ w_2 \circ ... \circ w_k$ we mean $((w_1 \circ w_2) \circ ...) \circ w_k$.

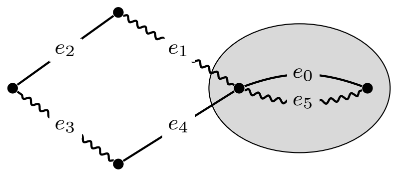

Figure 5: An example of an alternating circuit C of length 6 with indicator vector $(\pm\mathbb{1})_C = \sum_{i=0}^{5}(-1)^i \mathbb{1}_{e_i} = -\mathbb{1}_{e_1} + \mathbb{1}_{e_2} - \mathbb{1}_{e_3} + \mathbb{1}_{e_4}$ (since $\mathbb{1}_{e_0}$ and $\mathbb{1}_{e_5}$ cancel each other). Also note that $\langle (\pm\mathbb{1})_C, \mathbb{1}_{\delta(S)} \rangle = 0$ for the tight set S depicted in gray.

3 Alternating circuits and respecting a face

In this section we introduce two notions which are vital for our approach. Let us first motivate them.

Our argument is centered around *removing* even cycles. As discussed in Section 1.2 and Figure 2, the meaning of this term in the non-bipartite case needs to be more subtle than just "removing an edge of the cycle".

In order to deal with a cycle, we find a weight function which assigns it a nonzero circulation. Formally, given a cycle C with edges numbered in order, define a vector $(\pm\mathbb{1})_C \in \{-1, 0, 1\}^E$ as having 1 on even-numbered edges of C, -1 on odd-numbered edges of C, and 0 elsewhere; then we get that $\langle (\pm\mathbb{1})_C, w \rangle \neq 0$. Now, in the bipartite case, if such a cycle survived in the new face $F[w]$, that is, $C \subseteq E(F[w])$, then the vector $(\pm\mathbb{1})_C$ could be used to obtain a point in the face F with lower w-weight than the points in $F[w]$, a contradiction. This was possible because of the simple structure of the bipartite perfect matching polytope.

In the non-bipartite case, it is not enough that $C \subseteq E(F[w])$ in order to obtain such a point (and a contradiction); it is also required that, if the cycle C enters a tight odd-set S on an even-numbered edge, it exits it on an odd-numbered edge (and vice versa). This makes intuitive sense: if C were obtained from the symmetric difference of two perfect matchings which both have exactly one edge crossing S, then C would have this property. Formally, we require that $\langle (\pm\mathbb{1})_C, \mathbb{1}_{\delta(S)} \rangle = 0$ for each $S \in \mathcal{S}(F[w])$. If C meets these two conditions, which are exactly what is required to obtain a contradictory point as above (see the proof of Lemma 3.3), then we say that C *respects* the face $F[w]$.

To reiterate: if we assign a nonzero circulation to a cycle, then it will not respect the new face, and this is what is now meant by removing a cycle.

To deal with the second, more technical difficulty discussed in Section 1.2, we need to remove not only simple cycles of even length, but also walks with repeated edges. However, we would run into problems if we allowed all such walks (up to a given length). Consider for example a walk C of length 2; such a walk traverses an edge back and forth. It is impossible to assign a nonzero circulation to C, because its vector $(\pm\mathbb{1})_C$ is zero. However, for the same reason, such a walk C fails to induce a contradictory point, so its removal is unnecessary. Therefore we define alternating circuits to be those even walks whose vector $(\pm\mathbb{1})_C$ is nonzero (see Figure 5 for an example). To avoid further technical issues, we also formulate the definition of respect in terms of the vector $(\pm\mathbb{1})_C$.

Definition 3.1. *Let $C = (e_0, ..., e_{k-1})$ be a nonempty cyclic walk of even length k.*

- *We define the* alternating indicator vector $(\pm\mathbb{1})_C$ *of C to be $(\pm\mathbb{1})_C = \sum_{i=0}^{k-1}(-1)^i \mathbb{1}_{e_i}$ (where $\mathbb{1}_e \in \mathbb{R}^E$ is the indicator vector having 1 on position e and 0 elsewhere).*[10]

[10]Note that $(\pm\mathbb{1})_C$ does not need to have all entries -1, 0 or 1 since edges can repeat in C.

- *We say that C is an* alternating circuit *if its alternating indicator vector is nonzero. We also refer to C as an* alternating (simple) cycle *if it is an alternating circuit that visits every vertex at most once.*

- *When talking about a graph with node-weights, the node-weight of an alternating circuit is the sum of all node-weights of visited vertices (with multiplicities if visited multiple times).*

Definition 3.2. *We say that a vector $y \in \mathbb{Z}^E$* respects *a face F if:*

- $\operatorname{supp}(y) \subseteq E(F)$, *and*

- *for each $S \in \mathcal{S}(F)$ we have $\left\langle y, \mathbb{1}_{\delta(S)} \right\rangle = 0$.*

Furthermore, we say that an alternating circuit C respects a face F if its alternating indicator vector $(\pm\mathbb{1})_C$ respects F.

Clearly, if $F' \subseteq F$ are faces and a vector respects F', then it also respects F.

Now we argue that we can remove an alternating circuit by assigning it a nonzero circulation. The proof of this lemma (which generalizes Lemma 3.2 of [FGT16]) motivates Definition 3.2.

Lemma 3.3. *Let $y \in \mathbb{Z}^E$ be a vector and F a face. If $w : E \to \mathbb{R}$ is such that $\langle y, w \rangle \neq 0$, then y does not respect the face $F' = F[w]$.*

Proof. Suppose towards a contradiction that y respects F'. Assume that $\langle w, y \rangle < 0$ (otherwise use $-y$ in place of y). We pick $x \in F'$ to be the average of all extreme points of F', so that the constraints of PM which are tight for x are exactly those which are tight for F'. Select $\varepsilon > 0$ very small. Then $\langle x + \varepsilon y, w \rangle < \langle x, w \rangle$, which will contradict the definition of $F' = \operatorname{argmin}\{\langle w, x \rangle : x \in F\}$ once we show that $x + \varepsilon y \in F$. We show that $x + \varepsilon y \in F' \subseteq F$ by verifying:

- If $e \in E(F')$ is an edge with $x_e > 0$, then $(x + \varepsilon y)_e = x_e + \varepsilon y_e \geq 0$ if ε is chosen small enough.

- If $e \in E \setminus E(F')$ is an edge with $x_e = 0$, then from y respecting F' we get $e \notin \operatorname{supp}(y)$ and so $(x + \varepsilon y)_e = 0$.

- If $S \notin \mathcal{S}(F')$ is an odd set not tight for F', i.e., $\left\langle x, \mathbb{1}_{\delta(S)} \right\rangle > 1$, then $\left\langle x + \varepsilon y, \mathbb{1}_{\delta(S)} \right\rangle = \left\langle x, \mathbb{1}_{\delta(S)} \right\rangle + \varepsilon \left\langle y, \mathbb{1}_{\delta(S)} \right\rangle \geq 1$ if ε is chosen small enough.

- If $S \in \mathcal{S}(F')$ is an odd set tight for F' (this includes all singleton sets), then from y respecting F' we get $\left\langle y, \mathbb{1}_{\delta(S)} \right\rangle = 0$ and thus $\left\langle x + \varepsilon y, \mathbb{1}_{\delta(S)} \right\rangle = \left\langle x, \mathbb{1}_{\delta(S)} \right\rangle = 1$. □

The following lemma says that we can assign nonzero circulation to many vectors at once using an oblivious choice of weight function from $\mathcal{W}$. It is a minor generalization of Lemma 2.3 of [FGT16] and the proof remains similar; we give it for completeness.

Lemma 3.4. *For any number s and for any set of s vectors $y_1, ..., y_s \in \mathbb{Z}^E \setminus \{0\}$ with the boundedness property $\|y_i\|_1 \leq 4n^2$, there exists $w \in \mathcal{W}(n^3 s)$ with $\langle y_i, w \rangle \neq 0$ for each $i = 1, ..., s$.*

We usually invoke Lemma 3.4 with vectors y_i being the alternating indicator vectors of alternating circuits. Then the quantities $\langle y_i, w \rangle$ are the circulations of these circuits.

Proof. Let $w' : E \to \mathbb{Z}$ be given by $w'(e_j) = (4n^2+1)^j$ for $j = 1, ..., |E|$. Then we have $\langle y_i, w'\rangle \neq 0$ for each i because the highest nonzero coefficient dominates the expression.[11] Let $t = n^3 s$. We want to show that there exists $k = 2, ..., t$ such that for all $i = 1, ..., s$, $\langle y_i, w_k\rangle \neq 0$, that is, $|\langle y_i, w'\rangle| \neq 0 \mod k$ (w_k are as in Definition 2.3). This will be implied if there exists $k = 2, ..., t$ such that $\prod_i |\langle y_i, w'\rangle| \neq 0 \mod k$. So there should be some $k = 2, ..., t$ not dividing $\prod_i |\langle y_i, w'\rangle|$ – equivalently, $\mathrm{lcm}(2, ..., t)$ should not divide $\prod_i |\langle y_i, w'\rangle|$. Knowing that $\prod_i |\langle y_i, w'\rangle| \neq 0$, this will be implied if we have $\prod_i |\langle y_i, w'\rangle| < \mathrm{lcm}(2, ..., t)$. This is true because

$$\prod_{i=1}^{s} |\langle y_i, w'\rangle| < \left((4n^2+1)^{|E|+1}\right)^s < (4n^2+1)^{n^2 s} = 2^{n^2 s \log(4n^2+1)} < 2^{n^3 s} = 2^t < \mathrm{lcm}(2, ..., t)$$

where we used that $\mathrm{lcm}(2, ..., t) > 2^t$ for $t \geq 7$ [Nai82]. □

Lemmas 3.3 and 3.4 together imply the following:

Corollary 3.5. *Let F be a face. For any finite set of vectors $\mathcal{Y} \subseteq \mathbb{Z}^E \setminus \{0\}$ with the boundedness property $\|y\|_1 \leq 4n^2$ for $y \in \mathcal{Y}$, there exists $w \in \mathcal{W}(n^3|\mathcal{Y}|)$ such that each $y \in \mathcal{Y}$ does not respect the face $F' = F[w]$.* □

4 Contractible sets and λ-goodness

We will make progress by ensuring that larger and larger parts of the graph are "isolated" in our current face F. By "parts of the graph" we mean sets S which are tight for F. As discussed in Section 1.3, for such a set S, the following isolation property is desirable: once the (only) edge of a matching which lies on the boundary of S is fixed, the entire matching inside S is uniquely determined. This motivates the following definition:

Definition 4.1. *Let F be a face and let $S \in \mathcal{S}(F)$ be a tight set for F. We say that S is* F-contractible *if for every $e \in \delta(S)$ there are no two perfect matchings in F which both contain e and are different inside S.*

Intuitively, a contractible set can be thought of as a single vertex with respect to the structure of the current face of the perfect matching polytope.

The notion of contractibility enjoys the following two natural monotonicity properties:

Fact 4.2. *Let $F' \subseteq F$ be two faces. If S is F-contractible, then it is also F'-contractible.* □

Lemma 4.3. *Let F be a face and $S \subseteq T$ two sets tight for F (i.e., $S, T \in \mathcal{S}(F)$). If T is F-contractible, then so is S.*

Proof. Let $e \in \delta(S)$. Suppose that M_1 and M_2 are two perfect matchings in F which contain e but are different inside S. We will argue that in that case there also exist two perfect matchings M_1 and M_{12} in F which contain e, are different inside S, and are equal outside of S.

Once we have that, we conclude as follows. Let f be the (only) edge in $\delta(T) \cap M_1$ (perhaps $f = e$); then also $f \in M_{12}$. Then M_1 and M_{12} are two perfect matchings in F which contain $f \in \delta(T)$ but are different inside T, contradicting that T is F-contractible.

[11] Let j' be maximum with $y_i(e_{j'}) \neq 0$; suppose $y_i(e_{j'}) > 0$. Then because $\|y_i\|_\infty \leq \|y_i\|_1 \leq 4n^2$, we have

$$\langle y_i, w'\rangle = y_i(e_{j'})(4n^2+1)^{j'} + \sum_{j<j'} y_i(e_j)(4n^2+1)^j > (4n^2+1)^{j'} + \sum_{j=-\infty}^{j'-1} (-4n^2)(4n^2+1)^j = 0.$$

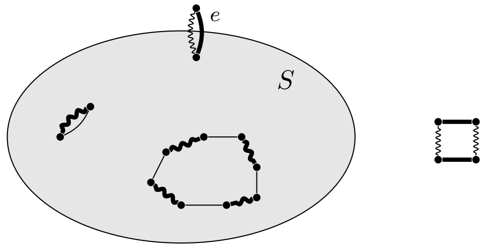

Figure 6: Illustration of the matching M_{12} constructed in the proof of Lemma 4.3. Straight and swirly edges denote M_1 and M_2 respectively. The thick edges denote M_{12}, which agrees with M_1 outside S and with M_2 inside S.

