

Smoothed Complexity of SWAP in Local Graph Partitioning

Xi Chen^{*}

Columbia University

xichen@cs.columbia.edu

Chenghao Guo[†]

MIT

chenghao@mit.edu

Emmanouil V. Vlatakis-Gkaragkounis[‡]

University California, Berkeley

emvlatakis@berkeley.edu

Mihalis Yannakakis[§]

Columbia University

mihalis@cs.columbia.edu

May 26, 2023

Abstract

We give the first quasipolynomial upper bound $\phi n^{\text{polylog}(n)}$ for the smoothed complexity of the SWAP algorithm for local Graph Partitioning (also known as Bisection Width), where n is the number of nodes in the graph and ϕ is a parameter that measures the magnitude of perturbations applied on its edge weights. More generally, we show that the same quasipolynomial upper bound holds for the smoothed complexity of the 2-FLIP algorithm for any binary Maximum Constraint Satisfaction Problem, including local Max-Cut, for which similar bounds were only known for 1-FLIP. Our results are based on an analysis of cycles formed in long sequences of double flips, showing that it is unlikely for every move in a long sequence to incur a positive but small improvement in the cut weight.

1 Introduction

Local search has been a powerful machinery for a plethora of problems in combinatorial optimization, from the classical Simplex algorithm for linear programming to the gradient descent method for modern machine learning problems, to effective heuristics (e.g. Kernighan-Lin) for basic combinatorial problems such as the Traveling Salesman Problem and Graph Partitioning. A local search algorithm begins with an initial candidate solution and then follows a path by iteratively moving to a better neighboring solution until a local optimum in its neighborhood is reached. The quality of the obtained solutions depends of course on how rich is the neighborhood structure that is explored by the algorithm. Local search is a popular approach to optimization because of the general applicability of the method and the fact that the algorithms typically run fast in practice. In contrast to their empirical fast convergence, however, many local search algorithms have exponential running time in the worst case due to delicate pathological instances that one may never encounter in practice. To bridge this striking discrepancy, Spielman and

^{*}Supported by NSF IIS-1838154, CCF-2106429 and CCF-2107187.

[†]Supported by NSF TRIPODS program award DMS-2022448 and by NSF Career Award CCF-1940205, CCF- 2131115.

[‡]Supported by Postdoctoral FODSI Simons-Fellowship.

[§]Supported by NSF CCF-2107187 and CCF-2212233.

Teng [1] proposed the framework of *smoothed analysis*, a hybrid of the classical worst-case and average-case analyses. They used it to provide rigorous justifications for the empirical performance of the Simplex algorithm by showing its smoothed complexity to be polynomial. Since then, the smoothed analysis of algorithms and problems from combinatorial optimization [2, 3, 4], among many other research areas such as numerical methods [5, 6, 4, 7], machine learning [8, 9, 10, 11] and algorithmic game theory [12, 13, 14, 15], has been studied extensively.

In this paper we study the smoothed complexity of local search algorithms for the classical problem of *Graph Partitioning* (also known as *Bisection Width* in the literature). In the problem we are given edge weights $X = (X_e : e \in E_{2n})$ of a complete graph $K_{2n} = (V_{2n}, E_{2n})$ with $X_e \in [-1, 1]$, and the goal is to find a *balanced* partition (U, V) of V_{2n} into two equal-size subsets U and V to minimize the weight of the corresponding cut (i.e., the sum of weights of edges with one node in U and the other node in V). Graph Partitioning has been studied extensively, especially in practice. It forms the basis of divide and conquer algorithms and is used in various application domains, for example in laying out circuits in VLSI. It has also served as a test bed for algorithmic ideas [16].

Given its NP-completeness [17], heuristics have been developed to solve Graph Partitioning in practice. A commonly used approach is based on local search: starting with an initial balanced partition, local improvements on the cut are made iteratively until a balanced partition that minimizes the cut within its *neighborhood* is reached. The simplest neighborhood is the SWAP neighborhood, where two balanced partitions are neighbors if one can be obtained from the other by swapping two nodes, one from each part. A locally optimal solution under the SWAP neighborhood can be found naturally by the SWAP algorithm, which keeps swapping two nodes as long as the swap improves the cut. A more sophisticated neighborhood structure, which yields much better locally optimal solutions in practice, is that of the Kernighan-Lin (KL) algorithm which performs in each move a sequence of swaps [18].

These local search algorithms for Graph Partitioning typically converge fast in practice. (For a thorough experimental analysis of their performance, and comparison with simulated annealing, regarding both the quality of solutions and the running time, see [16].) In contrast, it is also known that the worst-case complexity is exponential. (Finding a locally optimal solution for Graph Partitioning under the sophisticated Kernighan-Lin neighborhood, and even under the SWAP neighborhood is complete in PLS [19, 20]. The hardness reductions give instances on which these algorithms take exponential time to converge.) *This significant gap in our understanding motivates us to work on the smoothed complexity of the SWAP algorithm for Graph Partitioning in this paper.*

We work on the full perturbation model, under which edge weights are drawn independently from a collection of distributions $\mathcal{X} = (\mathcal{X}_e : e \in E_{2n})$. Each \mathcal{X}_e is supported on $[-1, 1]$, and has its density function bounded from above by a parameter $\phi > 0$. Our goal is to understand the expected number of steps the SWAP algorithm takes, as a function of n and ϕ , against any edge weight distributions \mathcal{X} .¹ Note that the SWAP algorithm, similar to the Simplex algorithm, is a family of algorithms since one can implement it using different pivoting rules, deterministic or randomized, to pick the next pair of nodes to swap when more than one pairs can improve the cut. We would like to establish upper bounds that hold for any implementation of the SWAP algorithm.

¹Note that any upper bound under the full perturbation model applies to the alternative, simpler model where an adversary commits to a weight w_e for each edge and then all edge weights are perturbed independently by a random noise \mathcal{Z}_e (for example, drawn uniformly from a small interval), i.e. the weights are $\mathcal{X}_e = w_e + \mathcal{Z}_e$. The parameter ϕ in the full perturbation model is a bound on the pdf of the perturbations \mathcal{Z}_e .

1.1 Related work: Smoothed analysis of 1-FLIP for Max-Cut

There has not been any previous analysis on SWAP under the smoothed setting, as far as we are aware. In contrast, much progress has been made on the smoothed analysis of the *1-FLIP algorithm for Max-Cut* [21, 22, 23, 24, 25]. The major challenge for the analysis of SWAP, as we discuss in more details in Section 1.3, is to overcome substantial new obstacles posed by the richer neighborhood structure of SWAP, which are not present in the simpler *1-change neighborhood* behind 1-FLIP.

Recall in Max-Cut, we are given edge weights $X = (X_e : e \in E_n)$ of a complete graph $K_n = (V_n, E_n)$ with $X_e \in [-1, 1]$ and the goal is to find a (*not necessarily balanced*) partition of V_n to maximize the cut.² The simplest neighborhood structure for local search on Max-Cut is the so-called *1-change neighborhood*, where two partitions are neighbors if one can be obtained from the other by moving a single node to the other side. The 1-FLIP algorithm finds such a locally optimal solution by keeping moving nodes, one by one, as long as each move improves the cut. For the structured perturbation model, where a graph G (not necessarily a complete graph) is given and only weights of edges in G are perturbed, [22] showed that the expected number of steps 1-FLIP takes to terminate is at most $\phi n^{\log n}$. Subsequently, the bound was improved by [23] to $\phi \cdot \text{poly}(n)$ for the full perturbation model, with further improvements in [24] on the polynomial part of n . The upper bound of [22] for the structured model was recently improved to $\phi n^{\sqrt{\log n}}$ in [25].

1.2 Our Contributions

We present the first smoothed analysis of the SWAP algorithm for Graph Partitioning. Our main result for SWAP is a quasipolynomial upper bound on its expected running time:³

Theorem 1.1. *Let $\mathcal{X} = (\mathcal{X}_e : e \in E_{2n})$ be distributions of edge weights such that each \mathcal{X}_e is supported on $[-1, 1]$ and has its density function bounded from above by a parameter $\phi > 0$. Then with probability at least $1 - o_n(1)$ over the draw of edge weights $X \sim \mathcal{X}$, any implementation of SWAP takes at most $\phi n^{O(\log^{10} n)}$ steps to terminate.*

The proof of Theorem 1.1 for SWAP is based on techniques we develop for a more challenging problem: the smoothed analysis of *2-FLIP for Max-Cut*. Starting with an initial partition (not necessarily balanced), in each round, 2-FLIP can move either one node (like 1-FLIP) or two nodes (not necessarily in different parts) as long as the cut is improved. If we restrict the algorithm to only use double flips in every move, then we call this variant *pure 2-FLIP*. Feasible moves in SWAP are clearly feasible in pure 2-FLIP as well but not vice versa. Thus, an improving sequence of SWAP in the Graph Partitioning problem is also an improving sequence of pure 2-FLIP in the Max-Cut problem on the same instance.

We do not make again any assumption on the pivoting rule used by 2-FLIP (i.e., which move is selected in each step if there are multiple improving moves), except that if both single and double flips are allowed, then the algorithm never moves a pair of nodes when moving only one of the two nodes would yield a better cut. Clearly, any reasonable implementation of 2-FLIP satisfies this property. Our main result on 2-FLIP is a similar quasipolynomial upper bound on its expected running time. The same result holds also for any implementation of the pure 2-FLIP algorithm that performs only 2-flips. This is the first smoothed analysis of 2-FLIP:

²Since we allow weights in $[-1, 1]$, maximizing the cut is the same as minimizing the cut after negating all edge weights. Hence the only difference of Max-Cut, from Graph Partitioning, is that the partition does not have to be balanced.

³We did not make an attempt to optimize the constant 10 in the polylog exponent.

Theorem 1.2. Let $\mathcal{X} = (\mathcal{X}_e : e \in E_n)$ be distributions of edge weights such that each \mathcal{X}_e is supported on $[-1, 1]$ and has its density function bounded from above by a parameter $\phi > 0$. Then with probability at least $1 - o_n(1)$ over the draw of edge weights $X \sim \mathcal{X}$, any implementation of the 2-FLIP algorithm takes at most $\phi n^{O(\log^{10} n)}$ steps to terminate.

A more general class of problems that is related to Max-Cut is the class of *Maximum Binary Constraint Satisfaction Problems* (MAX 2-CSP). In a general Max-2CSP, the input consists of a set of Boolean variables and a set of weighted binary constraints over the variables; the problem is to find an assignment to the variables that maximizes the weight of the satisfied constraints. Max-Cut is the special case when every constraint is a \neq (XOR) constraint. Other examples are Max 2SAT and Max Directed Cut (i.e., the Max Cut problem on weighted directed graphs). More generally, in a *Binary Function Optimization Problem* (BFOP), instead of binary constraints the input has a set of weighted binary functions on the variables, and the objective is to find an assignment that maximizes the sum of the weights of the functions (see Section 7 for the formal definitions). It was shown in [25] that the results for 1-FLIP for Max-Cut generalize to all Max 2-CSP and BFOP problems. We prove that this is the case also with 2-FLIP.

We say an instance of Max 2-CSP or BFOP is *complete* if it includes a constraint or function for every pair of variables.

Theorem 1.3. Let I be an arbitrary complete instance of a MAX 2-CSP (or BFOP) problem with n variables and m constraints (functions) with independent random weights in $[-1, 1]$ with density at most $\phi > 0$. Then, with probability at least $1 - o_n(1)$ over the draw of the weights, any implementation of 2-FLIP takes at most $m\phi n^{O(\log^{10} n)}$ steps to terminate.