To get the outstanding claim, we define

$$M_{12} = (M_1 \setminus E(S)) \cup (M_2 \cap E(S))$$

to be the perfect matching that agrees with M_1 on all edges not in $E(S)$ and agrees with M_2 on all edges in $E(S)$ (see Figure 6). To see that M_{12} is a perfect matching, notice that both M_1 and M_2 are in F and contain e. Furthermore, as $e \in \delta(S)$ for the tight set $S \in \mathcal{S}(F)$, we have that $M_1 \cap E(S)$ and $M_2 \cap E(S)$ are both perfect matchings on the vertex set S where we ignore the vertex incident to e. We can thus "replace" $M_1 \cap E(S)$ by $M_2 \cap E(S)$ to obtain the perfect matching M_{12}.

We now show that M_{12} is in the face F. Suppose the contrary. Since M_1 and M_2 are both in F, we have $M_{12} \subseteq E(F)$. Therefore, if M_{12} is not in F, we must have $|\delta(R) \cap M_{12}| > 1$ for some tight set $R \in \mathcal{S}(F)$. Since $|R|$ is odd, also $|\delta(R) \cap M|$ is odd for any perfect matching M. In particular, $|\delta(R) \cap M_{12}| \geq 3$, which contradicts

$$|\delta(R) \cap M_{12}| \leq |\delta(R) \cap M_1| + |\delta(R) \cap M_2| = 2\,,$$

where the equality holds because M_1 and M_2 are perfect matchings in F and $R \in \mathcal{S}(F)$ is a tight set. □

In our proof, we will be working with faces and laminar families which are compatible in the following sense:

Definition 4.4. *Let F be a face and $\mathcal{L}$ a laminar family. If $\mathcal{L} \subseteq \mathcal{S}(F)$, i.e., all sets $S \in \mathcal{L}$ are tight for F, then we say that $(F, \mathcal{L})$ is a* face-laminar pair.

Given a face-laminar pair $(F, \mathcal{L})$, we will often work with a multigraph obtained from G by contracting all small sets, i.e., those with size being at most some parameter λ (which is a measure of our progress). This multigraph will be called the *contraction* (see Figure 7 for an example).

In the contraction, we will also remove all boundaries of larger sets (i.e., those with size larger than λ). This is done to simulate working inside each such large set independently, because the contraction then decomposes into a collection of disconnected components, one per each large set. Because, in the contraction, each set in $\mathcal{L}$ has either been contracted or has had its boundary removed, our task is reduced to dealing with instances having no laminar sets.

Moreover, we only include those edges which are still in the support of the current face F, i.e., the set $E(F)$.

Definition 4.5. *Given a face-laminar pair $(F, \mathcal{L})$ and a parameter λ (with $1 \leq \lambda \leq 2n$), we define the $(F, \mathcal{L}, \lambda)$*-contraction *of G as a node-weighted multigraph as follows:*

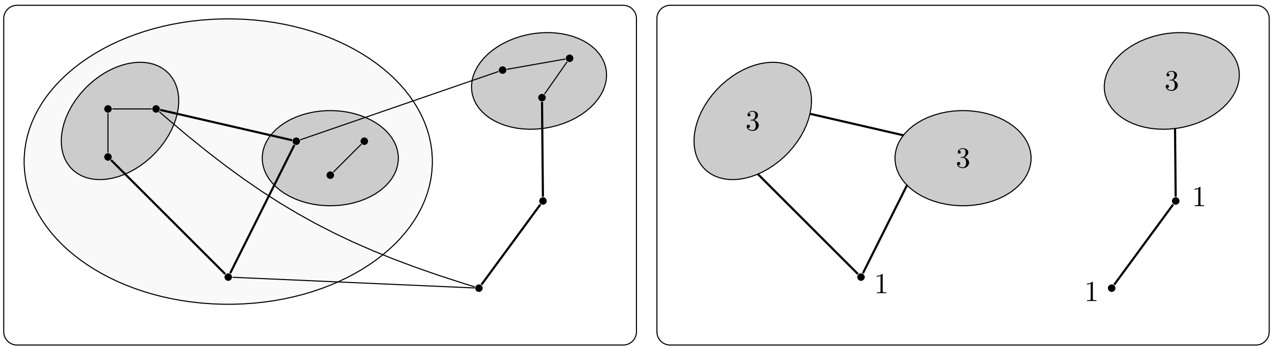

(a) A graph G and a laminar family $\mathcal{L}$. We only draw the edges in $E(F)$. We also do not draw ellipses for the singleton sets in $\mathcal{L}$. The dark-gray sets are F-contractible.

(b) The $(F, \mathcal{L}, 4)$-contraction of G. Its vertices are labeled by their node-weight.

Figure 7: An example of the $(F, \mathcal{L}, \lambda)$-contraction of G.

- *the node set is the set of maximal sets of size (cardinality) at most λ in $\mathcal{L}$,*
- *each node has a node-weight equal to the size of the corresponding set,*
- *the edge set is obtained from $E(F) \setminus \bigcup_{T \in \mathcal{L}: |T| > \lambda} \delta(T)$ by contracting each of these maximal sets.*[12]

In the $(F, \mathcal{L}, \lambda)$-contractions arising in our arguments, we will always only contract sets $S \in \mathcal{L}$ which are F-contractible (i.e., the vertices of a contraction will always correspond to F-contractible sets). Then, a very useful property is that alternating circuits in the contraction can be lifted to alternating circuits in the entire graph G in a canonical way. (This is done in the proofs of Lemmas 5.3 and 5.8.)

Finally, we need the following extension of Definition 3.2 for vectors defined on the contraction:

Definition 4.6. *Denote the $(F, \mathcal{L}, \lambda)$-contraction of G as H, and let $z \in \mathbb{Z}^{E(H)}$ be a vector on the edges of H. We say that z respects a subface $F' \subseteq F$ if*

- $\operatorname{supp}(z) \subseteq E(F')$, *and*
- *for each $S \in \mathcal{S}(F')$ which is a union of sets in $V(H)$,*[13] *we have $\langle z, \mathbb{1}_{\delta(S)} \rangle = 0$.*

As before, we say that an alternating circuit C in H respects a subface F' if its alternating indicator vector $(\pm\mathbb{1})_C \in \mathbb{Z}^{E(H)}$ respects F'.

Now we are able to define our measure of progress. On one hand, we want to make larger and larger laminar sets contractible. On the other hand, there could very well be no laminar sets, so we also proceed as in the bipartite case: remove longer and longer alternating circuits.

Definition 4.7. *Let $(F, \mathcal{L})$ be a face-laminar pair and λ a parameter (with $1 \leq \lambda \leq 2n$). We say that this pair is* λ-good *if:*

(i) each $S \in \mathcal{L}$ with $|S| \leq \lambda$ is F-contractible,

[12] That is, an edge of G maps to an edge of the contraction if it is in $E(F)$, it is not inside any of the contracted sets and it does not cross any cut defined by a set $T \in \mathcal{L} : |T| > \lambda$. Sometimes we identify edges of the contraction with their preimages in G.

[13] That is, the maximal sets of size at most λ in $\mathcal{L}$.

(ii) in the $(F, \mathcal{L}, \lambda)$-contraction of G, there is no alternating circuit of node-weight at most λ.

We begin with $\lambda = 1$, which is trivial, and then show that by concatenating enough weight functions we can obtain face-laminar families which are 2-good, 4-good, 8-good, and so on. We are done once we have a λ-good family with $\lambda \geq n$. The components of this proof strategy are given in the following three claims. The first step is clear:

Fact 4.8. *The face-laminar pair* $(\mathrm{PM}, \{\{v\} : v \in V\})$ *is* 1*-good.* ☐

We can then proceed iteratively in $\log_2 n$ rounds using the following theorem. Its proof is given in Section 5; it constitutes the bulk of our argument.

Theorem 4.9. *Let $(F, \mathcal{L})$ be a λ-good face-laminar pair. Then there exists a weight function $w \in \mathcal{W}^{\log_2 n+1}$ and a laminar family $\mathcal{L}' \supseteq \mathcal{L}$ such that $(F[w], \mathcal{L}')$ is a 2λ-good face-laminar pair.*

We are done once λ exceeds n:

Lemma 4.10. *Suppose $(F, \mathcal{L})$ is λ-good for some $\lambda \geq n$. Then $|F| = 1$.*

We think that the proof of this lemma is instructive. It serves to understand and motivate Definition 4.7, and more involved versions of this argument appear in the sequel.

Proof. Let H be the $(F, \mathcal{L}, \lambda)$-contraction of G. Also let $S_1, S_2, ..., S_k$ be all maximal sets in $\mathcal{L}$. Their disjoint union is V and we have $V(H) = \{S_1, ..., S_k\}$ and $E(H) = \bigcup_{i=1}^{k} \delta(S_i)$.[14] Since $(F, \mathcal{L})$ is λ-good, each set $S \in \mathcal{L}$ in F-contractible, and H contains no alternating circuit of node-weight at most λ – in particular, H contains no alternating simple cycle[15].

Now we show that there is only one perfect matching in F. One direction is easy: since F is a face, it is nonempty by definition. For the other direction, let M_1 and M_2 be two perfect matchings in F; we show that $M_1 = M_2$.

Because the sets $S_1, ..., S_k$ are tight for F, any perfect matching in G induces a perfect matching in H. If the matchings induced by M_1 and M_2 are different, then their symmetric difference contains an alternating simple cycle in H, which is impossible. So the induced matchings must be equal, i.e., $M_1 \cap \bigcup_i \delta(S_i) = M_2 \cap \bigcup_i \delta(S_i)$. Moreover, the sets $S_1, ..., S_k$ are F-contractible, which means that, given the boundary edges, there is a unique perfect matching in F inside each S_i. This yields $M_1 = M_2$. ☐

Before we proceed to the proof of Theorem 4.9, let us see how Fact 4.8, Theorem 4.9, and Lemma 4.10 together give our desired claim:

Theorem 4.11. *There exists an isolating weight function $w \in \mathcal{W}^{(\log_2 n+1)\log_2 n}$, i.e., one with $|\mathrm{PM}[w]| = 1$.*

Proof. We iteratively construct a sequence of face-laminar pairs $(F_i, \mathcal{L}_i)$ for $i = 0, 1, ..., \log_2 n$ such that $(F_i, \mathcal{L}_i)$ is 2^i-good and $F_i = F_{i-1}[w_i]$ for some weight function $w_i \in \mathcal{W}^{\log_2 n+1}$. We begin by setting $F_0 = \mathrm{PM}$ and $\mathcal{L}_0 = \{\{v\} : v \in V\}$; by Fact 4.8, $(F_0, \mathcal{L}_0)$ is 1-good. Then for $i = 1, ..., \log_2 n$ we use Theorem 4.9 to obtain the wanted weight function w_i along with a laminar family $\mathcal{L}_i \supseteq \mathcal{L}_{i-1}$. Finally, we have that $(F_{\log_2 n}, \mathcal{L}_{\log_2 n})$ is $2^{\log_2 n}$-good, so that by Lemma 4.10, $|F_{\log_2 n}| = 1$.

It remains to argue that $F_{\log_2 n} = \mathrm{PM}[w]$ for some $w \in \mathcal{W}^{(\log_2 n+1)\log_2 n}$. To do this, we proceed as in Section 2.3: define the concatenation $w' \bullet w'' := n^{21(\log_2 n+1)} w' + w''$ for two

[14] Recall that $\mathcal{L}$ contains all singletons, so that every vertex is covered by a set in $\mathcal{L}$.

[15] This is since an upper bound on the node-weight of any alternating simple cycle is $|S_1| + ... + |S_k| = n \leq \lambda$.

weight functions w' and w'', where $w'' \in \mathcal{W}^{\log_2 n+1}$.[16] Using a version of Fact 2.6 we get that $F_{\log_2 n} = \mathrm{PM}[w_1][w_2]...[w_{\log_2 n}] = \mathrm{PM}[w_1 \bullet w_2 \bullet ... \bullet w_{\log_2 n}]$. We put $w = w_1 \bullet w_2 \bullet ... \bullet w_{\log_2 n}$.$\square$

Theorem 4.11 implies Theorem 1.2 because we have $|\mathcal{W}^{(\log_2 n+1)\log_2 n}| = |\mathcal{W}|^{(\log_2 n+1)\log_2 n} \leq n^{20(\log_2 n+1)\log_2 n}$, the values of any $w \in \mathcal{W}^{(\log_2 n+1)\log_2 n}$ are bounded by $n^{21(\log_2 n+1)\log_2 n}$, and the functions $w \in \mathcal{W}$ can be generated obliviously (using only the number of vertices n).