For all the aforementioned problems, by controlling the tail-bound of the failure probability, we can strengthen our analysis to derive the same bound for the expected number of steps needed to terminate as in the standard smoothed analysis prototype (See Corollary 3.5).

1.3 Our Approach

Here, we give an overview of our proof approach, focusing on the analysis of the 2-FLIP algorithm for Max-Cut (Theorem 1.2). Many details are omitted in this subsection, to help the reader get an overall view of some of the key ideas and the structure of the proof. Note that 2-FLIP clearly subsumes 1-FLIP, since it explores a much larger neighborhood structure. For example, a 2-FLIP algorithm could apply improving 1-flips as long as possible, and only when the partition is locally optimal with respect to the 1-flip neighborhood apply an improving 2-flip. Therefore, the complexity (whether smoothed or worst-case) of 2-FLIP is clearly at least as large as the complexity of 1-FLIP, and could potentially be much larger. Similarly, the analysis of 2-FLIP has to subsume the analysis of 1-FLIP, but it needs to address many more challenges, in view of the larger space of possible moves in each step (quadratic versus linear). In a sense, it is analogous to the difference between a two-dimensional and a one-dimensional problem.

First, let's briefly review the approach of previous work [22] on the simpler 1-FLIP problem. Since the edge weights are in $[-1, 1]$, the weight of any cut is in $[-n^2, n^2]$. For the execution of the FLIP algorithm to be long, it must have many moves where the gain in the cut weight is very small, in $(0, \epsilon]$ for some small $\epsilon > 0$. It is easy to see that any single move by itself has small probability ($\phi\epsilon$) of this

being the case. If different moves were uncorrelated, then the probability that a sequence increases the weight of the cut by no more than ϵ would go down exponentially with the length of the sequence. Of course, different moves are correlated. However, the same effect holds if the improvements of the moves are linearly independent in the following sense. For any sequence of the FLIP algorithm, the *improvement vector* of one move is the vector indexed by the edges with entries in $\{-1, 0, 1\}$ indicating whether each edge is added or removed from the cut as a result of the move. Most work along this line of research is based on the following fact (see Corollary 2.1 for the formal statement): If the rank of the set of improvement vectors is rank , then the sequence has improvement at most ϵ with probability at most $(\phi\epsilon)^{\text{rank}}$. On the other hand, if all sequences with length at most $\Theta(n)$ have an improvement of at least ϵ , then the number of steps of FLIP is bounded by $\Theta(n) \cdot (2n^2/\epsilon) = \text{poly}(n)/\epsilon$, as the total improvement cannot exceed $2n^2$. So a natural approach is to union bound over all possible sequences of length $\Theta(n)$ and all 2^n possible initial configurations, which yields a probability upper bound of $2^n n^{\Theta(n)} (\phi\epsilon)^{\text{rank}}$.

Getting a quasi-polynomial complexity bound using the union bound above requires the rank of any sequence of length $\Theta(n)$ to be at least $\Omega(n/\log n)$. However, this is not always true (consider, e.g., a sequence in which only $n^{0.1}$ distinct nodes moved). One key idea of [22] is to avoid union bound over all initial configurations and only union bound over initial configurations of *active* nodes (nodes that move at least once in the sequence) by looking at *arcs*. An arc is defined to be two adjacent moves of the same node. By taking the sum of improvement vectors of the two moves of an arc, edges that involve inactive nodes are cancelled, so the union bound over sequences of length ℓ becomes $2^\ell n^\ell (\phi\epsilon)^{\text{rank}_{\text{arcs}}}$.

To lower bound the rank of arcs of sequence S , $\text{rank}_{\text{arcs}}(S)$, they proved it is at least half of the number of nodes that appear more than once in the sequence, denoted $V_2(S)$. The essential combinatorial claim made by [22] is that for any sequence of length $\Omega(n)$, there exists a substring of length ℓ with $V_2(S)$ at least $\Omega(\ell/\log n)$. This can be shown by bucketing arcs by length into buckets $[2^i, 2^{i+1})$ and picking the largest bucket as length of the substring. On average, a random substring would contain $\Omega(\ell/\log n)$ arcs with similar length, and therefore, $\Omega(\ell/\log n)$ arcs with distinct nodes. The similar idea is used in Case 1 of our Section 6 to handle 1-moves (moves that flip a single node).

Now let's return to the case of the 2-FLIP algorithm. A step now can move two nodes at the same time, and this fact poses qualitatively new challenges to the proof framework. Now we have to deal not just with sets (e.g., the set of nodes that move more than once) but instead with relations (graphs). Define an *auxiliary graph* H for the sequence of moves that contains K_n as vertices and an edge for each 2-move of the sequence. If we still want to eliminate the influence of inactive nodes in the improvement vector by summing or subtracting two moves as in the 1-FLIP case, the moves have to contain the *exact* same pair of nodes. This happens too sparsely in the improving sequence of 2-FLIP to provide enough rank. To this end, we generalize the notion of arcs to *cycles*. A cycle is a set of 2-moves of the sequence whose corresponding edges form a cycle in H . But not all cycles of H are useful. We are interested only in cycles for which there is a linear combination of the improvement vectors of the moves of the cycle that cancels all edges of K_n that involve an inactive node (i.e., the corresponding entry in the linear combination is 0); these are the cycles that are useful to the rank and we call them *dependent cycles*.

So the goal is to find a substring S of length ℓ where we can lower bound $\text{rank}_{\text{cycles}}(S)$ by $\ell/\text{polylog}(n)$. The ideal case would be the case where all nodes have $O(\text{polylog}(n))$ but at least 2 appearances in the substring, i.e., all nodes have degree between 2 and $O(\text{polylog}(n))$ in H . In this case, we can repeat

the following process to find enough cycles. Find a dependent cycle in H , pick an edge in K_n that is non-zero in the improvement vector of the cycle (we call this the *witness* of the cycle) and delete both nodes of the witness from H . This way the improvement vector of cycles of H we pick in the future will not contain witnesses from previous cycles, and improvement vectors of cycles we pick would form a triangular matrix that has full rank. Since any node has $O(\text{polylog}(n))$ degree in H , each iteration deletes $O(\text{polylog}(n))$ edges. So the process can be repeated at least $\Omega(\ell/\text{polylog}(n))$ times.

However, it is not hard to construct sequences with polynomial length, such that any substring consists mostly of moves involving one high-degree node (with degree even $\Omega(\ell)$) and one degree-1 node, so deleting the high-degree node would have a significant impact on the graph and the process can only repeat for a few rounds and lead to a few cycles. So the challenge is to run a similar process, but reuse high-degree nodes carefully without repeating witnesses found in previous cycles. Suppose we find a cycle C with witness edge (u, v) . To avoid including the edge in another cycle C' , a sufficient condition is that: (1) u is not included in C' . (2) For any two adjacent edges (edges in H , not K_n) of v in C' , u never moved between the two corresponding moves of the edges. To meet condition 1, we can delete u from H . To meet condition 2, we can make multiple copies of v in H where each copy corresponds to moves in S where u doesn't move between them. We call this operation *splitting* since the new graph is generated by deleting and splitting the original H . The new graph after splitting is denoted by *splitted auxiliary graph*. Our algorithm for finding a large number of linearly independent cycles can be described as repeatedly performing the following process. Find a cycle in the splitted auxiliary graph with witness (u, v) by a tree-growing argument, delete u and split the graph by creating multiple copies of v . We have to choose carefully the witness edges (u, v) and do the splitting, so that the number of nodes does not proliferate in this process.

Compared to the original auxiliary graph, the number of edges deleted and the number of new nodes introduced is proportional to the degree of u , so the number of cycles for a sequence of length ℓ we can find in the algorithm is bounded by $\ell/\deg(u)$. To find $\ell/\text{poly}(\log n)$ cycles, we need a window where decent amount of moves involve a node u that has $\deg(u)$ bounded by $\text{poly}(\log n)$. The existence of such window in an arbitrary sequence that is long enough can be proven via a sophisticated bucketing and counting argument.

The overall argument then for 2-FLIP is that, given a sufficiently long sequence of improving moves (specifically, of length $n \cdot \text{poly}(\log n)$), we can find a window (a substring) such that the rank of the arcs and cycles in the window is within a $\text{poly}(\log n)$ factor of the length of the window. As a consequence, with high probability the weight of the cut improves by a nontrivial amount ϵ (1/quasi-polynomial) during this window. This can happen at most n^2/ϵ times, hence the length of the execution sequence of 2-FLIP is at most quasi-polynomial.

1.4 Organization of the paper.

The rest of the paper is organized as follows. Section 2 gives basic definitions of the problems and the smoothed model, defines the central concepts of arcs and cycles, their improvement vectors, and proves a set of basic lemmas about them that are used throughout in the subsequent analysis. Section 3 states the main lemma on the existence of a nice window in the move sequence such that the arcs and cycles in the window have high rank, and shows how to derive the main theorem from this lemma. Sections

[4](#) and [5](#) prove the main lemma in the case that all the moves are 2-moves (this is the more challenging case). First we show in Section [4](#) the existence of a nice window (in fact a large number of nice windows, since this is needed in the general case) such that many moves in the window have the property that both nodes of the move appear a substantial number of times in the window (at least $\text{polylog}(n)$ times), and one of them does not appear too many times (at most a higher $\text{polylog}(n)$). In Section [5](#) we show how to find in such a nice window a large number of cycles whose improvement vectors are linearly independent. Section [6](#) extends the proof of the main lemma to the general case where the sequence of moves generated by 2-FLIP contains both 1- and 2-moves. Finally, in Section [7](#) we extend the results to the class of Maximum Binary Constraint Satisfaction and Function Optimization problems.

2 Preliminaries

We write $[n]$ to denote $\{1, \dots, n\}$. Given two integers $a \leq b$, we write $[a : b]$ to denote $\{a, \dots, b\}$. Given $\gamma, \gamma' \in \{\pm 1\}^n$ we use $d(\gamma, \gamma')$ to denote the Hamming distance between γ and γ' , i.e., the number of entries $i \in [n]$ such that $\gamma_i \neq \gamma'_i$.

2.1 Local Max-Cut and the FLIP Algorithm

Let $K_n = (V_n, E_n)$ with $V_n = [n]$ be the complete undirected graph over n nodes. Given edge weights $X = (X_e : e \in E_n)$ with $X_e \in [-1, 1]$, the k -local Max-Cut problem is to find a partition of V_n into two sets V_1 and V_2 such that the weight of the corresponding cut (the sum of weights of edges with one node in V_1 and the other in V_2) cannot be improved by moving no more than k nodes to the other set. Formally, the objective function of our interest is defined as follows: Given any configuration $\gamma \in \{\pm 1\}^n$ (which corresponds to a partition V_1, V_2 with $V_1 = \{u \in V_n : \gamma(u) = -1\}$ and $V_2 = \{u \in V_n : \gamma(u) = 1\}$), the objective function is

$$\text{obj}_X(\gamma) := \sum_{(u,v) \in E_n} X_{(u,v)} \cdot \mathbf{1}\{\gamma(u) \neq \gamma(v)\} = \frac{1}{2} \sum_{(u,v) \in E_n} X_{(u,v)} \cdot (1 - \gamma(u)\gamma(v)). \quad (1)$$

Our goal is to find a configuration $\gamma \in \{\pm 1\}^n$ that is a k -local optimum, i.e., $\text{obj}_X(\gamma) \geq \text{obj}_X(\gamma')$ for every configuration $\gamma' \in \{\pm 1\}^n$ with Hamming distance no more than k from γ .