5 Proof of the key Theorem 4.9: from λ-good to 2λ-good

In this section we show how to make progress (measured by the λ parameter of λ-goodness) by applying a new weight function to the current face. Our objective is to make larger sets contractible (by doubling the size threshold from λ to 2λ) and to ensure that in the new contracted graph, alternating circuits of an increased node-weight are not present. We do this by moving from the current face-laminar pair, which we call $(F_{\mathsf{in}}, \mathcal{L}_{\mathsf{in}})$, to a new face-laminar pair $(F_{\mathsf{out}}, \mathcal{L}_{\mathsf{out}})$. Both pairs have the property that the laminar family is a maximal laminar family of sets tight for the face. The new family extends the previous, i.e., $\mathcal{L}_{\mathsf{out}} \supseteq \mathcal{L}_{\mathsf{in}}$.

Our main technical tools are Theorem 5.1 and Lemma 5.6. Theorem 5.1 is used to ensure that certain alternating circuits are not present in the new contraction. It says that, if our current contraction has no alternating circuits of at most some node-weight, then a single weight function $w \in \mathcal{W}$ is enough to guarantee that all alternating circuits of at most twice that node-weight do not respect the new face obtained by applying w. We call this *removing* these circuits. Lemma 5.6 is used so as to make sure that sets in our laminar family which are of the appropriate size (regulated by λ) become contractible. Later, new sets will be added to the laminar family (in Lemma 5.8) in such a way that these properties are maintained and that the removed alternating circuits indeed do not survive in the new contraction.

The formal structure of the proof is as follows. We begin from a λ-good face-laminar pair $(F_{\mathsf{in}}, \mathcal{L}_{\mathsf{in}})$. First, in Theorem 5.5, we show the existence of a weight function $w_{\mathsf{out}} \in \mathcal{W}^{\log_2(n)+1}$ such that the face $F_{\mathsf{out}} = F_{\mathsf{in}}[w_{\mathsf{out}}]$ satisfies two conditions which make progress on conditions (i) and (ii) of λ-goodness:

(i)' For each $S \in \mathcal{L}_{\mathsf{in}}$ with $|S| \leq 2\lambda$, S is F_{out}-contractible.

(ii)' In the $(F_{\mathsf{in}}, \mathcal{L}_{\mathsf{in}}, 2\lambda)$-contraction of G, there is no F_{out}-respecting alternating circuit of node-weight at most 2λ.

The proof of Theorem 5.5 involves Theorem 5.1 and Lemma 5.6. It gives us the wanted face F_{out} and weight function w_{out}. Then, in Lemma 5.8, we show that extending the laminar family $\mathcal{L}_{\mathsf{in}}$ to a maximal laminar family $\mathcal{L}_{\mathsf{out}}$ (of the new tight sets) yields a 2λ-good pair $(F_{\mathsf{out}}, \mathcal{L}_{\mathsf{out}})$. This finishes the proof of Theorem 4.9.

5.1 Removing alternating circuits

This section is devoted to the proof of Theorem 5.1, which is a technical tool we use to remove alternating circuits of size between λ and 2λ from the contraction.

Theorem 5.1. *Consider a face-laminar pair $(F, \mathcal{L})$ such that each $S \in \mathcal{L}$ with $|S| \leq \beta$ is F-contractible (for a parameter β). Denote by H the $(F, \mathcal{L}, \beta)$-contraction of G. If H has no alternating circuit of node-weight at most λ, then there exists $w \in \mathcal{W}$ such that H has no $F[w]$-respecting alternating circuit of node-weight at most 2λ.*

[16] We need to use a padding term $n^{21(\log_2 n+1)}$ which is larger than the n^{21} of Definition 2.5 because the right-hand weight fuctions are now from $\mathcal{W}^{\log_2 n+1}$ rather than from $\mathcal{W}$.

We begin with a simple technical fact: to verify that a vector respects a face, it is enough to check this for a maximal laminar family of tight constraints.

Lemma 5.2. *Consider a face F and let $\mathcal{L}$ be a maximal laminar subset of $\mathcal{S}(F)$. Then a vector y respects F if $\operatorname{supp}(y) \subseteq E(F)$ and for each $S \in \mathcal{L}$ we have $\langle y, \mathbb{1}_{\delta(S)} \rangle = 0$.*

Proof. We need to prove that $\langle y, \mathbb{1}_{\delta(S)} \rangle = 0$ for all $S \in \mathcal{S}(F)$ assuming that this holds for all $S \in \mathcal{L}$. As $\mathcal{L}$ is a maximal laminar subset of $\mathcal{S}(F)$, Lemma 2.2 says that $\operatorname{span}(\mathcal{L}) = \operatorname{span}(\mathcal{S}(F))$. In other words, for any $S \in \mathcal{S}(F)$, we can write $\mathbb{1}_{\delta(S)}$ as a linear combination $\sum_{L \in \mathcal{L}} \lambda_L \mathbb{1}_{\delta(L)}$ for some coefficients $(\lambda_L)_{L \in \mathcal{L}}$. Hence

$$\langle y, \mathbb{1}_{\delta(S)} \rangle = \left\langle y, \sum_{L \in \mathcal{L}} \lambda_L \mathbb{1}_{\delta(L)} \right\rangle = \sum_{L \in \mathcal{L}} \lambda_L \langle y, \mathbb{1}_{\delta(L)} \rangle = 0. \qquad \square$$

Now we state a lemma which reduces the task of removing an alternating circuit in H to that of removing a vector defined on the edges of G, which we can do using Corollary 3.5. Throughout this section, F, $\mathcal{L}$ and H are as in the statement of Theorem 5.1.

Lemma 5.3. *Let $z \in \mathbb{Z}^{E(H)} \setminus \{0\}$ be a vector on the edges of H satisfying $\langle z, \mathbb{1}_{\delta(S)} \rangle = 0$ for each $S \in V(H)$.*[17] *Then there exists a vector $y \in \mathbb{Z}^E \setminus \{0\}$ such that for any face $F' \subseteq F$ we have: if z respects F', then y respects F'. We also have $\|y\|_1 \le n\|z\|_1$.*

We remark that y does not depend on F'.

Proof. We consider z as a vector $z \in \mathbb{Z}^E$ (by inserting zeroes at the coordinates in $E \setminus E(H)$). We may assume $\operatorname{supp}(z) \subseteq E(F)$; otherwise z cannot respect F' and thus we are done by outputting any y.

The proof idea is to extend z to a vector $y \in \mathbb{Z}^E$ which resembles an alternating indicator vector. We do this in a canonical way, so that if this extension does not respect F', then it must be because z itself does not respect F'.

To this end, we do the following for each $S \in V(H)$: pair up the boundary edges $e \in \delta(S)$ which have $z_e > 0$ with boundary edges e which have $z_e < 0$, respecting their multiplicities as given by z.[18] This is possible because $\langle z, \mathbb{1}_{\delta(S)} \rangle = 0$. Let $\{(e_i^+, e_i^-)\}_i$ be the set of pairs of edges obtained in this way across all $S \in V(H)$, and let $S_i \in V(H)$ be the set for which the pair (e_i^+, e_i^-) has been introduced. Also denote by v_i^+, v_i^- the S_i-endpoints of edges e_i^+, e_i^-.

Now, for each i we have $e_i^+, e_i^- \in \operatorname{supp}(z) \subseteq E(F)$, and S is F-contractible, so there is a unique perfect matching M_i^+ on the vertex-induced subgraph $(S_i \setminus \{v_i^+\}, E(S_i \setminus \{v_i^+\}))$ in F, as well as a unique perfect matching M_i^- on $(S_i \setminus \{v_i^-\}, E(S_i \setminus \{v_i^-\}))$ in F. We let

$$y := z + \sum_i \left(\mathbb{1}_{M_i^+} - \mathbb{1}_{M_i^-} \right).$$

Claim 1. *Let $F' \subseteq F$ be such that z respects F'. Then y respects F'.*

Proof. Since z respects F', we have $\operatorname{supp}(z) \subseteq E(F')$. This implies that for each i, $M_i^+ \subseteq E(F')$. Indeed, since $e_i^+ \in \operatorname{supp}(z) \subseteq E(F')$, there is a perfect matching on $(S_i \setminus \{v_i^+\}, E(S_i \setminus \{v_i^+\}))$ in F'. However, S_i is F-contractible and thus M_i^+ is the *only* such matching in F (thus also in F'). Therefore $M_i^+ \subseteq E(F')$ and analogously $M_i^- \subseteq E(F')$.

Now we check the conditions for y to respect F':

[17] Recall that vertices of H are elements of $\mathcal{L}$, i.e., sets of vertices.

[18] For example, if we had $\delta(S) = \{e_1, e_2, e_3\}$ with $z(e_1) = 3$, $z(e_2) = -2$ and $z(e_3) = -1$, we would get pairs $(e_1, e_2), (e_1, e_2), (e_1, e_3)$.

- We have $\mathrm{supp}(y) = \mathrm{supp}(z) \cup \bigcup_i \left(M_i^+ \cup M_i^-\right) \subseteq E(F')$.
- Let $T \in \mathcal{S}(F')$. We need to verify that $\left\langle y, \mathbb{1}_{\delta(T)} \right\rangle = 0$. Let $\mathcal{L}'$ be a maximal laminar subfamily of $\mathcal{S}(F')$ extending $\mathcal{L}$, i.e., $\mathcal{L} \subseteq \mathcal{L}' \subseteq \mathcal{S}(F')$. By Lemma 5.2, it is enough to verify that $\left\langle y, \mathbb{1}_{\delta(T)} \right\rangle = 0$ for $T \in \mathcal{L}'$. For a set T belonging to a laminar family which extends $\mathcal{L}$, it is not hard to see that there are two possibilities: either $T \subsetneq S$ for some $S \in V(H)$, or T is a union of sets in $V(H)$. In the latter case, $\left\langle y, \mathbb{1}_{\delta(T)} \right\rangle = \left\langle z, \mathbb{1}_{\delta(T)} \right\rangle = 0$ because y equals z on edges crossing sets in $V(H)$ and because z respects F'. In the former case, we have

$$\begin{aligned}\left\langle y, \mathbb{1}_{\delta(T)} \right\rangle &= \left\langle \sum_{i:S_i=S} \mathbb{1}_{e_i^+} - \mathbb{1}_{e_i^-} + \mathbb{1}_{M_i^+} - \mathbb{1}_{M_i^-}, \mathbb{1}_{\delta(T)} \right\rangle \\ &= \sum_{i:S_i=S} \left\langle \mathbb{1}_{e_i^+} - \mathbb{1}_{e_i^-} + \mathbb{1}_{M_i^+} - \mathbb{1}_{M_i^-}, \mathbb{1}_{\delta(T)} \right\rangle\end{aligned}$$

  (because these are the only edges of y which can possibly cross T). Now it is enough to show that each summand is 0.

  For this, we know that $M_i^+ \cup \{e_i^+\}$ and $M_i^- \cup \{e_i^-\}$ are (partial) matchings in F' and that T is tight for F'. Therefore we have $|\delta(T) \cap \left(M_i^+ \cup \{e_i^+\}\right)| = 1$,[19] and the same holds for $M_i^- \cup \{e_i^-\}$. Therefore $\left\langle \mathbb{1}_{M_i^+ \cup \{e_i^+\}}, \mathbb{1}_{\delta(T)} \right\rangle = \left\langle \mathbb{1}_{M_i^- \cup \{e_i^-\}}, \mathbb{1}_{\delta(T)} \right\rangle$. □

Regarding the norm, every edge (with multiplicity) in z causes at most $n/2$ new edges to appear in y. Therefore $\|y\|_1 \le (n/2 + 1)\|z\|_1 \le n\|z\|_1$. □

Our second lemma gives a bound on the number of alternating circuits we need to remove. Its proof resembles that of Lemma 3.4 in [FGT16], but it is slightly more complex, as we are dealing with a node-weighted multigraph, as well as with alternating circuits instead of simple cycles (see Section 3). (We have made no attempt to minimize the exponent 17.)

Lemma 5.4. *There are polynomially many alternating circuits of node-weight at most 2λ in H, up to identifying circuits with equal alternating indicator vectors. More precisely, the cardinality of the set*

$$\{(\pm\mathbb{1})_C : C \text{ is an alternating circuit in } H \text{ of node-weight at most } 2\lambda\}$$

is at most n^{17}.