A simple local search algorithm for k -Local Max-Cut is the following k -FLIP algorithm:

Start with some initial configuration $\gamma = \gamma_0 \in \{\pm 1\}^n$. While there exists a configuration γ' with $d(\gamma', \gamma) \leq k$ such that $\text{obj}_X(\gamma') > \text{obj}_X(\gamma)$, select one such configuration γ' (according to some pivoting criterion), set $\gamma = \gamma'$ and repeat, until no such configuration γ' exists.

The execution of k -FLIP on K_n with edge weights X depends on both the initial configuration γ_0 and the pivoting criterion used to select the next configuration in each iteration. The larger the value of k , the larger the neighborhood structure that is being explored, hence the better the quality of solutions that is expected to be generated. However, the time complexity of each iteration grows rapidly with k : there are $\Theta(n^k)$ candidate moves, and with suitable data structures we can determine in $O(n^k)$ if there is an improving move and select one. Thus, the algorithm is feasible only for small values of k . For $k = 1$, it is the standard FLIP algorithm. Here we are interested in the case $k = 2$. We will not make any assumption on the pivoting criterion in our results, except that we assume that the algorithm

does not choose to flip in any step two nodes when flipping only one of them would produce a strictly better cut. This is a natural property satisfied by any reasonable implementation of 2-FLIP. For example, one approach (to optimize the time of each iteration) is to first check if there is an improving 1-flip (n possibilities), and only if there is none, proceed to search for an improving 2-flip ($O(n^2)$ possibilities). Clearly any implementation that follows this approach satisfies the above property. Also, the greedy approach, that examines all $O(n^2)$ possible 1-flips and 2-flips and chooses one that yields the maximum improvement, obviously satisfies the above property.

Our results hold also for the variant of 2-FLIP that uses only 2-flips (no 1-flips). We refer to this variant as *Pure 2-FLIP*.

2.2 Graph Partitioning and the SWAP Algorithm

In the Graph Partitioning (or Bisection Width) problem, we are given a graph G on $2n$ nodes with weighted edges; the problem is to find a partition of the set V of nodes into two equal-sized subsets V_1, V_2 to minimize the weight of the cut.⁴ As in the Max Cut problem, in this paper we will assume the graph is complete and the edge weights are in $[-1, 1]$. A simple local search algorithm is the SWAP algorithm: Starting from some initial partition (V_1, V_2) with n nodes in each part, while there is a pair of nodes $u \in V_1, v \in V_2$ whose swap (moving to the other part) decreases the weight of the cut, swap u and v . We do not make any assumption on the pivoting rule, i.e. which pair is selected to swap in each iteration if there are multiple pairs whose swap improves the cut. At the end, when the algorithm terminates it produces a locally optimal balanced partition, i.e. one that cannot be improved by swapping any pair of nodes. The SWAP algorithm is clearly a restricted version of Pure 2-FLIP (restricted because the initial partition is balanced, and in each step the 2-flip must involve two nodes from different parts of the partition).

The SWAP algorithm is the simplest local search algorithm for the Graph Partitioning problem, but it is a rather weak one, in the sense that the quality of the locally optimal solutions produced may not be very good. For this reason, more sophisticated local search algorithms have been proposed and are typically used, most notably the Kernighan-Lin algorithm [18], in which a move from a partition to a neighboring partition involves a sequence of swaps. If a partition has a profitable swap, then Kernighan-Lin (KL) will perform the best swap; however, if there is no profitable swap then KL explores a sequence of n greedy steps, selecting greedily in each step the best pair of nodes to swap that have not changed sides before in the current sequence, and if this sequence of swaps produces eventually a better partition, then KL moves to the best such partition generated during this sequence. A related variant, to reduce the time cost of each iteration, was proposed by Fiduccia and Matheyses [26]. This idea of guided deep neighborhood search is a powerful method in local search that was introduced first in the [18] paper of Kernighan and Lin on Graph Partitioning, and was applied subsequently successfully to the Traveling Salesman Problem and other problems.

⁴Since the weights can be positive or negative, there is no difference between maximization and minimization. The Graph Partitioning problem is usually stated as a minimization problem.

2.3 Smoothed Analysis

We focus on the 2-FLIP algorithm from now on. Under the smoothed complexity model, there is a family $\mathcal{X} = (\mathcal{X}_e : e \in E_n)$ of probability distributions, one for each edge in $K_n = (V_n, E_n)$. The edge weights $X = (X_e : e \in E_n)$ are drawn independently with $X_e \sim \mathcal{X}_e$. We assume that each \mathcal{X}_e is a distribution supported on $[-1, 1]$ and its density function is bounded from above by a parameter $\phi > 0$. (The assumption that the edge weights are in $[-1, 1]$ is no loss of generality, since they can be always scaled to lie in that range.) Our goal is to bound the number of steps the 2-FLIP algorithm takes to terminate when running on K_n with edge weights $X \sim \mathcal{X}$, in terms of n and the parameter ϕ .

2.4 Move Sequences

We introduce some of the key definitions that will be used in the smoothed analysis of 2-FLIP.

A *move sequence* $\mathcal{S} = (\mathcal{S}_1, \dots, \mathcal{S}_\ell)$ is an ℓ -tuple for some $\ell \geq 1$ such that \mathcal{S}_i is a subset of V_n of size either one or two. We will refer to the i -th move in \mathcal{S} as a *1-move* if $|\mathcal{S}_i| = 1$ and a *2-move* if $|\mathcal{S}_i| = 2$, and write $\text{len}(\mathcal{S}) := \ell$ to denote its length. Additionally, let $1\text{-move}(\mathcal{S})$ and $2\text{-move}(\mathcal{S})$ denote the corresponding subsequence of single flip or double flips correspondingly. We say a node $u \in V_n$ is *active* in \mathcal{S} if u appears in \mathcal{S}_i for some i , and is *inactive* otherwise. We write $V(\mathcal{S}) \subseteq V_n$ to denote the set of active nodes in \mathcal{S} .

Given $\gamma_0 \in \{\pm 1\}^n$ as the initial configuration, a move sequence $\mathcal{S} = (\mathcal{S}_1, \dots, \mathcal{S}_\ell)$ naturally induces a sequence of configurations $\gamma_0, \gamma_1, \dots, \gamma_\ell \in \{\pm 1\}^n$, where γ_{i+1} is obtained from γ_i by flipping the nodes in \mathcal{S}_{i+1} . We say (γ_0, \mathcal{S}) is *improving* with respect to edge weights X if

$$\text{obj}_X(\gamma_i) > \text{obj}_X(\gamma_{i-1}), \quad \text{for all } i \in [\ell]$$

and is ϵ -*improving* with respect to edge weights X , for some $\epsilon > 0$, if

$$\text{obj}_X(\gamma_i) - \text{obj}_X(\gamma_{i-1}) \in (0, \epsilon], \quad \text{for all } i \in [\ell].$$

For each $i \in [\ell]$, the change $\text{obj}_X(\gamma_i) - \text{obj}_X(\gamma_{i-1})$ from the i -th move \mathcal{S}_i can be written as follows:

1. When $\mathcal{S}_i = \{u\}$,

$$\text{obj}_X(\gamma_i) - \text{obj}_X(\gamma_{i-1}) = \sum_{w \in V_n : w \neq u} \gamma_{i-1}(w) \gamma_{i-1}(u) X_{(u,w)}. \quad (2)$$

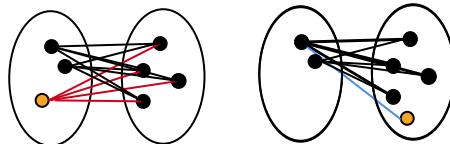


Figure 1: Example of a *1-move*, showing edges in the cut only.

2. When $\mathcal{S}_i = \{u, v\}$,

$$\text{obj}_X(\gamma_i) - \text{obj}_X(\gamma_{i-1}) = \sum_{w \in V_n : w \notin \{u, v\}} (\gamma_{i-1}(w) \gamma_{i-1}(u) X_{(w,u)} + \gamma_{i-1}(w) \gamma_{i-1}(v) X_{(w,v)}). \quad (3)$$

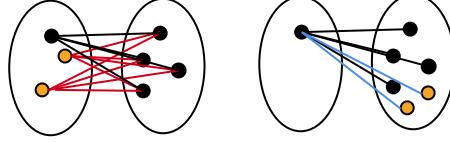


Figure 2: Example of a 2-move, showing edges in the cut only.

For each $i \in [\ell]$, we write $\text{imprv}_{\gamma_0, \mathcal{S}}(i)$ to denote the improvement vector in $\{0, \pm 1\}^{E_n}$ such that

$$\text{obj}_X(\gamma_i) - \text{obj}_X(\gamma_{i-1}) = \text{imprv}_{\gamma_0, \mathcal{S}}(i) \cdot X. \quad (4)$$

Next, let $E(\mathcal{S})$ denote the set of edges $(u, v) \in E_n$ such that both u and v are active in \mathcal{S} . We write $\text{imprv}_{\gamma_0, \mathcal{S}}^*(i) \in \{0, \pm 1\}^{E(\mathcal{S})}$ to denote the projection of $\text{imprv}_{\gamma_0, \mathcal{S}}(i)$ on entries that correspond to edges in $E(\mathcal{S})$. We note that $\text{imprv}_{\gamma_0, \mathcal{S}}^*(i)$ only depends on the initial configuration of active nodes $V(S)$ in γ_0 . Given a (partial) configuration $\tau_0 \in \{\pm 1\}^{V(S)}$ of $V(S)$, we let

$$\text{imprv}_{\tau_0, \mathcal{S}}(i) := \text{imprv}_{\gamma_0, \mathcal{S}}^*(i) \in \{0, \pm 1\}^{E(\mathcal{S})},$$

where $\gamma_0 \in \{\pm 1\}^n$ is an arbitrary (full) configuration that is an extension of τ_0 . (To aid the reader we will always use γ to denote a full configuration and τ to denote a partial configuration in the paper.)

Note that if \mathcal{S} is a sequence of moves generated by an execution of the 2-FLIP algorithm then \mathcal{S} must be improving, because every move must increase the weight of the cut and therefore every 1– or 2– move is improving. On the other hand, if every move in \mathcal{S} increases the cut weight by no more than ϵ then we can not directly guarantee that after $\text{poly}(|\mathcal{S}|, n, 1/\epsilon)$ steps the algorithm would certainly terminate. From probabilistic perspective, in order to provide a smoothed upper bound on the running time of 2-FLIP method, it suffices to show that it is exponentially small probability for every move in a long enough sequence to incur only a $o(1/\text{poly}(n))$ improvement in our objective.

Indeed, in an idealized scenario where the improvements of different moves of a sequence were disentangled, the event for a linear-length sequence to be at most ϵ –improving would have exponentially small probability. Unfortunately, going back to the 2-FLIP algorithm, there could be improving steps that are strongly correlated (as an extreme situation there could be two flips with almost the same improvement vector). Thus, as one may expect the probability exponential decay holds only for linearly independent $\text{imprv}_{\tau_0, \mathcal{S}}(\cdot)$, introducing the necessity of analysis of the rank $(\{\text{imprv}_{\tau_0, \mathcal{S}}(i) | i \in \mathcal{S}'\})$, for some neatly chosen subset \mathcal{S}' of moves from the sequence \mathcal{S} .

Corollary 2.1 ([22]). *Let X_1, \dots, X_m be independent real random variables and let $f_i : \mathbb{R} \rightarrow [0, \phi]$ for some $\phi > 0$ denote the density of X_i for each $i \in [m]$. Additionally, let \mathcal{C} be a collection of k not necessarily linearly independent integer row vectors, namely $\mathcal{C} = \{V_1, \dots, V_k\}$. Then it holds that for any interval $I \subset \mathbb{R}$*

$$\Pr[F_\epsilon] = \Pr \left[\bigcap_{i \in [k]} \{V_i \cdot X \in I\} \right] \leq (\phi \text{len}(I))^{r_{\text{rank}}(\mathcal{C})}$$

However, one standard issue, which typically occurs with the direct usage of improvement vectors of sequence's moves, is their dependence also on the initial configuration γ of inactive nodes that do not appear in the sequence \mathcal{S} . Their number may be much larger than the rank of the active nodes, and thus

considering all their possible initial values in a union-bound will overwhelm the probability $(\phi\epsilon)^r$. For these reasons, in the literature [22, 24, 25] more complex combinatorial structures have been proposed, like pairs of (consecutive) moves of the same node. Interestingly, for the case of 2-FLIP, new challenges have to be overcome due to the 2-move case. To alleviate these harnesses, we introduce the idea of dependent cycles whose role will be revealed in the case that our sequence abounds with 2-moves.

2.5 Arcs and Cycles

Definition 2.2. An arc α in a move sequence $\mathcal{S} = (\mathcal{S}_1, \dots, \mathcal{S}_\ell)$ is an ordered pair (i, j) with $i < j \in [\ell]$ such that $\mathcal{S}_i = \mathcal{S}_j = \{u\}$ for some node $u \in V_n$ and for any $i < k < j$, $\mathcal{S}_k \neq \{u\}$.

Let $\tau_0 \in \{\pm 1\}^{V(\mathcal{S})}$ be a configuration of active nodes in \mathcal{S} , and let $\tau_0, \tau_1, \dots, \tau_\ell \in \{\pm 1\}^{V(\mathcal{S})}$ be the sequence of configurations induced by \mathcal{S} , i.e., τ_i is obtained from τ_{i-1} by flipping nodes in \mathcal{S}_i . We make the following observation:

Lemma 2.3. For any configuration $\gamma_0 \in \{\pm 1\}^n$ that is an extension of τ_0 , letting $\gamma_0, \gamma_1, \dots, \gamma_\ell \in \{\pm 1\}^n$ be the sequence of configurations induced by \mathcal{S} and letting $w[u, i, j] := \gamma_i(u) \cdot \text{imprv}_{\gamma_0, \mathcal{S}}(i) - \gamma_j(u) \cdot \text{imprv}_{\gamma_0, \mathcal{S}}(j)$, we have that

$$(w[u, i, j]_e)_{e \in E_n} = \begin{cases} (\tau_i(u) \cdot \text{imprv}_{\tau_0, \mathcal{S}}(i) - \tau_j(u) \cdot \text{imprv}_{\tau_0, \mathcal{S}}(j))_e & \text{for every entry } e \in E(\mathcal{S}), \\ 0 & \text{otherwise.} \end{cases}$$

for any arbitrary choice of $u \in V(\mathcal{S})$.

Motivated by Lemma 2.3, we define for an arc $\alpha = (i, j)$ of a node u ,

$$\text{imprv}_{\tau_0, \mathcal{S}}(\alpha) := \tau_i(u) \cdot \text{imprv}_{\tau_0, \mathcal{S}}(i) - \tau_j(u) \cdot \text{imprv}_{\tau_0, \mathcal{S}}(j) \in \mathbb{Z}^{E(\mathcal{S})}. \quad (5)$$

Let $\text{arcs}(\mathcal{S})$ denote the set of all arcs in \mathcal{S} . We will be interested in the rank of

$$Q_{\text{arcs}} := \{\text{imprv}_{\tau_0, \mathcal{S}}(\alpha) : \alpha \in \text{arcs}(\mathcal{S})\} \quad (6)$$

It is easy to show that the rank does not depend on the choice of τ_0 so we will denote it by $\text{rank}_{\text{arcs}}(\mathcal{S})$.

Lemma 2.4. The rank of the set of vectors in (6) does not depend on the choice of τ_0 .

Definition 2.5. A cycle C in a move sequence $\mathcal{S} = (\mathcal{S}_1, \dots, \mathcal{S}_\ell)$ is an ordered tuple $C = (c_1, \dots, c_t)$ for some $t \geq 2$ such that c_1, \dots, c_t are distinct, and $\mathcal{S}_{c_j} = \{u_j, u_{j+1}\}$ for all $j \in [t-1]$ and $\mathcal{S}_{c_t} = \{u_t, u_1\}$ for some nodes $u_1, \dots, u_t \in V_n$. (Every \mathcal{S}_{c_j} is a 2-move. The same vertex may appear in multiple \mathcal{S}_{c_j} 's).

Definition 2.6. Given a configuration $\tau_0 \in \{\pm 1\}^{V(\mathcal{S})}$, we say a cycle $C = (c_1, \dots, c_t)$ in \mathcal{S} is dependent with respect to τ_0 if there exists $b \in \{\pm 1\}^t$ such that

For all $j \in [t-1]$ we have that $b_j \cdot \tau_{c_j}(u_{j+1}) + b_{j+1} \cdot \tau_{c_{j+1}}(u_{j+1}) = 0$ and $b_t \cdot \tau_{c_t}(u_1) + b_1 \cdot \tau_{c_1}(u_1) = 0$,

where $\tau_0, \tau_1, \dots, \tau_\ell$ are configurations induced by \mathcal{S} starting from τ_0 .

We note that such a vector b , if it exists, it has the form $b = b_1 \cdot (1, \dots, (-1)^{k-1} \prod_{i \in [2:k]} \tau_{c_{i-1}}(u_i) \tau_{c_i}(u_i), \dots)^\top$ and hence it is unique if we further require $b_1 = 1$. After elimination of the above equations, we see the following equivalent criterion:

Remark 2.7 (Dependence Criterion). *A cycle C is dependent $\Leftrightarrow (-1)^t = \tau_{c_t}(u_1) \tau_{c_1}(u_1) \cdot \prod_{i=2}^t \tau_{c_{i-1}}(u_i) \tau_{c_i}(u_i)$.*

We will refer to the unique vector $b \in \{\pm 1\}^t$ as the *cancellation vector* of C . Notice that whether a cycle C in \mathcal{S} is dependent or not actually does not depend on the choice of τ_0 .

Lemma 2.8. *If a cycle C of \mathcal{S} is dependent with respect to some $\tau_0 \in \{\pm 1\}^{V(\mathcal{S})}$ using b as its cancellation vector, then it is dependent with respect to every configuration $\tau'_0 \in \{\pm 1\}^{V(\mathcal{S})}$ using the same b as its cancellation vector.*

As a result, we can refer to cycles of \mathcal{S} as dependent cycles without specifying a configuration τ_0 ; the same holds for cancellation vectors. Next we prove a lemma that is similar to Lemma 2.3 for arcs:

Lemma 2.9. *Let $C = (c_1, \dots, c_t)$ be a dependent cycle of \mathcal{S} and let b be its cancellation vector. Then for any configurations $\tau_0 \in \{\pm 1\}^{V(\mathcal{S})}$ and $\gamma_0 \in \{\pm 1\}^n$ such that γ_0 is an extension of τ_0 , letting $w[C] := \sum_{j \in [t]} b_j \cdot \text{imprv}_{\gamma_0, \mathcal{S}}(c_j)$, we have that*

$$(w[C]_e)_{e \in E_n} = \begin{cases} (\sum_{j \in [t]} b_j \cdot \text{imprv}_{\tau_0, \mathcal{S}}(c_j))_e & \text{for every entry } e \in E(\mathcal{S}), \\ 0 & \text{otherwise.} \end{cases}$$

Given $\tau_0 \in \{\pm 1\}^{V(\mathcal{S})}$ and a dependent cycle C of \mathcal{S} with b as its cancellation vector, we define

$$\text{imprv}_{\tau_0, \mathcal{S}}(C) := \sum_{j \in [t]} b_j \cdot \text{imprv}_{\tau_0, \mathcal{S}}(c_j). \quad (7)$$

Let $\text{cycles}(\mathcal{S})$ denote the set of all dependent cycles in \mathcal{S} . We will be interested in the rank of

$$Q_{\text{cycles}} := \{\text{imprv}_{\tau_0, \mathcal{S}}(C) : C \in \text{cycles}(\mathcal{S})\} \quad (8)$$

Similarly we note that the rank does not depend on the choice of τ_0 so we denote it by $\text{rank}_{\text{cycles}}(\mathcal{S})$.

Lemma 2.10. *The rank of the set of vectors in (8) does not depend on the choice of τ_0 .*

For the sake of readability we defer the proofs of initial configuration invariance for the rank of improvement vectors of arcs and cycles to Appendix B. Having defined the sets of $\text{arcs}(\mathcal{S})$ and $\text{cycles}(\mathcal{S})$, we conclude this section by showing that for a fixed parameter $\epsilon > 0$, a move sequence \mathcal{S} and an initial configuration $\tau_0 \in \{\pm 1\}^{V(\mathcal{S})}$, if either $\text{rank}_{\text{arcs}}(\mathcal{S})$ or $\text{rank}_{\text{cycles}}(\mathcal{S})$ is high, then most likely (over $X \sim \mathcal{X}$) (γ_0, \mathcal{S}) is not ϵ -improving for every $\gamma_0 \in \{\pm 1\}^n$ that is an extension of τ_0 .

Lemma 2.11. *Let $\epsilon > 0$. With probability at least*

$$1 - (2\text{len}(\mathcal{S}) \cdot \phi\epsilon)^{\max(\text{rank}_{\text{arcs}}(\mathcal{S}), \text{rank}_{\text{cycles}}(\mathcal{S}))}$$

over $X \sim \mathcal{X}$, we have that (γ_0, \mathcal{S}) is not ϵ -improving for every $\gamma_0 \in \{\pm 1\}^n$ that is an extension of τ_0 .

Proof. Let $\mathcal{E}_{\text{moves}}$ be the event of a given (γ_0, \mathcal{S}) being ϵ -improving with respect to edge weights X , for some fixed $\epsilon > 0$:

$$\mathcal{E}_{\text{moves}} : \left\{ \text{imprv}_{\gamma_0, \mathcal{S}}(i) \cdot X \in (0, \epsilon], \quad \text{for all } i \in [\ell] \right\}$$

where $\text{imprv}_{\gamma_0, \mathcal{S}}(i)$ correspond to the improvement vector of \mathcal{S}_i move (See (2), (3)). Now, notice that the improvement vector of an arc (See (5)) or of a dependent cycle (See (7)) can be written as the $\{-1, 0, 1\}$ sum of all the improvement vectors of either 1 or 2-moves in \mathcal{S} . Thus, we define the corresponding event for cycles and arcs for a given sequence (γ_0, \mathcal{S}) with respect to edge weights X :

$$\mathcal{E}_{\text{arcs/cycles}} : \left\{ \text{imprv}_{\gamma_0, \mathcal{S}}(\beta) \cdot X \in [-\text{len}(\mathcal{S})\epsilon, \text{len}(\mathcal{S})\epsilon], \quad \text{for any } \beta \in \text{arcs}(\mathcal{S})/\text{cycles}(\mathcal{S}) \right\}$$

So it is easy to see that $\mathcal{E}_{\text{moves}}$ implies $\mathcal{E}_{\text{arcs/cycles}}$, or equivalently $\Pr[\mathcal{E}_{\text{moves}}] \leq \min\{\Pr[\mathcal{E}_{\text{arcs}}], \Pr[\mathcal{E}_{\text{cycles}}]\}$. Thus, by leveraging Corollary 2.1 for vectors in Q_{arcs} and Q_{cycles} , we get that:

$$\Pr[(\gamma_0, \mathcal{S}) \text{ being an } \epsilon\text{-improving sequence}] \leq (2\text{len}(\mathcal{S}) \cdot \phi\epsilon)^{\max(\text{rank}_{\text{arcs}}(\mathcal{S}), \text{rank}_{\text{cycles}}(\mathcal{S}))}$$

This finishes the proof of the lemma. \square

3 Main Lemma and the Proof of Theorem 1.2 and Theorem 1.1

We start with the definition of *valid* move sequences:

Definition 3.1. We say a move sequence $\mathcal{S} = (\mathcal{S}_1, \dots, \mathcal{S}_\ell)$ is *valid* if it satisfies the following property: For every $i < j \in [\ell]$, at least one node $w \notin \mathcal{S}_i$ appears an odd number of times in $\mathcal{S}_i, \dots, \mathcal{S}_j$.