Proof. We will associate a small *signature* with each alternating circuit in H of node-weight at most 2λ, with the property that alternating circuits with different alternating indicator vectors are assigned different signatures. This will prove that the considered cardinality is at most the number of possible signatures, which will be polynomially bounded.

Let $C = (e_0, e_1, ..., e_{k-1})$ be an alternating circuit in H of node-weight at most 2λ; we want to define its signature $\sigma(C)$. To streamline notation, we let v_i be the tail of e_i for $i = 0, \ldots, k-1$ (where we direct e_i according to the walk C). Thus C is of the form

$$v_0 \xrightarrow{e_0} v_1 \xrightarrow{e_1} \ldots \xrightarrow{e_{k-2}} v_{k-1} \xrightarrow{e_{k-1}} v_0\,.$$

[19] Formally, consider a perfect matching M^+ (on G) in F' which is a superset of $M_i^+ \cup \{e_i^+\}$. Then we have $|\delta(T) \cap M^+| = 1$. But $\delta(T) \cap \left(M_i^+ \cup \{e_i^+\}\right) = \delta(T) \cap M^+$ because $T \subseteq S$.

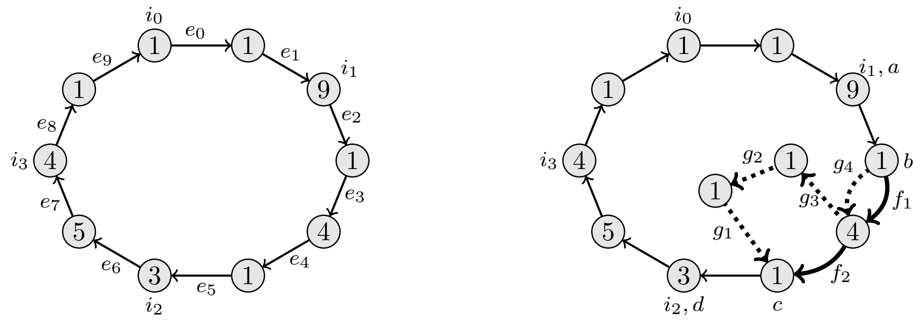

Figure 8: Intuition of the signature vector definition and the proof of Lemma 5.4.
On the left, each vertex is labeled by its node-weight, and the corresponding selection of i_0, i_1, i_2, i_3 is shown for $\lambda = 16$. Notice that the selected vertices partition the alternating circuit into paths; the total node-weight of internal vertices on each path is at most $\lambda/2$.
On the right we see two different alternating circuits with the same signature: they differ in that one uses f_2 and the other uses g_3, g_2, g_1. The thick edges illustrate the alternating circuit $B = (f_1, f_2, g_1, g_2, g_3, g_4)$ of node-weight at most λ which leads to the contradiction. We walk the dashed path (P_D) in reverse.

(See also Figure 8 for an example.) We also let $\mathrm{NW}(v_i)$ denote the node-weight of vertex v_i, and $\mathrm{in}(v_i) = e_{(i-1) \bmod k}$ and $\mathrm{out}(v_i) = e_i$ be the incoming and outgoing edges of v_i in C.[20] We now define the signature $\sigma(C)$ as the output of the following procedure:

- Let $i_0 = 0$ be the index of the first vertex in C.
- For $j = 1, 2, 3$, select $i_j \leq k$ to be the largest index satisfying $\sum_{i=i_{j-1}+1}^{i_j - 1} \mathrm{NW}(v_i) \leq \lambda/2$.
- Let $t = \max\{j : i_j < k\}$ and output the signature $\sigma(C) = ((-1)^{i_j}, \mathrm{in}(v_{i_j}), \mathrm{out}(v_{i_j}))_{j=0,1,\ldots,t}$.

The intuition of the signature is as follows (see also the left part of Figure 8). The procedure starts at the first vertex $v_{i_0} = v_0$. It then selects the farthest (according to C) vertex v_{i_1} while guaranteeing that the total node-weight of the vertices visited in-between v_{i_0} and v_{i_1} is at most $\lambda/2$. Similarly, v_{i_2} is selected to be the farthest vertex such that the total node-weight of the vertices $v_{i_1+1}, \ldots, v_{i_2-1}$ is at most $\lambda/2$, and i_3 is selected in the same fashion. The indices $i_0, i_1, \ldots, i_t$ thus partition C into edge-disjoint paths

$$
\begin{aligned}
C_0 &= v_{i_0} \xrightarrow{e_{i_0}} v_{i_0+1} \xrightarrow{e_{i_0+1}} \ldots \xrightarrow{e_{i_1-2}} v_{i_1-1} \xrightarrow{e_{i_1-1}} v_{i_1} \\
C_1 &= v_{i_1} \xrightarrow{e_{i_1}} v_{i_1+1} \xrightarrow{e_{i_1+1}} \ldots \xrightarrow{e_{i_2-2}} v_{i_2-1} \xrightarrow{e_{i_2-1}} v_{i_2} \\
&\vdots \\
C_t &= v_{i_t} \xrightarrow{e_{i_t}} v_{i_t+1} \xrightarrow{e_{i_t+1}} \ldots \xrightarrow{e_{i_0-2}} v_{i_0-1} \xrightarrow{e_{i_0-1}} v_{i_0}
\end{aligned}
$$

so that the total node-weight of the internal vertices on each path is at most $\lambda/2$. Indeed, for C_j with $j < 3$ this follows from the selection of i_j. For C_3 (in the case $t = 3$), by maximality of

[20] The functions $\mathrm{in}(v)$ and $\mathrm{out}(v)$ are not formally well-defined since they depend on the considered alternating circuit and on which occurrence of v in the circuit we are considering, but their values will be clear from the context.

i_1, i_2 and i_3 we have

$$\underbrace{\sum_{i=i_0+1}^{i_1} \mathrm{NW}(v_i)}_{\geq \lambda/2} + \underbrace{\sum_{i=i_1+1}^{i_2} \mathrm{NW}(v_i)}_{\geq \lambda/2} + \underbrace{\sum_{i=i_2+1}^{i_3} \mathrm{NW}(v_i)}_{\geq \lambda/2} \geq \frac{3}{2}\lambda$$

and so the internal vertices of C_3, which are disjoint from the vertices in the above sums, can have node-weight at most $\lambda/2$ (the total node-weight of C being at most 2λ).

We now count the number of possible signature vectors. As for each j there are at most n^2 ways of choosing the incoming edge, at most n^2 ways of choosing the outgoing edge, and i_j can have two different parities, the number of possible signatures is at most (summing over the choices of $t = 0, 1, 2, 3$) $(2n^2 \cdot n^2) + (2n^2 \cdot n^2)^2 + (2n^2 \cdot n^2)^3 + (2n^2 \cdot n^2)^4 < n^{17}$.

It remains to be shown that any two alternating circuits C and D in H of node-weight at most 2λ have different signatures if $(\pm\mathbb{1})_C \neq (\pm\mathbb{1})_D$. Suppose that $(\pm\mathbb{1})_C \neq (\pm\mathbb{1})_D$ but $\sigma(C) = \sigma(D)$. We would like to derive a contradiction with the assumption (in Theorem 5.1) that H contains no alternating circuit of node-weight at most λ. This will finish the proof.

As described above, C can be partitioned into disjoint paths $C_0, \ldots, C_t$ using its indices $i_0, \ldots, i_t$. Similarly we partition D into $D_0, \ldots, D_t$. Since these are disjoint unions, $(\pm\mathbb{1})_C \neq (\pm\mathbb{1})_D$ implies that at least one of the four subpaths must be different between C and D, in the sense that the part of the alternating indicator vector arising from that subpath is different. More formally, there is $j \in \{0, \ldots, t\}$ such that $(\pm\mathbb{1})_{C_j} \neq (\pm\mathbb{1})_{D_j}$. We will "glue" together the paths C_j and D_j to obtain another alternating circuit B. First notice that both C_j and D_j are paths of the form $v_{i_j} = a \to b \to \cdots \to c \to d = v_{i_{j+1 \bmod t}}$, where the segment from b to c differs between them.[21] This follows from the assumption that $\sigma(C) = \sigma(D)$. Let P_C denote the path from b to c in C_j and let P_D denote the path from b to c in D_j. As the parity fields of the signatures agree, we have that $|P_C| + |P_D|$ is even. Now let B be the cyclic walk of even length obtained by walking from b to c along path P_C and back from c to b along the path P_D (in reverse). That is, B is of the form (see also the right part of Figure 8)

$$b \xrightarrow{f_1} \ldots \xrightarrow{f_{|P_C|}} c \xrightarrow{g_1} \ldots \xrightarrow{g_{|P_D|}} b\,,$$

where we let $f_1, \ldots, f_{|P_C|}$ denote the edges of P_C and $g_1, \ldots, g_{|P_D|}$ denote the edges of the reversed path P_D. To verify that B is an alternating circuit we need to show that its alternating indicator vector is nonzero:

$$\begin{aligned}
-(\pm\mathbb{1})_B &= \sum_{i=1}^{|P_C|} (-1)^i \mathbb{1}_{f_i} + \sum_{i=1}^{|P_D|} (-1)^{|P_C|+i} \mathbb{1}_{g_i} \\
&= \underbrace{\left((-1)^0 \mathbb{1}_{\mathrm{out}(a)} + \sum_{i=1}^{|P_C|} (-1)^i \mathbb{1}_{f_i} + (-1)^{|P_C|+1} \mathbb{1}_{\mathrm{in}(d)} \right)}_{=(\pm\mathbb{1})_{C_j}} \\
&\qquad + \underbrace{\left((-1)^{|P_C|+2} \mathbb{1}_{\mathrm{in}(d)} + \sum_{i=1}^{|P_D|} (-1)^{|P_C|+2+i} \mathbb{1}_{g_i} + (-1)^{|P_C|+|P_D|+3} \mathbb{1}_{\mathrm{out}(a)} \right)}_{=-(\pm\mathbb{1})_{D_j}}.
\end{aligned}$$

[21] Here again we slightly abuse notation since i_j might differ between C and D; however, the vertex v_{i_j} does not, because it is the tail of e_{i_j}, which is part of the signature $\sigma(C) = \sigma(D)$. The same applies to $v_{i_{j+1 \bmod t}}$.

The second equality is easiest to see by mentally extending B from a circuit $b \to \dots \to c \to \dots \to b$ to $a \to b \to \dots \to c \to d \to c \to \dots \to b \to a$; also recall that $|P_C| + |P_D|$ is even. Thus we get $(\pm\mathbb{1})_B = -(\pm\mathbb{1})_{C_j} + (\pm\mathbb{1})_{D_j}$, which is nonzero by the choice of j. Finally, the node-weight of B is at most the node-weight of the internal nodes of path C_j plus the node-weight of the internal nodes of path D_j and thus at most $\lambda/2 + \lambda/2 = \lambda$.

We have thus shown that B is a nonempty cyclic walk of even length whose alternating indicator vector is nonzero (thus an alternating circuit) and whose node-weight is at most λ. This contradicts our assumption on H. $\square$

Now we have all the tools needed to prove the main result of this section.

Proof of Theorem 5.1. Let us fix some $w \in \mathcal{W}$. We want to articulate conditions on w which will make sure that the statement is satisfied. Then we show that some $w \in \mathcal{W}$ satisfies these conditions.

Let C be any alternating circuit in H of node-weight at most 2λ. Our condition on w will be that all such circuits C should not respect $F[w]$, i.e., that all vectors from the set

$$\mathcal{Z} := \{(\pm\mathbb{1})_C : C \text{ is an alternating circuit in } H \text{ of node-weight at most } 2\lambda\}$$

should not respect $F[w]$. We use Lemma 5.3 to transform each $z \in \mathcal{Z}$ ($z \in \mathbb{Z}^{E(H)}$) to a vector $y = y(z) \in \mathbb{Z}^E$ such that if $y(z)$ does not respect $F[w]$, then z does not respect $F[w]$. Let $\mathcal{Y} = \{y(z) : z \in \mathcal{Z}\}$. Clearly $|\mathcal{Y}| \le |\mathcal{Z}|$,[22] and $|\mathcal{Z}| \le n^{17}$ by Lemma 5.4. Moreover, since the alternating circuits C were of node-weight at most $2\lambda \le 4n$, we have $\|z\|_1 \le 4n$ for $z \in \mathcal{Z}$ and $\|y\|_1 \le 4n^2$ for $y \in \mathcal{Y}$. Now it is enough to apply Corollary 3.5 to obtain that there exists $w \in \mathcal{W}(n^3 \cdot n^{17}) = \mathcal{W}$ such that each $y \in \mathcal{Y}$ does not respect the face $F[w]$. $\square$

5.2 The existence of a good weight function

In this section, we use Theorem 5.1 to prove the existence of a weight function defining a face F_{out} with the desired properties (so as to be the face in our 2λ-good face-laminar pair).