Lemma 3.2. The move sequence generated by 2-FLIP (or by pure 2-FLIP), for any pivoting rule and any instance, is *valid*.

Proof. Let \mathcal{S} be a move sequence generated by 2-FLIP (or pure 2-FLIP). If there are two moves $\mathcal{S}_i, \mathcal{S}_j$, $i \leq j$, such that no node appears an odd number of times in $\mathcal{S}_i, \dots, \mathcal{S}_j$, then the configurations before \mathcal{S}_i and after \mathcal{S}_j are the same, contradicting the fact that all the moves increase the weight of the cut. Therefore, the set O of nodes that appear an odd number of times in $\mathcal{S}_i, \dots, \mathcal{S}_j$ is nonempty. Suppose that $i < j$ and $O \subseteq \mathcal{S}_i$. If $O = \mathcal{S}_i$, then the set of nodes that appear an odd number of times in $\mathcal{S}_{i+1}, \dots, \mathcal{S}_j$ would be empty, a contradiction to the above property. Therefore, $O \neq \mathcal{S}_i$.

In the case of pure 2-FLIP, since all moves are 2-flips, O has even size, and hence $O \neq \emptyset$ and $O \neq \mathcal{S}_i$ imply the claim.

In the case of 2-FLIP, $O \neq \emptyset$, $O \neq \mathcal{S}_i$ and $O \subseteq \mathcal{S}_i$ imply that \mathcal{S}_i has size 2, say $\mathcal{S}_i = \{u, v\}$ and $O = \{u\}$ or $O = \{v\}$. If $O = \{u\}$ then the configuration γ_j differs from γ_{i-1} only in that node u is flipped. Thus, at configuration γ_{i-1} , flipping node u results in configuration γ_j which has strictly greater cut than the configuration γ_i that results by flipping the pair $\{u, v\}$, contradicting our assumption about 2-FLIP. A similar argument holds if $O = \{v\}$. In either case we have a contradiction to $O \subseteq \mathcal{S}_i$. The claim follows. \square

Given a move sequence \mathcal{S} , a *window* W of \mathcal{S} is a substring of \mathcal{S} , i.e., $W = (\mathcal{S}_i, \dots, \mathcal{S}_j)$ for some $i < j \in [\ell]$ (so W itself is also a move sequence). Our main technical lemma below shows that every

long enough valid move sequence has a window W such that either $\text{rank}_{\text{arcs}}(W)$ or $\text{rank}_{\text{cycles}}(W)$ is large relative to $\text{len}(W)$.

Lemma 3.3. *Let \mathcal{S} be a valid move sequence with $\text{len}(\mathcal{S}) \geq n \log^{10} n$. Then \mathcal{S} has a window W such that*

$$\max(\text{rank}_{\text{arcs}}(W), \text{rank}_{\text{cycles}}(W)) \geq \Omega\left(\frac{\text{len}(W)}{\log^{10} n}\right). \quad (9)$$

We prove Lemma 3.3 when \mathcal{S} consists of 2-moves only in Section 4 and 5, and then generalize the proof to work with general move sequences in Section 6. Assuming Lemma 3.3, we use it to establish our main theorem, restated below:

Theorem 1.2. *Let $\mathcal{X} = (\mathcal{X}_e : e \in E_n)$ be distributions of edge weights such that each \mathcal{X}_e is supported on $[-1, 1]$ and has its density function bounded from above by a parameter $\phi > 0$. Then with probability at least $1 - o_n(1)$ over the draw of edge weights $X \sim \mathcal{X}$, any implementation of the 2-FLIP algorithm takes at most $\phi n^{O(\log^{10} n)}$ steps to terminate.*

Proof. Let $\epsilon > 0$ be specified as follows:

$$\epsilon := \frac{1}{\phi n^{c_1 \log^{10} n}}$$

for some large enough constant $c_1 > 0$ to be specified later. We write \mathcal{F} to denote the following event on the draw of edge weights $X \sim \mathcal{X}$:

Event \mathcal{F} : For every move sequence W of length at most $n \log^{10} n$ such that (letting $a > 0$ be the constant hidden in (9))

$$\max(\text{rank}_{\text{arcs}}(W), \text{rank}_{\text{cycles}}(W)) \geq \frac{a}{\log^{10} n} \cdot \text{len}(W). \quad (10)$$

and every configuration $\gamma_0 \in \{\pm 1\}^n$, (γ_0, W) is not ϵ -improving with respect to X .

We break the proof of the theorem into two steps. First we show that \mathcal{F} occurs with probability at least $1 - o_n(1)$. Next we show that when \mathcal{F} occurs, any implementation of 2-FLIP must terminate in at most $\phi n^{O(\log^{10} n)}$ many steps.

For the first step, we apply Lemma 2.11 on every move sequence W of length at most $n \log^{10} n$ that satisfies (10) and every configuration $\tau_0 \in \{\pm 1\}^{V(W)}$. It then follows from a union bound that \mathcal{F} occurs with probability at least

$$1 - \sum_{\ell \in [n \log^{10} n]} n^{2\ell} \cdot 2^{2\ell} \cdot (\ell \phi \epsilon)^{\frac{a\ell}{\log^{10} n}} = 1 - \sum_{\ell \in [n \log^{10} n]} \left((2n)^{\frac{2\log^{10} n}{a}} \cdot \ell \phi \epsilon \right)^{\frac{a\ell}{\log^{10} n}} = 1 - o_n(1),$$

where the factor $n^{2\ell}$ is an upper bound for the number of W of length ℓ and $2^{2\ell}$ is an upper bound for the number of configurations τ_0 since $|V(W)| \leq 2\ell$. The last equation follows by setting the constant c_1 sufficiently large.

For the second step, we assume that the event \mathcal{F} occurs, and let $\gamma_0, \dots, \gamma_N \in \{\pm 1\}^n$ be a sequence of N configurations that is the result of the execution of some implementation of 2-FLIP under X . Let

$\mathcal{S} = (\mathcal{S}_1, \dots, \mathcal{S}_N)$ denote the move sequence induced by $\gamma_0, \dots, \gamma_N$. So (γ_0, \mathcal{S}) is improving with respect to edge weights X . By Lemma 3.2, \mathcal{S} is a valid sequence.

We use the event \mathcal{F} to bound the length of \mathcal{S} . Because of \mathcal{F} and that \mathcal{S} is a valid move sequence, it follows from Lemma 3.3 that the objective function gets improved by at least ϵ for every $n \log^{10} n$ consecutive moves in \mathcal{S} . Given that the objective function lies between $[-n^2, n^2]$, we have

$$\text{len}(\mathcal{S}) \leq n \log^{10} n \cdot \frac{2n^2}{\epsilon} \leq \phi n^{O(\log^{10} n)}. \quad (11)$$

This finishes the proof of the theorem. \square

Corollary 3.4. *Under the same setting of Theorem 1.2, the same result holds for Pure 2-FLIP.*

Proof. The only property of the sequence of moves used in the proof of Theorem 1.2 is that it is a valid sequence, and this property holds for Pure 2-FLIP as well. \square

Notice that by twining the constant c_1 in the exponent, we can control the tail-bound of the failure probability. Thus, we can strengthen our proof to get the same bound for the expected number of steps needed to terminate as in the standard smoothed analysis prototype :

Corollary 3.5. *Under the same setting of Theorem 1.2, any implementation of the 2-FLIP algorithm (or Pure 2-FLIP) takes at most $\phi n^{O(\log^{10} n)}$ many steps to terminate on expectation.*

Proof. We let \mathcal{F}_ϵ denote the event \mathcal{F} in the proof of Theorem 1.2 with a specified $\epsilon > 0$. Let $\epsilon_0 = 1/(\phi \cdot n^{c_1 \log^{10} n})$, where $c_1 > 0$ is a constant to be fixed shortly. For any $\epsilon < \epsilon_0$, we have that

$$\begin{aligned} \Pr[\neg \mathcal{F}_\epsilon] &\leq \sum_{\ell \in [n \log^{10} n]} n^{2\ell} \cdot 2^{2\ell} \cdot (\ell \phi \epsilon)^{\max(ranks_{arcs}(W), ranks_{cycles}(W))} \leq \sum_{\ell \in [n \log^{10} n]} n^{2\ell} \cdot 2^{2\ell} \cdot (\ell \phi \epsilon \cdot \frac{\epsilon_0}{\epsilon_0})^{\lceil \frac{a\ell}{\log^{10} n} \rceil} \\ &\leq \sum_{\ell \in [n \log^{10} n]} n^{2\ell} \cdot 2^{2\ell} \cdot (\ell \phi \epsilon_0)^{\lceil \frac{a\ell}{\log^{10} n} \rceil} \cdot (\frac{\epsilon}{\epsilon_0})^{\lceil \frac{a\ell}{\log^{10} n} \rceil} \leq \sum_{\ell \in [n \log^{10} n]} \frac{1}{n^3} \left(\frac{\epsilon}{\epsilon_0}\right)^{\lceil \frac{a\ell}{\log^{10} n} \rceil} \leq \frac{\epsilon}{c_2 n \epsilon_0} \end{aligned}$$

where $c_1 = 2 + 6/a$ (letting $a > 0$ be the constant hidden in (9)) and $c_2 = 10^8$. From the proof of Theorem 1.2, conditionally to the event \mathcal{F}_ϵ for any $\epsilon \leq \epsilon_0$, $\text{len}(\mathcal{S}) \leq L(\epsilon) := \frac{2n^3 \log^{10} n}{\epsilon}$. Notice that for any $\epsilon(\rho) := \epsilon_0/\rho$, $L(\epsilon(\rho)) = \rho L(\epsilon_0)$. Thus, the probability that $\text{len}(\mathcal{S})$ is larger than $cL(\epsilon)$ for any $\rho \geq 1$ is

$$\Pr[\text{len}(\mathcal{S}) > \rho L(\epsilon)] \leq \Pr[\neg \mathcal{F}_{\epsilon_0/\rho}] \leq \frac{1/\rho}{c_2 \cdot n}.$$

Note that L is always trivially bounded by the total number of configurations, 2^n . Therefore, we have

$$\begin{aligned} \mathbb{E}[\text{len}(\mathcal{S})] &= \sum_{s=1}^{2^n} \Pr[\text{len}(\mathcal{S}) \geq s] \leq \sum_{s=1}^{\lceil L(\epsilon_0) \rceil} \Pr[L \geq s] + \sum_{s=\lceil L(\epsilon_0) \rceil}^{2^n} \Pr[L \geq s] \leq L(\epsilon_0) + \sum_{s=\lceil L(\epsilon_0) \rceil}^{2^n} \Pr[L \geq s] \\ &\leq L(\epsilon_0) + \sum_{s=\lceil L(\epsilon_0) \rceil}^{2^n} \Pr\left[L \geq \frac{s}{L(\epsilon_0)} \cdot L(\epsilon_0)\right] \leq L(\epsilon_0) + \sum_{s=\lceil L(\epsilon_0) \rceil}^{2^n} \frac{L(\epsilon_0)/s}{c_2 n} = O(n) \cdot L(\epsilon_0) = \phi n^{O(\log^{10} n)}. \end{aligned}$$

This finishes the proof of Corollary 3.5. \square

The same results hold for the Graph Partitioning problem and the SWAP neighborhood.

Theorem 1.1. *Let $\mathcal{X} = (\mathcal{X}_e : e \in E_{2n})$ be distributions of edge weights such that each \mathcal{X}_e is supported on $[-1, 1]$ and has its density function bounded from above by a parameter $\phi > 0$. Then with probability at least $1 - o_n(1)$ over the draw of edge weights $X \sim \mathcal{X}$, any implementation of SWAP takes at most $\phi n^{O(\log^{10} n)}$ steps to terminate.*

Proof. Every move sequence \mathcal{S} generated by SWAP (for any pivoting rule, any weights, and any initial balanced partition) is also a legal move sequence for Pure 2-FLIP on the same instance, except that the sequence may be incomplete for Pure 2-FLIP, that is, the final partition may not be locally optimal for Pure 2-FLIP, since there may be a 2-move (but not a swap) that improves the weight of the cut (the resulting partition would not be balanced), and Pure 2-FLIP would continue and produce a longer sequence. Hence, the number of steps of SWAP is upper bounded by the number of steps of Pure 2-FLIP, and thus it is at most $\phi n^{O(\log^{10} n)}$ with probability $1 - o_n(1)$, as well as in expectation. \square

4 Windows in a Valid Sequence of 2-Moves

We will start with the proof of Lemma 3.3 for the case when \mathcal{S} consists of 2-moves only in Section 4 and 5, and generalize it to deal with general move sequences in Section 6.

We start with a combinatorial argument about sets and subsequences of $[N]$, where $N = \text{poly}(n)$ for any polynomial at n . Let I be a subset of $[N]$ with $|I| \geq \log^{10} n$. Intuitively, later in this section I will be chosen to be the set I_u representing the appearances of some frequently appeared active node $u \in V(\mathcal{S})$ in a move sequence \mathcal{S} . We will write $\text{order}(i)$ to denote the order of $i \in I$. In other words, the smallest index in I has order 1 and the largest index in I has order $|I|$. To give an example if $I = \{2, 5, 9, 11\}$ then $\text{order}(2) = 1, \text{order}(5) = 2$ and so on. Let $\delta = 0.01$. We start by quantifying how much large windows centered around an index $i \in I$ should be to cover the majority of a set I . Afterwards, we present the combinatorial lemmas about subset I .

Definition 4.1. *Let $I \subseteq [N]$. We say an index $i \in I$ is ℓ -good for some positive integer ℓ if*

$$[i - \lceil (1 + 2\delta)L' \rceil : i + \lceil (1 + 2\delta)L' \rceil] \subseteq [N], \quad \text{where } L' = \lceil (1 + \delta)^{\ell-1} \rceil$$

and I satisfies

$$|I \cap [i - L' : i + L']| \geq \log^3 n \quad \text{and} \quad |I \cap [i - \lceil (1 + 2\delta)L' \rceil : i + \lceil (1 + 2\delta)L' \rceil]| \leq \log^7 n. \quad (12)$$

If there exists no such constant ℓ , we call the corresponding index bad.

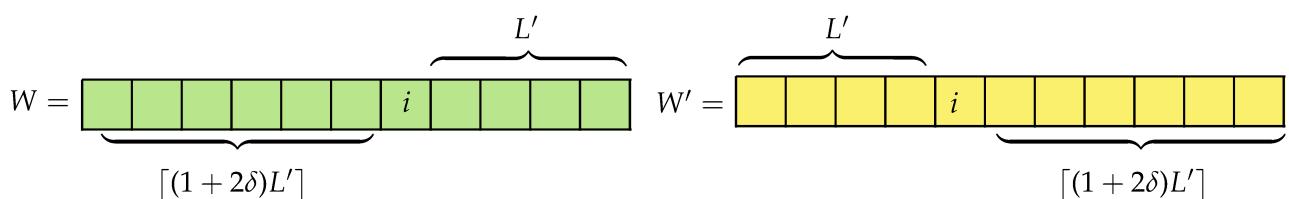


Figure 3: Two examples of windows whose intersection in I is between $[\log^3 n, \log^7 n]$.

Remark 4.2. Some motivation behind the definition 4.1: When $i \in I$ is ℓ -good (letting $L = L' + \lceil(1+2\delta)L'\rceil + 1$ and $L' = \lceil(1+\delta)^{\ell-1}\rceil$), it implies that all the $\lceil(1+2\delta)L'\rceil - L' \geq 2\delta L' = \Omega(L)$ ⁵ windows W of length L , i.e., those start at $i - \lceil(1+2\delta)L'\rceil, \dots, i - L'$, satisfy

$$i \in W, \quad |I \cap W| \geq \log^3 n \quad \text{and} \quad |I \cap W| \leq \log^7 n.$$

Remark 4.3. By definition 4.1, ℓ can get at most $\log_{1+\delta} N = \Theta(\log n)$, for any $N = \text{poly}(n)$.

Lemma 4.4. Suppose I is a subset of $[N]$ with $|I| \geq \log^8 n$. Then at least a $(1 - O(1/\log n))$ -fraction of $i \in I$ is ℓ_i -good for some nonnegative integer ℓ_i .

Proof. We start by defining an ℓ_i for each $i \in I$ (except for the smallest $\lceil \log^7 n \rceil$ indices and the largest $\lceil \log^7 n \rceil$ indices in I , which are negligible since $|I| \geq \log^8 n$) and then show that most $i \in I$ is ℓ_i -good. Let I' be the subset of I after removing the smallest $\lceil \log^7 n \rceil$ indices and the largest $\lceil \log^7 n \rceil$ indices in I . For each $i \in I'$, let

- $j \in I'$ be the index in I' of order $\text{order}(i) - \lfloor \log^7 n/2 \rfloor + 1$,
- $k \in I'$ be the index in I' of order $\text{order}(i) + \lfloor \log^7 n/2 \rfloor - 1$,
- Δ be the minimum distance between index i and indices j, k , $\Delta = \min(i - j, k - i)$ and
- ℓ_i be the largest integer such that $\lceil(1+2\delta) \cdot (1+\delta)^{\ell_i-1}\rceil \leq \Delta - 2$.

Using the fact that for any real positive number x , it holds that $0 \leq \lceil(1+2\delta) \cdot \lceil x \rceil\rceil - \lceil(1+2\delta) \cdot x\rceil \leq 2$, we get that:

$$\begin{cases} \lceil(1+2\delta) \cdot \lceil(1+\delta)^{\ell_i-1}\rceil\rceil \leq \Delta \\ \lceil(1+2\delta) \cdot (1+\delta)^{\ell_i}\rceil > \Delta - 2 \end{cases}.$$

For the rest of the proof, let $L'_i = \lceil(1+\delta)^{\ell_i-1}\rceil$. It follows from the choice of ℓ_i that

$$|I \cap [i - \lceil(1+2\delta)L'_i\rceil : i + \lceil(1+2\delta)L'_i\rceil]| \leq 2(\lfloor \log^7 n/2 \rfloor - 1) + 1 \leq \log^7 n \quad (13)$$

$$|I \cap [i - \lceil(1+\delta) \cdot (1+2\delta)L'_i\rceil : i + \lceil(1+\delta) \cdot (1+2\delta)L'_i\rceil]| \geq (\lfloor \log^7 n/2 \rfloor - 1) + 1 - 4 \quad (14)$$

For (14), we use the observation that left-hand side is larger than $|I \cap [i - (\Delta - 2) : i + (\Delta - 2)]| \geq |I \cap [i - \Delta : i + \Delta]| - 4$. Using $(1+\delta)(1+2\delta) \leq 1+4\delta$ with $\delta = 0.01$, the second inequality implies

$$|I \cap [i - \lceil(1+4\delta)L'_i\rceil : i + \lceil(1+4\delta)L'_i\rceil]| \geq \lfloor \log^7 n/2 \rfloor - 4. \quad (15)$$

On the other hand, (13) implies that $i \in I'$ is ℓ_i -good unless

$$|I \cap [i - L'_i : i + L'_i]| \leq \log^3 n. \quad (16)$$

Assume now for a contradiction that the number of $i \in I'$ that are bad is at least $|I|/\log n$. Addition-

⁵Indeed, by definition $L = L' + 1 + \lceil(1+2\delta)L'\rceil \leq 2((1+\delta)L' + 1) \leq 2((1+\delta)L' + 1) \leq 2(2+\delta)L'$, since $L' \geq 1$

ally, for any possible exponent $\ell \in [\log_{1+\delta} N]$, let \mathcal{R}_ℓ be the set of the indices i that are not ℓ -good:

$$\mathcal{R}_\ell := \{i \in I' \text{ s.t } i \text{ is not } \ell\text{-good}\} \text{ & } \ell^* = \operatorname{argmax}_{\ell \in [\log_{1+\delta} N]} |\mathcal{R}_\ell|$$

Then, it holds that

$$|\mathcal{R}_{\ell^*}| \geq \frac{|I|/\log n}{\log_{1+\delta} N} = \Omega(|I|/\log^2 n),$$

where we use the facts that $|I| \geq \log^8 n$ and $N = \text{poly}(n)$. We define then $L^* = \lceil (1+\delta)^{\ell^*-1} \rceil$ and for each $\rho \in \mathcal{R}_{\ell^*}$ we let

$$B_\rho = (I \cap [\rho - \lceil (1+4\delta)L^* \rceil, \rho - L^*]) \cup (I \cap [\rho + L^*, \rho + \lceil (1+4\delta)L^* \rceil]).$$

Note now that when an index i is not ℓ -good, we have from (16) and (15) that

$$\left| (I \cap [i - \lceil (1+4\delta)L'_i \rceil : i - L'_i]) \cup (I \cap [i + L'_i : i + \lceil (1+4\delta)L'_i \rceil]) \right| \geq \lfloor \log^7 n / 2 \rfloor - 4 - \log^3 n = \Omega(\log^7 n).$$

Hence, we have that $|B_\rho| \geq \Omega(\log^7 n)$ for every $\rho \in \mathcal{R}_{\ell^*}$ and thus,

$$\sum_{\rho \in \mathcal{R}_{\ell^*}} |B_\rho| \geq \Omega\left(\frac{|I|}{\log^2 n}\right) \cdot \Omega(\log^7 n) = \Omega(|I| \log^5 n).$$

On the other hand, we can prove the following claim:

Claim 4.5. *For any $i \in I$, the number of $\rho \in \mathcal{R}_{\ell^*}$ such that $i \in B_\rho$ is at most $O(\log^3 n)$.*