Theorem 5.5. *Let $(F_{\mathsf{in}}, \mathcal{L}_{\mathsf{in}})$ be a λ-good face-laminar pair. There exists a weight function $w_{\mathsf{out}} \in \mathcal{W}^{\log_2(n)+1}$ such that the face $F_{\mathsf{out}} = F_{\mathsf{in}}[w_{\mathsf{out}}]$ satisfies:*

(i)' For each $S \in \mathcal{L}_{\mathsf{in}}$ with $|S| \le 2\lambda$, S is F_{out}-contractible.

(ii)' In the $(F_{\mathsf{in}}, \mathcal{L}_{\mathsf{in}}, 2\lambda)$-contraction of G, there is no F_{out}-respecting alternating circuit of node-weight at most 2λ.

Throughout this section, F_{in} and $\mathcal{L}_{\mathsf{in}}$ are as in the statement of Theorem 5.5. The proof of Theorem 5.5 is based on the following technical lemma.

Lemma 5.6. *There exists a weight function $w_{\mathsf{mid}} \in \mathcal{W}^{\log_2 n}$ such that the face $F_{\mathsf{mid}} = F_{\mathsf{in}}[w_{\mathsf{mid}}]$ satisfies:*

(i)' For each $S \in \mathcal{L}_{\mathsf{in}}$ with $|S| \le 2\lambda$, S is F_{mid}-contractible.

Before giving the proof of Lemma 5.6, let us see how it together with Theorem 5.1 rather easily implies Theorem 5.5.

[22] Actually $|\mathcal{Y}| = |\mathcal{Z}|$ since the mapping $z \mapsto y(z)$ is one-to-one.

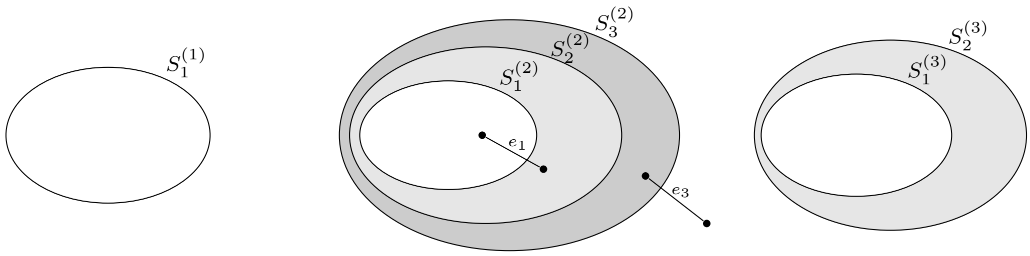

Figure 9: An example of the laminar family $\mathcal{L}$, which consists of disjoint chains. The different shades of gray depict the sets $U_p^{(i)}$.

Proof of Theorem 5.5. Lemma 5.6 says that there is a weight function $w_{\text{mid}} \in \mathcal{W}^{\log_2(n)}$ such that the face $F_{\text{mid}} = F_{\text{in}}[w_{\text{mid}}]$ satisfies that every $S \in \mathcal{L}_{\text{in}}$ with $|S| \leq 2\lambda$ is F_{mid}-contractible. We have thus proved point $(i)'$ of Theorem 5.5, as any set that is F_{mid}-contractible remains contractible in any subface of F_{mid} (Fact 4.2).

By the above, we have that every vertex in the $(F_{\text{in}}, \mathcal{L}_{\text{in}}, 2\lambda)$-contraction of G corresponds to an F_{mid}-contractible set. Moreover, by the assumption that the face-laminar pair $(F_{\text{in}}, \mathcal{L}_{\text{in}})$ is λ-good, the $(F_{\text{in}}, \mathcal{L}_{\text{in}}, \lambda)$-contraction of G does not have any alternating circuits of node-weight at most λ. This implies that the $(F_{\text{in}}, \mathcal{L}_{\text{in}}, 2\lambda)$-contraction of G does not have any alternating circuits of node-weight at most λ. For suppose C were such an alterating circuit. Let $S_1, ..., S_k$ be maximal sets of size at most 2λ in $\mathcal{L}_{\text{in}}$, i.e., the vertices of the $(F_{\text{in}}, \mathcal{L}_{\text{in}}, 2\lambda)$-contraction of G. Note that C cannot cross a set S_i with $|S_i| > \lambda$, because then its node-weight would be larger than λ. Therefore C only crosses sets S_i with $|S_i| \leq \lambda$. Thus C also appears in the $(F_{\text{in}}, \mathcal{L}_{\text{in}}, \lambda)$-contraction of G, with the same node-weight. This is a contradiction.

We can thus apply Theorem 5.1 (with $\beta = 2\lambda$) to get a weight function $w \in \mathcal{W}$ such that the face $F_{\text{out}} = F_{\text{mid}}[w]$ satisfies $(ii)'$. Selecting $w_{\text{out}} = w_{\text{mid}} \circ w \in \mathcal{W}^{\log_2(n)+1}$ completes the proof (Fact 2.6). $\square$

The rest of this section is devoted to the proof of Lemma 5.6. Recall that we need to prove the existence of a weight function $w_{\text{mid}} \in \mathcal{W}^{\log_2(n)}$ satisfying $(i)'$, i.e., that

$$\text{for each } S \in \mathcal{L}_{\text{in}} \text{ with } |S| \leq 2\lambda,\ S \text{ is } F_{\text{mid}}\text{-contractible},$$

where $F_{\text{mid}} = F_{\text{in}}[w_{\text{mid}}]$. First note that the statement will be true for every $S \in \mathcal{L}_{\text{in}}$ with $|S| \leq \lambda$ regardless of the choice of the weight function w_{mid}. Indeed, by assumption, $(F_{\text{in}}, \mathcal{L}_{\text{in}})$ is λ-good and so S is F_{in}-contractible. Thus, by Fact 4.2, S remains F_{mid}-contractible for any subface $F_{\text{mid}} \subseteq F_{\text{in}}$.

It remains to deal with the sets $S \in \mathcal{L}_{\text{in}}$ with $\lambda < |S| \leq 2\lambda$. Let $\mathcal{L} = \{S \in \mathcal{L}_{\text{in}} : \lambda < |S| \leq 2\lambda\}$ be the laminar family $\mathcal{L}_{\text{in}}$ restricted to these sets. Notice that any set in $\mathcal{L}$ can have at most one child in the laminar family $\mathcal{L}$ due to the cardinality constraints. In other words, $\mathcal{L}$ consists of disjoint chains, as depicted in Figure 9.

Notation. We refer to the sets in the i-th chain of $\mathcal{L}$ by $\mathcal{L}_i$. Let $\ell_i = |\mathcal{L}_i|$ and index the sets of the chain $\mathcal{L}_i = \{S_1^{(i)}, S_2^{(i)}, \dots, S_{\ell_i}^{(i)}\}$ so that $S_1^{(i)} \subseteq S_2^{(i)} \subseteq \dots \subseteq S_{\ell_i}^{(i)}$. Let $U_1^{(i)} = S_1^{(i)}$ and $U_p^{(i)} = S_p^{(i)} \setminus S_{p-1}^{(i)}$ for $p = 2, 3, \dots, \ell_i$. Also define $U_{p,r}^{(i)} = U_p^{(i)} \cup U_{p+1}^{(i)} \cup \dots \cup U_r^{(i)}$.

Recall that $U_{1,r}^{(i)} = S_r^{(i)}$ is defined to be F-contractible if for every $e_r \in \delta(U_{1,r}^{(i)})$ there are no two perfect matchings in F which both contain e_r and are different inside $U_{1,r}^{(i)}$ (Definition 4.1).

For the proof of Lemma 5.6, we generalize this definition to also include sets $U^{(i)}_{p,r}$ with $p \geq 2$.

Definition 5.7. *Consider a face F. We say that a set $U^{(i)}_{p,r}$ with $2 \leq p \leq r \leq \ell_i$ is F-contractible if for every $e_{p-1} \in \delta(S^{(i)}_{p-1})$ and $e_r \in \delta(S^{(i)}_r)$ there are no two perfect matchings in F which both contain e_{p-1} and e_r and are different inside $U^{(i)}_{p,r}$.*[23]

The intuition of this definition is similar to that of Definition 4.1. Consider the second chain of Figure 9. If we, for example, restrict our attention to perfect matchings that must contain edges $e_1 \in \delta(S^{(2)}_1)$ and $e_3 \in \delta(S^{(2)}_3)$, then, as $S^{(2)}_1$ are $S^{(2)}_3$ are tight sets, the task of selecting such a matching decomposes into two *independent* problems: the problem of selecting a perfect matching in $U^{(2)}_{2,3}$ (ignoring the vertices incident to e_1 and e_3) and the problem of selecting a perfect matching in $V \setminus U^{(2)}_{2,3}$ (again ignoring the vertices incident to e_1 and e_3).

The proof now proceeds iteratively as follows.

- First we select $w_1 \in \mathcal{W}$ such that $F_1 = F_{\mathsf{in}}[w_1]$ satisfies:

$$U^{(i)}_p \text{ is } F_1\text{-contractible for all chains } i \text{ and } 1 \leq p \leq \ell_i. \tag{1}$$

- For $t = 2, 3, \ldots, \log_2(n)$ we select $w_t \in \mathcal{W}$ such that $F_t = F_{t-1}[w_t]$ satisfies:

$$U^{(i)}_{p,r} \text{ is } F_t\text{-contractible for all chains } i \text{ and } 1 \leq p \leq r \leq \ell_i \text{ with } r - p \leq 2^{t-1} - 1. \tag{2}$$

We remark that, having selected $w_1, w_2, \ldots, w_{\log_2(n)}$ as above, if we let $w_{\mathsf{mid}} = w_1 \circ w_2 \circ \ldots \circ w_{\log_2(n)}$, then the face $F_{\mathsf{mid}} = F_{\mathsf{in}}[w_{\mathsf{mid}}]$ equals $F_{\log_2(n)}$ (Fact 2.6). To see that this completes the proof of Lemma 5.6, note that $\ell_i < n/2$ for any chain i since $|S^{(i)}_1| > |S^{(i)}_{\ell_i}|/2$ and $|S^{(i)}_{\ell_i}| \leq n$. We have thus that any set $S^i_r \in \mathcal{L}$ has $r \leq n/2$ and so, by (2), $S^i_r = U^i_{1,r}$ is F_{mid}-contractible. In what follows, we complete the proof of Lemma 5.6 with a description of how to select w_1, followed by the selection of w_t in the iterative case.

5.2.1 The selection of w_1

The following claim allows us to use Theorem 5.1 to show the existence of a weight function w_1 satisfying (1).

Claim 2. *If the $(F_{in}, \mathcal{L}_{in}, \lambda)$-contraction of G has no F_1-respecting alternating circuit of node-weight at most 2λ, then every $U^{(i)}_p$ is F_1-contractible.*

Proof. This proof resembles that of Lemma 4.10. Fix $U^{(i)}_p$, and let $e_{p-1} \in \delta(S^{(i)}_{p-1})$ and $e_p \in \delta(S^{(i)}_p)$; suppose that M_1 and M_2 are two perfect matchings in F_1 which both contain e_{p-1} and e_p.[24] We want to show that M_1 and M_2 are equal inside $U^{(i)}_p$.

Let $S_1, ..., S_k$ be all maximal sets $S \in \mathcal{L}_{\mathsf{in}}$ with $S \subseteq U^{(i)}_p$. (They are those vertices of the $(F_{\mathsf{in}}, \mathcal{L}_{\mathsf{in}}, \lambda)$-contraction which lie in $U^{(i)}_p$, and we have $S_1 \cup ... \cup S_k = U^{(i)}_p$.) Because these sets (as well as $S^{(i)}_{p-1}$ and $S^{(i)}_p$) are tight for F_1, any perfect matching (on G) in F_1 containing e_{p-1} and e_p induces an almost-perfect matching on $S_1, ..., S_k$, that is, one where only the (up to two)

[23] It is possible that $e_{p-1} = e_r$, in which case neither endpoint of this edge lies in $U^{(i)}_{p,r}$.

[24] Here and in Section 5.2.2, we abuse notation and assume $p \geq 2$; the only difference is that, given a set $U^{(i)}_{p,r}$ with $p = 1$ (in this section $p = r$), we consider matchings containing one edge e_r ($e_r \in \delta(S^{(i)}_r)$) instead of matchings containing two edges e_{p-1} and e_r ($e_{p-1} \in \delta(S^{(i)}_{p-1})$ and $e_r \in \delta(S^{(i)}_r)$) (since if $p = 1$, the set $S^{(i)}_{p-1}$ is not defined).