It follows then from the claim that

$$\text{For any } i \in I, \text{ we get } |\{\rho \in \mathcal{R}_{\ell^*} | i \in B_\rho\}| = O(\log^3 n) \Rightarrow \sum_{\rho \in \mathcal{R}_{\ell^*}} |B_\rho| \leq |I| \cdot O(\log^3 n),$$

which leads to a contradiction. \square

Proof of Claim 4.5. Fix any $i \in I$. Let us assume then that ρ be a $\rho \in \mathcal{R}_{\ell^*}$ such that $i \in B_\rho$. We prove that the number of $\rho \in \mathcal{R}_{\ell^*}$ with $\rho > i$ and $i \in B_\rho$ is at most $O(\log^3 n)$; the case with $\rho < i$ is symmetric. If no such ρ exists then the claim is trivially true. Hence, let's assume that such one exists with $\rho > i$. Given that $i \in B_\rho$, we have that $i \in [\rho - \lceil (1+4\delta)L^* \rceil, \rho - L^*]$ and we also have

$$|I \cap [\rho - L^*, \rho + L^*]| \leq \log^3 n. \tag{17}$$

On the other hand, every other $\rho' \in \mathcal{R}_{\ell^*}$ that satisfies $\rho' > i$ and $i \in B_{\rho'}$ also has the property that $i \in [\rho' - \lceil (1+4\delta)L^* \rceil, \rho' - L^*]$ and thus, $\rho' \in [i + L^*, i + \lceil (1+4\delta)L^* \rceil]$. But combining this with $i \in [\rho - (1+4\delta)L^*, \rho - L^*]$ we have

$$\rho' \leq i + \lceil (1+4\delta)L^* \rceil \leq \rho + \lceil (1+4\delta)L^* \rceil - L^*$$

and

$$\rho' \geq i + L^* \geq \rho - \lceil (1 + 4\delta)L^* \rceil + L^*.$$

So $\rho' \in [\rho - \lceil (1 + 4\delta)L^* \rceil + L^*, \rho + \lceil (1 + 4\delta)L^* \rceil - L^*] \subseteq [\rho - L^*, \rho + L^*]$ and by (17) the number of such ρ' is no more than $\log^3 n$. \square

Now we return to work on our problem and an arbitrary move sequence $\mathcal{S} = (\mathcal{S}_1, \dots, \mathcal{S}_N)$. Let W be a window (move sequence) of \mathcal{S} . For each active node $u \in V(W)$, we write $\#_W(u)$ to denote the number of occurrences of u in W . The main result in this section is the following lemma:

Lemma 4.6. *Let \mathcal{S} be a move sequence of length $N = n \log^{10} n$ that consists of 2-moves only. There exists a positive integer L such that \mathcal{S} has at least $\Omega((N - L + 1)/\log n)$ many windows $W = (W_1, \dots, W_L)$ of length L such that at least $\Omega(L/\log n)$ moves $W_i = \{u, v\}$ of W satisfy*

$$\log^3 n \leq \#_W(u) \leq \log^7 n \quad \text{and} \quad \#_W(v) \geq \log^3 n \quad (18)$$

Proof. For each node $u \in V(\mathcal{S})$ we write $I_u \subseteq [N]$ to denote the set of $i \in [N]$ with $u \in \mathcal{S}_i$. We say the i -th move $\mathcal{S}_i = \{u, v\}$ is ℓ -good for some positive integer ℓ if i is ℓ_1 -good in I_u and i is ℓ_2 -good in I_v for some positive integers ℓ_1, ℓ_2 such that $\ell = \max(\ell_1, \ell_2)$.

Let $\mathcal{S}_i = \{u, v\}$. Then we consider the following cases:

1. Either $|I_u|$ or $|I_v|$ is smaller than $\log^8 n$: Given that no more than $n \log^8 n$ moves can contain a vertex that appears less than $\log^8 n$ times in the sequence, we have the number of such i is at most

$$n \log^8 n = o(N/\log n);$$

2. $|I_u|, |I_v| \geq \log^8 n$ but either u is not ℓ_1 -good for any ℓ_1 or v is not ℓ_2 -good for any ℓ_2 : By Lemma 4.4, the number of such i is at most (using $\sum_u |I_u| = 2N$)

$$\sum_{u: |I_u| \geq \log^8 n} \frac{|I_u|}{\log n} \leq \frac{2N}{\log n}.$$

3. Otherwise, \mathcal{S}_i is ℓ -good by setting $\ell = \max(\ell_1, \ell_2)$.

Thus, the number of $i \in [N]$ such that the \mathcal{S}_i is ℓ -good for some ℓ is at least $(1 - 3/\log n)N$.

Given that ℓ is at most $O(\log N) = O(\log n)$, there exists a positive integer ℓ such that the number of moves in \mathcal{S} that are ℓ -good is at least $\Omega(N/\log n)$. Let $L' = \lceil (1 + \delta)^{\ell-1} \rceil$ and

$$L = L' + \lceil (1 + 2\delta)L' \rceil + 1.$$

For any move $\mathcal{S}_i = \{u, v\}$ that is ℓ -good, it is easy to verify that there are $\Omega(L)$ windows W of length L that contain i and satisfy (18) (See Remark 4.2).

Let's pick a window W of \mathcal{S} of size L uniformly at random; note that there are $N - L + 1$ many such windows in total. Let X be the random variable that denotes the number of moves in W that satisfy

(18). Given that the number of moves that are ℓ -good is at least $\Omega(N/\log n)$, we have

$$\mathbf{E}[X] \geq \Omega\left(\frac{N}{\log n}\right) \cdot \frac{\Omega(L)}{N} = \Omega\left(\frac{L}{\log n}\right).$$

Let a be the constant hidden above. Given that we always have $X \leq L$, we have

$$\Pr\left[X \geq \frac{aL}{2\log n}\right] \geq \frac{a}{2\log n} \quad (19)$$

since otherwise,

$$\mathbf{E}[X] \leq \frac{a}{2\log n} \cdot L + \left(1 - \frac{a}{2\log n}\right) \cdot \frac{aL}{2\log n} < \frac{aL}{\log n}.$$

a contradiction. The lemma then follows directly from (19). \square

5 Finding Cycles

Let $\mathcal{S} = (\mathcal{S}_1, \dots, \mathcal{S}_N)$ be a valid move sequence of length $N = n \log^{10} n$ that consists of 2-moves only. By Lemma 4.6, \mathcal{S} has a window $W = (W_1, \dots, W_L)$ of length L such that the number of moves in W that satisfy (18) is at least $\Omega(L/\log n)$. We show in this section that such a W satisfies

$$\text{rank}_{\text{cycles}}(W) = \Omega\left(\frac{L}{\log^{10} n}\right). \quad (20)$$

This will finish the proof of Lemma 3.3 when \mathcal{S} consists of 2-moves only.

To this end, let $\tau_0 \in \{\pm 1\}^{V(W)}$ be the configuration with $\tau_0(u) = -1$ for all $u \in V(W)$ so that we can work on vectors $\text{imprv}_{\tau_0, W}(i)$ and $\text{imprv}_{\tau_0, W}(C)$ for dependent cycles of W (at the same time, recall from Lemma 2.10 that $\text{rank}_{\text{cycles}}(W)$ does not depend on the choice of τ_0). Let τ_0, \dots, τ_L denote the sequence of configurations induced by W .

Next, let us construct an auxiliary graph $H = (V(W), E)$, where every move $W_i = \{u, v\}$ adds an edge between u and v in E . Note that we allow parallel edges in H so $|E| = L$ and $\#_W(u)$ is exactly the degree of u in H . There is also a natural one-to-one correspondence between cycles of W and cycles of H . The following lemma shows the existence of a nice looking bipartite graph in H :

Lemma 5.1. *There are two disjoint sets of nodes $V_1, V_2 \subset V(W)$ and a subset of edges $E' \subseteq E$ such that*

1. *Every edge in E' has one node in V_1 and the other node in V_2 ;*
2. *$|V_1 \cup V_2| = O(L/\log^3 n)$ and $|E'| = \Omega(L/\log n)$;*
3. *$\#_W(u) \leq \log^7 n$ for every node $u \in V_1$.*

Proof. Let V be the set of vertices v such that $\#_W(v) \geq \log^3 n$. We start our proof with the size of V :

$$|V| \log^3 n \leq |V| \min_{v \in V} \#_W(v) = |V| \min_{v \in V} \deg_H(v) \leq \sum_{v \in V} \deg_H(v) = 2|E(H)| \leq 2L.$$

We further partition V into V_ℓ and V_h such that V_ℓ contains those in V with $\#_W(v) \leq \log^7 n$ and V_h contains those with $\#_W(v) > \log^7 n$. By Lemma 4.6 we can assume the number of edges incident to at least one vertex in V_ℓ (that is edges in $V_\ell \times V_\ell \cup V_\ell \times V_h$) is at least $\Omega(L/\log n)$. Suppose we construct V_1 and V_2 by randomly put each node in V_ℓ in V_1 or V_2 and put all nodes in V_h in V_2 . Any edge in $V_\ell \times V_\ell$ or $V_\ell \times V_h$ is between V_1 and V_2 with $1/2$ probability. Thus $\mathbb{E}[\text{EdgesInCut}(V_1, V_2)] = |E|/2 = \Omega(L/\log n)$. Thus, by standard probabilistic argument, there exist at least one assignment of V_ℓ to V_1 and V_2 such that at least half of the edges in $V_\ell \times V_\ell \cup V_\ell \times V_h$ are included. Hence, we get a bipartite graph between V_1 and V_2 with at least $\Omega(L/\log n)$ edges, and any node v in V_1 satisfies $\log^3 n \leq \#_W(v) \leq \log^7 n$. Notice that since $V = |V_1 \cup V_2| = |V_\ell \cup V_h|$, we get that $|V_1 \cup V_2| = O(L/\log^3 n)$. \square

Recall the definition of dependent cycles of W (and their cancellation vectors) from Section 2. Since we only care about the rank of vectors induced by dependent cycles, we give the following definition which classify each edge of H into two types and then use it to give a sufficient condition for a cycle of W to be dependent:

Definition 5.2. We say the i -th move $W_i = \{u, v\}$ of W is of the same sign if $\tau_i(u) = \tau_i(v)$, and is of different signs if $\tau_i(u) \neq \tau_i(v)$.

Lemma 5.3. Let $C = (c_1, \dots, c_t)$ be a cycle of W and assume that t is even. If all of W_{c_1}, \dots, W_{c_t} are of different signs, then C is a dependent cycle; If all of W_{c_1}, \dots, W_{c_t} are of the same sign, then C is a dependent cycle of W .

Proof. Recall the Dependence Criterion (Remark 2.7)

$$C \text{ is a dependent cycle of } W \Leftrightarrow (-1)^t = \tau_{c_t}(u_1)\tau_{c_1}(u_1) \cdot \prod_{i=2}^t \tau_{c_{i-1}}(u_i)\tau_{c_i}(u_i)$$

If all of W_{c_1}, \dots, W_{c_t} are of different signs, then $\tau_{c_j}(u_j)\tau_{c_j}(u_{j+1}) = \tau_{c_t}(u_1)\tau_{c_t}(u_t) = -1$, and the above expression equals to $(-1)^t = (-1)^t$. If all of W_{c_1}, \dots, W_{c_t} are of the same sign, then $\tau_{c_j}(u_j)\tau_{c_j}(u_{j+1}) = \tau_{c_t}(u_1)\tau_{c_t}(u_t) = 1$, the above expression is also $1 = (-1)^t$, which holds since t is even. \square

We assume in the rest of the proof that at least half of edges in E' are of the same sign; the case when at least half of E' are of different signs can be handled similarly. Let E'' be the subset of E' that consists of edges of the same sign, with $|E''| \geq |E'|/2$. In the following discussion, cycles in E'' always refer to cycles that do not use the same edge twice (parallel edges are counted as different edges, since they correspond to different moves in the window W).

The aforementioned discussion leads to the following corollary which reduces the existence of dependent cycle of W to a simple cycle in auxiliary graph H :

Corollary 5.4. Since every cycle in a bipartite graph has even length, every cycle in E'' corresponds to a dependent cycle of W . For convenience, given any cycle C of E'' we will write $\text{imprv}_{\tau_0, W}(C)$ to denote the vector of its corresponding dependent cycle of W .