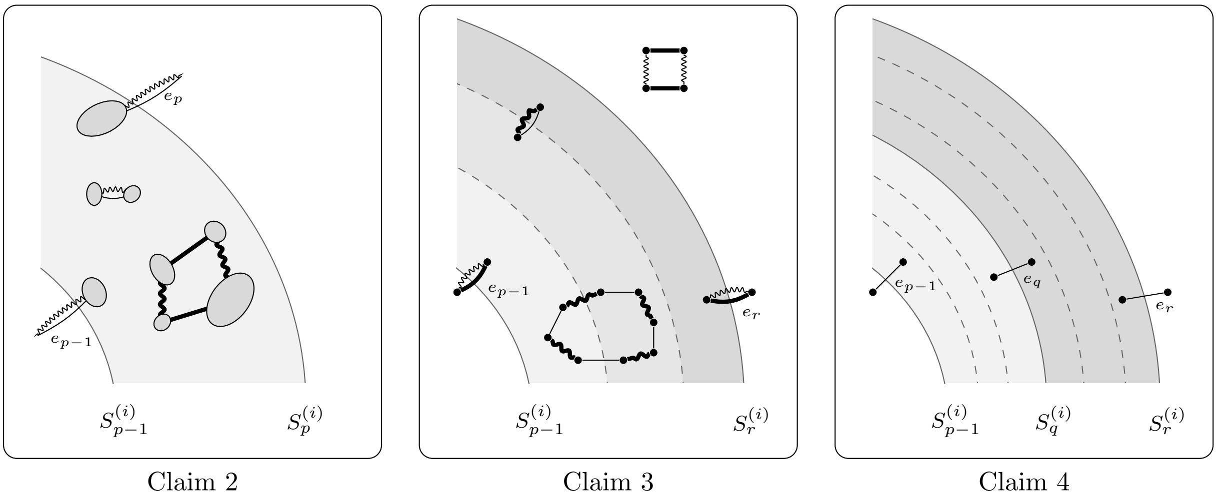

Figure 10: An illustration of the different claims used in the proof of Lemma 5.6.
Claim 2: Straight and swirly edges denote M_1 and M_2 respectively. The thick edges denote the alternating cycle. The dark-gray sets are $S_1, ..., S_k$.
Claim 3: Straight and swirly edges denote M_1 and M_2 respectively. The thick edges denote M_{12}, which agrees with M_1 outside $U_{p,r}^{(i)}$ and with M_2 inside $U_{p,r}^{(i)}$.
Claim 4: The divide-and-conquer argument is illustrated (only edges e_{p-1}, e_q, e_r are depicted). After fixing e_{p-1} and e_q, the matching in the light-gray area is unique in the face F_{t-1}. Similarly, after fixing e_q and e_r, the matching in the dark-gray area is unique in the face F_{t-1}. Therefore, for each choice of e_{p-1} and e_r, there can be at most one matching inside $U_{p,r}^{(i)}$ for each possible way of fixing e_q.

sets S_i containing endpoints of e_{p-1} and e_p are unmatched (see the left part of Figure 10). If the matchings induced by M_1 and M_2 were different, then their symmetric difference would contain an alternating simple cycle in the $(F_{\mathsf{in}}, \mathcal{L}_{\mathsf{in}}, \lambda)$-contraction. Since this cycle arises from two matchings in F_1, it respects F_1.[25] Moreover, since it is a simple cycle inside $U_p^{(i)}$, its node-weight is at most $|U_p^{(i)}| \leq 2\lambda$. This would contradict our assumption.

Therefore, the induced matchings must be equal. Moreover, the sets $S_1, ..., S_k$ are F_1-contractible (since they are vertices of the $(F_{\mathsf{in}}, \mathcal{L}_{\mathsf{in}}, \lambda)$-contraction and $F_1 \subseteq F_{\mathsf{in}}$), which means that, given the boundary edges (i.e., the induced matching plus e_{p-1} and e_p), there is a unique perfect matching in F_1 inside each S_i. It follows that M_1 and M_2 are equal inside $U_p^{(i)}$. $\square$

The claim together with Theorem 5.1 completes the selection of w_1 as follows. Since $(F_{\mathsf{in}}, \mathcal{L}_{\mathsf{in}})$ is λ-good, we can apply Theorem 5.1 (with $\beta = \lambda$) to obtain the existence of a weight function $w_1 \in \mathcal{W}$ such that, in the $(F_{\mathsf{in}}, \mathcal{L}, \lambda)$-contraction of G, there is no F_1-respecting alternating circuit of node-weight at most 2λ, where $F_1 = F_{\mathsf{in}}[w_1]$. Hence, by the above claim, every $U_p^{(i)}$ is F_1-contractible as required.

[25] This is in similar vein as the proof of Claim 1. Call the cycle C. Clearly, $\mathrm{supp}(C) \subseteq M_1 \cup M_2 \subseteq E(F_1)$. Let $T \in \mathcal{S}(F_1)$ be a set tight for F_1 which is a union of the vertices of the contraction; we want to show that $\langle (\pm \mathbb{1})_C, \mathbb{1}_{\delta(T)} \rangle = 0$. Because C is a cycle in the contraction and $|M_1 \cap \delta(T)| = |M_2 \cap \delta(T)| = 1$, either C has no edge in $\delta(T)$ or it has two, one from M_1 and one from M_2 (and they cancel out).

5.2.2 The selection of w_t for $t = 2, 3, \ldots, \log_2(n)$

In this section, we show the existence of a weight function $w_t \in \mathcal{W}$ satisfying (2), i.e.,

$$U_{p,r}^{(i)} \text{ is } F_t\text{-contractible for all chains } i \text{ and } 1 \le p \le r \le \ell_i \text{ with } r - p \le 2^{t-1} - 1,$$

where $F_t = F_{t-1}[w_t]$.

The proof outline is as follows. First, in Claim 3, we give sufficient conditions on w_t for $U_{p,r}^{(i)}$ being F_t-contractible. They are given as a system of linear non-equalities with coefficients in $\{-1, 0, 1\}$. Then, in Claim 4, we upper-bound the number of these non-equalities by n^{11}. This allows us to deduce the existence of $w_t \in \mathcal{W}$ by applying Lemma 3.4.

The following claim gives sufficient linear non-equalities on w_t for every $U_{p,r}^{(i)}$ to be F_t-contractible (one non-equality for each choice of $U_{p,r}^{(i)}$, e_{p-1}, e_r, M_1 and M_2).

Claim 3. *Consider $U = U_{p,r}^{(i)}$ for some chain i and $1 \le p \le r \le \ell_i$. Suppose that for every two edges $e_{p-1} \in \delta(S_{p-1}^{(i)})$ and $e_r \in \delta(S_r^{(i)})$ defining a face $F = \{x \in F_{t-1} : x_{e_{p-1}} = 1, x_{e_r} = 1\}$, we have:*

$$w_t(M_1 \cap E(U)) \neq w_t(M_2 \cap E(U)) \quad \textit{for any two matchings } M_1, M_2 \textit{ in } F \textit{ that differ inside } U.$$

Then U is F_t-contractible.

Proof. We prove the contrapositive. Suppose that U is not F_t-contractible. Then, by definition, there must be $e_{p-1} \in \delta(S_{p-1}^{(i)})$ and $e_r \in \delta(S_r^{(i)})$ that define a face $F' = \{x \in F_t : x_{e_{p-1}} = 1, x_{e_r} = 1\}$ such that there are two matchings M_1 and M_2 in F' that differ inside U. Notice that $F' \subseteq F = \{x \in F_{t-1} : x_{e_{p-1}} = 1, x_{e_r} = 1\}$; therefore M_1 and M_2 are also two matchings in F that differ inside U.

We complete the proof of the claim by showing that

$$w_t(M_1 \cap E(U)) = w_t(M_2 \cap E(U)) \,. \tag{3}$$

Define

$$M_{12} = (M_1 \setminus E(U)) \cup (M_2 \cap E(U))$$

to be the perfect matching that agrees with M_1 on all edges not in $E(U)$ and agrees with M_2 on all edges in $E(U)$ (see the central part of Figure 10 for an example). By the same argument as in the proof of Lemma 4.3, M_{12} is a perfect matching in F'. It differs from M_1 inside U and agrees with M_1 outside U.

We now use that M_1 and M_{12} are perfect matchings in F' to prove (3). As F_t is the convex-hull of matchings in F_{t-1} that minimize the objective function w_t, all matchings M in F_t and in its subface F' have the same weight $w_t(M)$. In particular,

$$w_t(M_1) = w_t(M_1 \setminus E(U)) + w_t(M_1 \cap E(U)) = w_t(M_1 \setminus E(U)) + w_t(M_2 \cap E(U)) = w_t(M_{12}) \,,$$

and thus $w_t(M_1 \cap E(U)) = w_t(M_2 \cap E(U))$ as required. □

The above claim says that it is sufficient to write down a non-equality for each choice of $U_{p,r}^{(i)}, e_{p-1}, e_r, M_1$ and M_2. It is easy to upper-bound the number of ways of choosing i, p, r, e_{p-1} and e_r. The following claim bounds the number of ways of choosing M_1 and M_2. The proof is based on a divide-and-conquer strategy (see the right part of Figure 10) and it uses the inductive assumption that $U_{p,r}^{(i)}$ is F_{t-1}-contractible for all chains i and $1 \le p \le r \le \ell_i$ with $r - p \le 2^{t-2} - 1$.

Claim 4. *Consider $U = U_{p,r}^{(i)}$ with $r - p \leq 2^{t-1} - 1$ and define $q = \lfloor (p+r)/2 \rfloor$. For any two edges $e_p \in \delta(S_{p-1}^{(i)})$ and $e_r \in \delta(S_r^{(i)})$ defining a face $F = \{x \in F_{t-1} : x_{e_{p-1}} = 1, x_{e_r} = 1\}$, we have*

$$|\{M \cap E(U) : M \text{ is a matching in } F\}| \leq |\delta(S_q^{(i)}) \cap E(F)| \leq n^2 \,.$$

We remark that the first inequality holds with equality, but we only need the inequality.

Proof. The second inequality in the statement is trivial. We prove the first.

As F_{t-1} and thus F is a subface of F_{in}, any matching M in F must satisfy $M \cap \delta(S_q^{(i)}) = \{e_q\}$ for some edge $e_q \in \delta(S_q^{(i)}) \cap E(F)$, where we use that $S_q^{(i)} \in \mathcal{L}_{\mathsf{in}} \subseteq \mathcal{S}(F_{\mathsf{in}})$ is a tight set in these faces.

We prove the statement by showing that for every choice of e_q, any matching M in the face $F_{e_q} = \{x \in F : x_{e_q} = 1\}$ matches the nodes in $U_{p,r}^{(i)}$ in a unique way. In other words, we show that $|\{M \cap E(U) : M \text{ is a matching in } F_{e_q}\}| \leq 1$ for every $e_q \in \delta(S_q^{(i)}) \cap E(F)$, which implies

$$\begin{aligned} |\{M \cap E(U) : M \text{ is a matching in } F\}| &\leq \sum_{e_q \in \delta(S_q^{(i)}) \cap E(F)} |\{M \cap E(U) : M \text{ is a matching in } F_{e_q}\}| \\ &\leq |\delta(S_q^{(i)}) \cap E(F)| \,. \end{aligned}$$

To prove that $|\{M \cap E(U) : M \text{ is a matching in } F_{e_q}\}| \leq 1$, suppose the contrary, i.e., that $|\{M \cap E(U) : M \text{ is a matching in } F_{e_q}\}| \geq 2$. Take two such matchings M_1 and M_2 that differ inside U. By the definition of F_{e_q}, $M_1 \cap \delta(S_q^{(i)}) = M_2 \cap \delta(S_q^{(i)}) = \{e_q\}$ and so M_1 and M_2 must differ inside $U_{p,q}^{(i)}$ or inside $U_{q+1,r}^{(i)}$; assume the former (the argument for the other case is the same). Notice that M_1 and M_2 are two matchings in $F_{e_q} \subseteq F_{t-1}$ which both contain e_{p-1} and e_q but differ inside $U_{p,q}^{(i)}$, which contradicts that $U_{p,q}^{(i)}$ is F_{t-1}-contractible. (Note that $q - p \leq (r-p)/2 \leq 2^{t-2} - 1/2$, which implies that $q - p \leq 2^{t-2} - 1$.) □

We now have all the needed tools to show the existence of a weight function $w_t \in \mathcal{W}$ such that the face $F_t = F_{t-1}[w_t]$ satisfies (2): for all chains i and $1 \leq p \leq r \leq \ell_i$ with $r - p \leq 2^{t-1} - 1$, $U_{p,r}^{(i)}$ is F_t-contractible. By Claim 3, this holds if for any $U = U_{p,r}^{(i)}$ with $r - p \leq 2^{t-1} - 1$ and for any $e_{p-1} \in \delta(S_{p-1}^{(i)})$ and $e_r \in \delta(S_r^{(i)})$ defining a face $F = \{x \in F_{t-1} : x_{e_{p-1}} = 1, x_{e_r} = 1\}$ we have the following:

$$w_t(M_1 \cap E(U)) - w_t(M_2 \cap E(U)) \neq 0 \quad \text{for any two matchings } M_1, M_2 \text{ in } F \text{ that differ inside } U.$$

There are at most n ways of choosing i, n ways of choosing p, n ways of choosing r, n^2 ways of choosing e_{p-1}, n^2 ways of choosing e_r, and by Claim 4 there are at most n^4 ways of choosing M_1 and M_2. In total, we can write the sufficient conditions on the weight function w_t as a system of at most n^{11} linear non-equalities with coefficients in $\{-1, 0, 1\}$. It follows by Lemma 3.4 that there is a weight function $w_t \in \mathcal{W}(n^{14}) \subseteq \mathcal{W}(n^{20}) = \mathcal{W}$ satisfying these conditions. This completes the selection of w_t and the proof of Lemma 5.6.

5.3 A maximal laminar family completes the proof

Theorem 5.5 demonstrates the existence of a weight function w_{out} that defines a face F_{out} with properties $(i)'$ and $(ii)'$. We now show that extending $\mathcal{L}_{\mathsf{in}}$ to a maximal laminar family $\mathcal{L}_{\mathsf{out}}$ of $\mathcal{S}(F_{\mathsf{out}})$ yields a 2λ-good face-laminar pair.[26] As explained at the beginning of Section 5, this will complete the proof of Theorem 4.9.

[26] Such an extension is possible because $\mathcal{L}_{\mathsf{in}}$ consists of sets which are tight for F_{in}, therefore also for F_{out}.

Why a maximal laminar family? Part of our argument so far was about removing certain alternating circuits C; in other words, we have made C not respect the new face F_{out}. This means either not having some edge from $\mathrm{supp}(C)$ in the support $E(F_{\mathsf{out}})$ of F_{out}, or introducing a new odd-set S which is tight for F_{out} and such that $\left\langle (\pm\mathbb{1})_C, \mathbb{1}_{\delta(S)} \right\rangle \neq 0$. In the latter case, we want C to have an odd-set with this property also in the new laminar family $\mathcal{L}_{\mathsf{out}}$, so that the removal of C is reflected in the new contraction (which is based on $\mathcal{L}_{\mathsf{out}}$). Lemma 5.2 guarantees that this will happen if we choose $\mathcal{L}_{\mathsf{out}}$ to be a maximal laminar subset of $\mathcal{S}(F_{\mathsf{out}})$.

Lemma 5.8. *Let $(F_{in}, \mathcal{L}_{in})$ be a λ-good face-laminar pair, and $F_{out} \subseteq F_{in}$ be the face guaranteed by Theorem 5.5. Then $(F_{out}, \mathcal{L}_{out})$ is a 2λ-good face-laminar pair, where $\mathcal{L}_{out}$ is any maximal laminar family with $\mathcal{L}_{in} \subseteq \mathcal{L}_{out} \subseteq \mathcal{S}(F_{out})$.*

Proof of Lemma 5.8. Recall that Theorem 5.5 guarantees that:

(i)' each $S \in \mathcal{L}_{\mathsf{in}}$ with $|S| \leq 2\lambda$ is F_{out}-contractible,

(ii)' in the $(F_{\mathsf{out}}, \mathcal{L}_{\mathsf{in}}, 2\lambda)$-contraction of G, there is no alternating circuit of node-weight at most 2λ which respects F_{out}.

(The statement of Theorem 5.5 refers to the $(F_{\mathsf{in}}, \mathcal{L}_{\mathsf{in}}, 2\lambda)$-contraction of G, but since we are considering F_{out}-respecting alternating circuits, this is equivalent.)

We want to show that the pair $(F_{\mathsf{out}}, \mathcal{L}_{\mathsf{out}})$ satisfies Definition 4.7, that is,

(i) each $S \in \mathcal{L}_{\mathsf{out}}$ with $|S| \leq 2\lambda$ is F_{out}-contractible,

(ii) in the $(F_{\mathsf{out}}, \mathcal{L}_{\mathsf{out}}, 2\lambda)$-contraction of G, there is no alternating circuit of node-weight at most 2λ.

Property (i). Fix a set $S \in \mathcal{L}_{\mathsf{out}}$ with $|S| \leq 2\lambda$. Let $S_1, ..., S_k$ be all maximal subsets of S in $\mathcal{L}_{\mathsf{in}}$ (we have $S = S_1 \cup ... \cup S_k$). If S is contained in a set from $\mathcal{L}_{\mathsf{in}}$ of size at most 2λ, then that set is F_{out}-contractible by $(i)'$, and thus S is F_{out}-contractible by Lemma 4.3. So assume that is not the case; therefore each S_i is a maximal set of size at most 2λ in $\mathcal{L}_{\mathsf{in}}$, that is, a vertex of the $(F_{\mathsf{out}}, \mathcal{L}_{\mathsf{in}}, 2\lambda)$-contraction of G. By $(i)'$, each S_i is F_{out}-contractible.

Now the proof proceeds as in Claim 2; we present it for completeness. Let M_1 and M_2 be two perfect matchings in F_{out} which both contain an edge $e \in \delta(S)$. We want to show that M_1 and M_2 are equal inside S. Because S and $S_1, ..., S_k$ are tight for F_{out}, any perfect matching (on G) in F_{out} containing e induces an almost-perfect matching on $S_1, ..., S_k$, that is, one where only the set S_i containing the endpoint of e is unmatched. If the matchings induced by M_1 and M_2 were different, then their symmetric difference would contain an alternating simple cycle in the $(F_{\mathsf{out}}, \mathcal{L}_{\mathsf{in}}, 2\lambda)$-contraction. Since this cycle arises from two matchings in F_{out}, it respects F_{out}. Moreover, since it is a simple cycle inside S, its node-weight is at most $|S| \leq 2\lambda$. This would contradict our assumption (ii)'.

Therefore, the induced matchings must be equal. Moreover, the sets $S_1, ..., S_k$ are F_{out}-contractible, which means that, given the boundary edges (i.e., the induced matching plus e), there is a unique perfect matching in F_{out} inside each S_i. It follows that M_1 and M_2 are equal inside S.

Property (ii). Let H_{out} be the $(F_{\mathsf{out}}, \mathcal{L}_{\mathsf{out}}, 2\lambda)$-contraction of G and let H_{in} be the $(F_{\mathsf{out}}, \mathcal{L}_{\mathsf{in}}, 2\lambda)$-contraction of G. Thus H_{out} can also be obtained by further contracting H_{in} (and this will be our perspective)[27]. Suppose towards a contradiction that there is an alternating circuit C_{out} in H_{out} of node-weight at most 2λ.

[27] We also need to remove further edges – those from the boundaries of sets $S \in \mathcal{L}_{\mathsf{out}} \setminus \mathcal{L}_{\mathsf{in}}$ with $|S| > 2\lambda$.

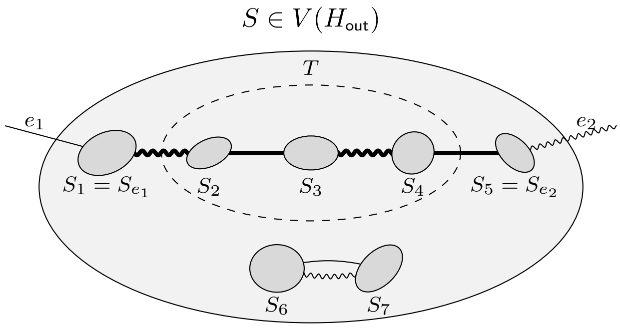

Figure 11: A depiction of the construction of paths $P_{e_1e_2}$ in the proof of Property (ii). The straight and swirly edges depict matchings M_1 and M_2, respectively. The path $P_{e_1e_2}$ is depicted by fat edges.

To obtain a contradiction, we are going to lift C_{out} back to an F_{out}-respecting alternating circuit C_{in} in H_{in}, which should not exist by $(ii)'$.[28] Namely, whenever C_{out} visits a vertex $S \in V(H_{\text{out}})$, we connect up the dangling endpoints of this visit inside S to obtain a walk in H_{in}. More precisely, let e_1 and e_2 be two consecutive edges of C_{out}, whose common endpoint in H_{out} is S; between them, we insert a simple path $P_{e_1e_2}$ inside the image of S in H_{in}, which is constructed as follows.

Since $e_1, e_2 \in \text{supp}(C_{\text{out}}) \subseteq E(F_{\text{out}})$, there exist matchings M_1 and M_2 (on G) in F_{out} containing e_1 and e_2, respectively. Let $S_1, ..., S_k$ be all maximal subsets of S in $\mathcal{L}_{\text{in}}$ (S_i are vertices of H_{in} and we have $S = S_1 \cup ... \cup S_k$). Denote by S_{e_1} and S_{e_2} the sets S_i which contain the S-endpoint of e_1 and e_2, respectively. The sets S and $S_1, ..., S_k$ are tight for F_{out}, so M_1 induces a perfect matching on $\{S_1, ..., S_k\} \setminus \{S_{e_1}\}$ (and similarly for M_2 and e_2). The symmetric difference of these two induced matchings contains a simple path $P_{e_1e_2}$ from S_{e_1} to S_{e_2} in H_{in} which has even length (possibly 0). For an example, see Figure 11. We obtain C_{in} by inserting such a path $P_{e_1e_2}$ between each two consecutive edges e_1, e_2 in C_{out}.

To obtain a contradiction, we need to prove that C_{in} is an alternating circuit of node-weight at most 2λ which respects F_{out}.

- That C_{in} is an alternating circuit follows by construction since each path $P_{e_1e_2}$ is of even length.

- For the node-weight, note that in C_{out}, the visit to S (on the edge e_1) incurs an increase of $|S|$, whereas in C_{in}, the visit to a certain subset of $\{S_1, ..., S_k\}$ (on e_1 and $P_{e_1e_2}$) incurs an increase of at most $|S_1| + ... + |S_k| \leq |S|$ because $P_{e_1e_2}$ is a simple path. Therefore the node-weight of C_{in} is at most that of C_{out}, thus at most 2λ.

- To see that C_{in} respects F_{out}, we use the assumption that $\mathcal{L}_{\text{out}}$ is a maximal laminar subset of $\mathcal{S}(F_{\text{out}})$ together with Lemma 5.2[29]. First, note that $\text{supp}(C_{\text{in}}) \subseteq E(F_{\text{out}})$ by construction. Second, let $T \in \mathcal{L}_{\text{out}}$ be a union of vertices of H_{in}; we need to show that $\left\langle (\pm\mathbb{1})_{C_{\text{in}}}, \mathbb{1}_{\delta(T)} \right\rangle = 0$.

  - If $|T| > 2\lambda$, then all boundary edges of T are absent from H_{out} (see Definition 4.5), so $\text{supp}(C_{\text{out}}) \cap \delta(T) = \emptyset$; in this case T is a union of vertices of H_{out} and so no path $P_{e_1e_2}$ contains any edges from $\delta(T)$ either, so that $\text{supp}(C_{\text{in}}) \cap \delta(T) = \emptyset$.

[28] This is in the same spirit as the proof of Lemma 5.3.

[29] Formally, note that $\mathcal{L}_{\text{out}}$ naturally maps to a laminar family $\mathcal{L}'_{\text{out}}$ of subsets of $V(H_{\text{in}})$, and that $\mathcal{L}'_{\text{out}}$ is also maximal; for if it were possible to add any set to $\mathcal{L}'_{\text{out}}$ while maintaining laminarity, then that set could also be used to enlarge $\mathcal{L}_{\text{out}}$. Therefore we can apply Lemma 5.2 to H_{in} and $\mathcal{L}'_{\text{out}}$.

– If $|T| \leq 2\lambda$, then T must be contained in a single set $S \in V(H_{\mathsf{out}})$ (as depicted in Figure 11).[30] For every path $P_{e_1 e_2}$ inside S, the path $e_1, P_{e_1 e_2}, e_2$ is a path from outside of S to outside of S which is part of the symmetric difference of two matchings in F_{out}. If this path enters T, it must also leave T. Suppose it entered T on an edge of the first matching; then it must exit T on an edge of the second matching (since $T \in \mathcal{L}_{\mathsf{out}} \subseteq \mathcal{S}(F_{\mathsf{out}})$ is tight for F_{out} and the matchings are in F_{out}) and the corresponding ± 1 terms cancel out. Abusing notation, we have $\left\langle (\pm \mathbb{1})_{e_1, P_{e_1 e_2}, e_2}, \mathbb{1}_{\delta(T)} \right\rangle = 0$. Since this holds for every path $P_{e_1 e_2}$ inside S, we get $\left\langle (\pm \mathbb{1})_{C_{\mathsf{in}}}, \mathbb{1}_{\delta(T)} \right\rangle = 0$ as required.

The existence of C_{in} contradicts $(ii)'$ and concludes the proof. □

A Proof of Lemma 2.2

The proof proceeds via the primal uncrossing technique; it is adapted from [LRS11]. Assume without loss of generality that $E = E(F)$.[31] We begin with an uncrossing lemma.

Lemma A.1 (uncrossing). *Let $S, T \in \mathcal{S}(F)$ be two sets which are crossing (i.e., $S \cap T, S \setminus T, T \setminus S \neq \emptyset$). Then:*

- *if $|S \cap T|$ is odd: then $S \cap T, S \cup T \in \mathcal{S}(F)$ and $\mathbb{1}_{\delta(S)} + \mathbb{1}_{\delta(T)} = \mathbb{1}_{\delta(S \cap T)} + \mathbb{1}_{\delta(S \cup T)}$,*
- *otherwise: $S \setminus T, T \setminus S \in \mathcal{S}(F)$ and $\mathbb{1}_{\delta(S)} + \mathbb{1}_{\delta(T)} = \mathbb{1}_{\delta(S \setminus T)} + \mathbb{1}_{\delta(T \setminus S)}$.*

Proof. **Case $|S \cap T|$ odd.** Note that we have

$$\mathbb{1}_{\delta(S)} + \mathbb{1}_{\delta(T)} = \mathbb{1}_{\delta(S \cap T)} + \mathbb{1}_{\delta(S \cup T)} + 2 \cdot \mathbb{1}_{\delta(S \setminus T, T \setminus S)}.$$

For any $x \in F'$, since $S, T \in \mathcal{S}(F)$ and because $S \cap T, S \cup T$ are nonempty odd sets, we have

$$1 + 1 = x(\delta(S)) + x(\delta(T)) = x(\delta(S \cap T)) + x(\delta(S \cup T)) + 2 \cdot x(\delta(S \setminus T, T \setminus S)) \geq 1 + 1 + 2 \cdot 0$$

where the inequality must be an equality, and thus $x(\delta(S \cap T)) = 1$, $x(\delta(S \cup T)) = 1$ (implying $S \cap T, S \cup T \in \mathcal{S}(F)$) and $x(\delta(S \setminus T, T \setminus S)) = 0$ for all $x \in F'$ (which, given that $E = E(F)$, implies that $\delta(S \setminus T, T \setminus S) = \emptyset$ and thus $\mathbb{1}_{\delta(S \setminus T, T \setminus S)} = 0$).

Case $|S \cap T|$ even. Now we have

$$\mathbb{1}_{\delta(S)} + \mathbb{1}_{\delta(T)} = \mathbb{1}_{\delta(S \setminus T)} + \mathbb{1}_{\delta(T \setminus S)} + 2 \cdot \mathbb{1}_{\delta(S \cap T, V \setminus (S \cup T))}.$$

The sets $S \setminus T$ and $T \setminus S$ are odd and nonempty, and we proceed as above. □

Define $\mathrm{cross}(S, \mathcal{L})$ to be the number of sets in $\mathcal{L}$ that cross S.

Proposition A.2. *If $S \notin \mathcal{L}$ and $T \in \mathcal{L}$ are crossing, then all four numbers* $\mathrm{cross}(S \cap T, \mathcal{L})$, $\mathrm{cross}(S \cup T, \mathcal{L})$, $\mathrm{cross}(S \setminus T, \mathcal{L})$ *and* $\mathrm{cross}(T \setminus S, \mathcal{L})$ *are smaller than* $\mathrm{cross}(S, \mathcal{L})$.

Proof. See Claim 9.1.6 in [LRS11]. □

Now we can prove Lemma 2.2. Towards a contradiction suppose that $\mathrm{span}(\mathcal{L}) \subsetneq \mathrm{span}(\mathcal{S}(F))$. Then there exists $S \in \mathcal{S}(F)$ with $\mathbb{1}_{\delta(S)} \notin \mathrm{span}(\mathcal{L})$. Pick such a set with minimum $\mathrm{cross}(S, \mathcal{L})$. Clearly $\mathrm{cross}(S, \mathcal{L}) \geq 1$ (otherwise $\mathcal{L} \cup \{S\}$ would be laminar, contradicting maximality of $\mathcal{L}$); let $T \in \mathcal{L}$ be a set crossing S. Assume that $|S \cap T|$ is odd; the other case is analogous. Then by Lemma A.1, $S \cap T, S \cup T \in \mathcal{S}(F)$ and

$$\mathbb{1}_{\delta(S)} + \mathbb{1}_{\delta(T)} = \mathbb{1}_{\delta(S \cap T)} + \mathbb{1}_{\delta(S \cup T)}. \tag{4}$$

By Proposition A.2 and our choice of S we have $\mathbb{1}_{\delta(S \cap T)}, \mathbb{1}_{\delta(S \cup T)} \in \mathrm{span}(\mathcal{L})$, and of course also $\mathbb{1}_{\delta(T)} \in \mathrm{span}(\mathcal{L})$. This and (4) implies that $\mathbb{1}_{\delta(S)} \in \mathrm{span}(\mathcal{L})$, a contradiction. □

[30] This is because the sets $S \in V(H_{\mathsf{out}})$ are maximal sets $S \in \mathcal{L}_{\mathsf{out}}$ with $|S| \leq 2\lambda$.

[31] We can do this since including the constraint $x_e = 0$ yields the same face as removing the edge e from G.

References

[AHT07] Manindra Agrawal, Thanh Minh Hoang, and Thomas Thierauf. The polynomially bounded perfect matching problem is in NC^2. In *STACS 2007, 24th Annual Symposium on Theoretical Aspects of Computer Science*, pages 489–499, 2007.

[AM08] Vikraman Arvind and Partha Mukhopadhyay. Derandomizing the isolation lemma and lower bounds for circuit size. In *APPROX and RANDOM*, pages 276–289, 2008.

[Bar92] D. A. M. Barrington. Quasipolynomial size circuit classes. In *Proceedings of the Seventh Annual Structure in Complexity Theory Conference*, pages 86–93, Jun 1992.

[BCH86] Paul W Beame, Stephen A Cook, and H James Hoover. Log depth circuits for division and related problems. *SIAM J. Comput.*, 15(4):994–1003, November 1986.

[Ber84] Stuart J. Berkowitz. On computing the determinant in small parallel time using a small number of processors. *Information Processing Letters*, 18(3):147–150, 1984.

[CNN89] Marek Chrobak, Joseph Naor, and Mark B. Novick. *Using bounded degree spanning trees in the design of efficient algorithms on claw-free graphs*, pages 147–162. Springer Berlin Heidelberg, Berlin, Heidelberg, 1989.

[Csa76] L. Csanky. Fast parallel inversion algorithm. *SIAM Journal of Computing*, 5:618–623, 1976.

[DHK93] E. Dahlhaus, P. Hajnal, and M. Karpinski. On the parallel complexity of Hamiltonian cycle and matching problem on dense graphs. *Journal of Algorithms*, 15(3):367 – 384, 1993.

[DK98] Elias Dahlhaus and Marek Karpinski. Matching and multidimensional matching in chordal and strongly chordal graphs. *Discrete Applied Mathematics*, 84(1-3):79–91, 1998.

[DKR10] Samir Datta, Raghav Kulkarni, and Sambuddha Roy. Deterministically isolating a perfect matching in bipartite planar graphs. *Theory Comput. Syst.*, 47(3):737–757, 2010.

[DS84] Eliezer Dekel and Sartaj Sahni. A parallel matching algorithm for convex bipartite graphs and applications to scheduling. *Journal of Parallel and Distributed Computing*, 1(2):185 – 205, 1984.

[Edm65a] Jack Edmonds. Maximum matching and a polyhedron with $0,1$ vertices. *Journal of Research of the National Bureau of Standards*, 69:125–130, 1965.

[Edm65b] Jack Edmonds. Paths, trees, and flowers. *Canadian Journal of Mathematics*, 17:449–467, 1965.

[FGT16] Stephen A. Fenner, Rohit Gurjar, and Thomas Thierauf. Bipartite perfect matching is in Quasi-NC. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 754–763, 2016.

[GK87] Dima Grigoriev and Marek Karpinski. The matching problem for bipartite graphs with polynomially bounded permanents is in NC. In *28th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 166–172, 1987.

[GT16] Rohit Gurjar and Thomas Thierauf. Linear matroid intersection is in quasi-NC. *Electronic Colloquium on Computational Complexity (ECCC)*, 23:182, 2016. To appear in STOC 2017.

[Har09] Nicholas J. A. Harvey. Algebraic algorithms for matching and matroid problems. *SIAM J. Comput.*, 39(2):679–702, 2009.

[Kas67] P. W. Kasteleyn. Graph theory and crystal physics. In F. Harary, editor, *Graph Theory and Theoretical Physics*, pages 43–110. Academic Press, 1967.

[KUW86] Richard M. Karp, Eli Upfal, and Avi Wigderson. Constructing a perfect matching is in random NC. *Combinatorica*, 6(1):35–48, 1986.

[KVV85] Dexter Kozen, Umesh V. Vazirani, and Vijay V. Vazirani. NC algorithms for comparability graphs, interval gaphs, and testing for unique perfect matching. In *Proceedings of the Fifth Conference on Foundations of Software Technology and Theoretical Computer Science*, pages 496–503, London, UK, UK, 1985. Springer-Verlag.

[Lov79] László Lovász. On determinants, matchings, and random algorithms. In *FCT*, pages 565–574, 1979.

[LPV81] G. F. Lev, N. Pippenger, and L. G. Valiant. A fast parallel algorithm for routing in permutation networks. *IEEE Transactions on Computers*, C-30(2):93–100, Feb 1981.

[LRS11] Lap Chi Lau, Ramamoorthi Ravi, and Mohit Singh. *Iterative methods in combinatorial optimization*, volume 46. Cambridge University Press, 2011.

[MN89] G. L. Miller and J. Naor. Flow in planar graphs with multiple sources and sinks. In *30th Annual Symposium on Foundations of Computer Science*, pages 112–117, Oct 1989.

[MS04] Marcin Mucha and Piotr Sankowski. Maximum matchings via Gaussian elimination. In *45th Symposium on Foundations of Computer Science (FOCS)*, pages 248–255, 2004.

[MV97] Meena Mahajan and V. Vinay. Determinant: Combinatorics, algorithms, and complexity. Technical report, 1997.

[MV00] Meena Mahajan and Kasturi R. Varadarajan. A new nc-algorithm for finding a perfect matching in bipartite planar and small genus graphs. In *Proceedings of the Thirty-second Annual ACM Symposium on Theory of Computing*, STOC '00, pages 351–357, New York, NY, USA, 2000. ACM.

[MVV87] Ketan Mulmuley, Umesh V. Vazirani, and Vijay V. Vazirani. Matching is as easy as matrix inversion. *Combinatorica*, 7(1):105–113, 1987.

[Nai82] M. Nair. On Chebyshev-type inequalities for primes. *The American Mathematical Monthly*, 89(2):126–129, 1982.

[NSV94] H. Narayanan, Huzur Saran, and Vijay V. Vazirani. Randomized parallel algorithms for matroid union and intersection, with applications to arboresences and edge-disjoint spanning trees. *SIAM J. Comput.*, 23(2):387–397, 1994.

[Par98] I. Parfenoff. An efficient parallel algorithm for maximum matching for some classes of graphs. *Journal of Parallel and Distributed Computing*, 52(1):96 – 108, 1998.

[PY82] Christos H. Papadimitriou and Mihalis Yannakakis. The complexity of restricted spanning tree problems. *J. ACM*, 29(2):285–309, April 1982.

[Sch03] Alexander Schrijver. *Combinatorial Optimization - Polyhedra and Efficiency.* Springer-Verlag, Berlin, 2003.

[Tut47] W. T. Tutte. The factorization of linear graphs. *Journal of the London Mathematical Society*, 22:107–111, 1947.

[TV12] Raghunath Tewari and N. V. Vinodchandran. Green's theorem and isolation in planar graphs. *Inf. Comput.*, 215:1–7, 2012.

[Vaz89] Vijay V. Vazirani. NC algorithms for computing the number of perfect matchings in $K_{3,3}$-free graphs and related problems. *Information and Computation*, 80(2):152 – 164, 1989.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il