We first deal with the case when E'' contains many parallel edges:

Lemma 5.5. Let D be the subset of nodes in V_1 that have parallel edges in E'' . Then $\text{rank}_{\text{cycles}}(W) \geq |D|/2$.

Proof. We prove the lemma even if the sequence contains both 1-moves and 2-moves so that we can use it also in the general case in the next section. We note first that if $S_i = \{u, v\}$, $S_j = \{u, v\}$, $i < j$ are two moves that involve the same two nodes, then there is at least one node $z \neq u, v$ that appears an odd number of times between the two moves. This follows from the definition of a valid move sequence.

We will construct a set Q of at least $|D|/2$ 2-cycles, where each 2-cycle consists of two parallel edges in E'' . We use the following procedure.

1. While there is a 2-cycle (u, v) with $u \in D$, $v \in V_2$, such that some node $z \neq u$ of V_1 moves an odd number of times between the two $\{u, v\}$ moves of the 2-cycle, pick any such 2-cycle (u, v) and add it to our set Q , pick any such node $z \neq u$ that moves an odd number of times between the two $\{u, v\}$ moves, and delete u and z from D (if z is in D).
2. Suppose now that there are no more 2-cycles as in step 1. While D is not empty, let u be any remaining node in D , take any two incident parallel edges $\{u, v\}$ in E'' , add the corresponding 2-cycle to Q , and delete u from D .

Firstly, notice that for every new entry at Q in the procedure, we delete at most 2 nodes from D . Hence, this procedure will generate clearly a set Q of at least $|D|/2$ 2-cycles. Let $(u_1, v_1), (u_2, v_2), \dots, (u_k, v_k)$ be the sequence of 2-cycles selected, where the first d were selected in step 1, and the rest in step 2. The nodes u_i are distinct, while the nodes v_i may not be distinct. For each $i = 1, \dots, d$, let z_i be the node in V_1 that appears an odd number of times between the two $\{u_i, v_i\}$ moves that was selected by the algorithm. Note that node $z_i \neq u_j$ for all $j \geq i$, since z_i was deleted from D when u_j was selected. For each $i = d + 1, \dots, k$, let z_i be any node, other than u_i, v_i , that appears an odd number of times between the two $\{u_i, v_i\}$ moves. Then z_i is not in V_1 because in step 2 there are no odd nodes in V_1 . For each $i = 1, \dots, k$, we view the edge $\{u_i, z_i\}$ of the complete graph as a witness for the 2-cycle (u_i, v_i) .

Consider the matrix with columns corresponding to the selected 2-cycles (u_i, v_i) , $i = 1, \dots, k$, and rows corresponding to the witness edges $\{u_i, z_i\}$. The entry for the corresponding witness edge $\{u_i, z_i\}$ is nonzero. Indeed, by definition 2.6, $\text{imprv}_{\tau_0, W}(C = (u_i, v_i))_{\{u_i, z_i\}} = -b_1(\tau_{c_1}(u_i)\tau_{c_1}(z_i)) - b_2(\tau_{c_2}(u_i)\tau_{c_2}(z_i))$

$$\text{and } \begin{cases} \tau_{c_1}(z_i) = -\tau_{c_2}(z_i) \\ b_1 = 1 \& b_2 = -\tau_{c_1}(v_i)\tau_{c_2}(v_i) \\ \tau_{c_m}(v_i) = \tau_{c_m}(u_i), \text{ for } m \in \{1, 2\} \end{cases} \quad \text{which yields } \text{imprv}_{\tau_0, W}(C = (u_i, v_i))_{\{u_i, z_i\}} = 2 \cdot \tau_{c_1}(z_i) \neq 0.$$

Consider the column for a 2-cycle (u_i, v_i) selected in step 1. The entry for any other witness edge $\{u_j, z_j\}$ with $j < i$ is 0 because $u_j, z_j \neq u_i, v_i$. (The entries for witness edges $\{u_j, z_j\}$ with $j > i$ could be nonzero.)

Consider the column for a 2-cycle (u_i, v_i) selected in step 2. The entry for any witness edge $\{u_j, z_j\}$ from step 1 (i.e. with $j \leq d$) is 0 because $u_j, z_j \neq u_i, v_i$. The entry for any witness edge $\{u_j, z_j\}$ from step 2 (i.e. with $j > d$) is also 0 because (1) $u_j \neq u_i, v_i$, (2) $z_j \notin V_1$ hence $z_j \neq u_i$, and (3), even if $z_j = v_i$, all nodes of V_1 –hence also u_j –occur an even number of times between the two $\{u_i, v_i\}$ moves, therefore

the entry for $\{u_j, v_i\}$ is 0.

$$\begin{cases} \mathcal{M}_{1 \rightarrow d}^{[\text{step 1}]} = \begin{bmatrix} \text{imprv}_{\tau_0, W}(C_1)_{\{x_1, y_1\}} \neq 0 & 0 & 0 \\ * & \ddots & \vdots \\ * & * & \text{imprv}_{\tau_0, W}(C_d)_{\{x_d, y_d\}} \neq 0 \end{bmatrix} \Rightarrow \mathcal{M} = \begin{bmatrix} \mathcal{M}_{1 \rightarrow d}^{[\text{step 1}]} & \mathbf{0} \\ * & \mathcal{M}_{d+1 \rightarrow k}^{[\text{step 2}]} \end{bmatrix} \\ \mathcal{M}_{d+1 \rightarrow k}^{[\text{step 2}]} = \text{diag}_{t \in (d+1) \rightarrow k}(\text{imprv}_{\tau_0, W}(C_t)_{\{x_t, y_t\}} \neq 0) \end{cases}$$

Thus, the matrix with columns corresponding to the selected 2-cycles (u_i, v_i) and rows corresponding to their witness edges $\{u_i, z_i\}$ is a lower triangular matrix with non-zero diagonal entries. It follows that the columns are linearly independent. \square

As a result, it suffices to deal with the case when $|D|$ is $o(L/\log^8 n)$. Let E^* denote the subset of edges obtained from E'' after deleting all nodes of D and their incident edges. The remaining bipartite graph has no parallel edges. Then we have

$$|E^*| \geq |E''| - |D| \cdot \log^7 n = \Omega(L/\log n).$$

We list all properties of the bipartite graph $H^* = (V_1 \cup V_2, E^*)$ we need as follows:

1. H^* is a bipartite graph with no parallel edges;
2. $|V_1 \cup V_2| \leq O(L/\log^3 n)$ and $|E^*| \geq \Omega(L/\log n)$; and
3. $\#_W(u) \leq \log^7 n$ for every node $u \in V_1$.
4. Every edge $e = \{u, v\} \in E^*$ corresponds to a move $W_i = \{u, v\}$ which is of the same sign.

Recall that $E(W)$ denotes the set of edges in K_n which have both nodes in $V(W)$. These edges are indices of $\text{imprv}_{\tau_0, S}(i)$ and $\text{imprv}_{\tau_0, S}(C)$ for a given dependent cycle C of W . Our main lemma is the following:

Lemma 5.6. *Fix an arbitrary $s \in [0 : L/\log^{10}(n)]$. Assume additionally that there exists a set of edges $\mathcal{E}_s = \{\{x_1, y_1\}, \dots, \{x_s, y_s\}\} \subseteq E(W)$ such that $x_i \in V_1$ for all $i \in [s]$. Then there exists a cycle C in H^* and an edge $\{u, v\} \in E(W)$ with $u \in V_1 \setminus \{x_1, y_1, \dots, x_s, y_s\}$ such that*

$$(\text{imprv}_{\tau_0, W}(C))_{\{u, v\}} \neq 0 \quad \text{and} \quad (\text{imprv}_{\tau_0, W}(C))_{\{x_i, y_i\}} = 0, \quad \text{for all } i \in [s]. \quad (21)$$

Proof of (20) Assuming Lemma 5.6. Start with $\mathcal{E}_{s=0} = \emptyset$, For integer s going from 0 to $\lfloor L/\log^{10} n \rfloor$, using Lemma 5.6, find cycle C_{s+1} and an edge $\{u, v\}$ satisfying (21), let $\mathcal{E}_{s+1} = \mathcal{E}_s \cup \{x_{s+1}, y_{s+1}\} = \{u, v\}$ and repeat the above process.

In the end, we get a set of cycles C_1, \dots, C_k where $k = \lfloor L/\log^{10} n \rfloor$. And for any $j \in [k]$, we have

$$(\text{imprv}_{\tau_0, W}(C_j))_{\{x_j, y_j\}} \neq 0 \quad \text{and} \quad (\text{imprv}_{\tau_0, W}(C_j))_{\{x_i, y_i\}} = 0, \quad \text{for all } i \in [j-1].$$

Let \mathcal{M} be the $k \times k$ square matrix where $\mathcal{M}_{ij} = (\text{imprv}_{\tau_0, W}(C_j))_{\{x_i, y_j\}}.$

$$\mathcal{M} = \begin{bmatrix} \text{imprv}_{\tau_0, W}(C_1)_{\{x_1, y_1\}} \neq 0 & 0 & 0 & \cdots & 0 \\ * & \text{imprv}_{\tau_0, W}(C_2)_{\{x_2, y_2\}} \neq 0 & 0 & \cdots & 0 \\ \vdots & * & \text{imprv}_{\tau_0, W}(C_3)_{\{x_3, y_3\}} \neq 0 & \cdots & 0 \\ \vdots & \vdots & * & \ddots & 0 \\ * & * & * & * & \text{imprv}_{\tau_0, W}(C_k)_{\{x_k, y_k\}} \neq 0 \end{bmatrix}$$

As we can see, the matrix is lower triangular with non-zero diagonal entries, so it has full rank k . Note that \mathcal{M} is a submatrix of the matrix formed by taking $\text{imprv}_{\tau_0, W}(C_j)$ as column vectors, therefore we have $\text{rank}_{\text{cycles}}(W) \geq k \geq L / \log^{10} n$. \square

5.1 Proof of Lemma 5.6

Given a cycle C in H^* , we say $\{u, v\} \in E(W)$ is a *witness* of C if

$$(\text{imprv}_{\tau_0, W}(C))_{\{u, v\}} \neq 0.$$

So the goal of Lemma 5.6 is to find a cycle C of H^* such that none of $(u_i, v_i) \in \mathcal{E}_s$ are witnesses of C and at the same time, C has a witness edge $\{u, v\}$ with u being a new node in V_1 not seen in \mathcal{E}_s before. The proof consists of two steps. First we introduce a so-called *split auxiliary graph* G using H^* and \mathcal{E}_s , by deleting certain nodes and creating extra copies of certain nodes in H^* . We show in Lemma 5.7 that certain simple cycles in G correspond to cycles in H^* that don't have any edge in \mathcal{E}_s as witnesses. Next we show in Lemma 5.9 how to find such a simple cycle in G that has a new witness (u, v) such that $u \in V_1$ and does not appear in \mathcal{E}_s .

Let $\text{wit}_1(\mathcal{E}_s)$ be the set of $u \in V_1$ that appear in \mathcal{E}_s and let $\text{wit}_2(\mathcal{E}_s)$ be the set of $v \in V_2$ that appear in \mathcal{E}_s . For each $v \in \text{wit}_2(\mathcal{E}_s)$, we write $\text{wit}_1(v) \neq \emptyset$ to denote the set of nodes $u \in \text{wit}_1(\mathcal{E}_s)$ such that $(u, v) \in \mathcal{E}_s$, and let k_v denote the number of moves in W that involve at least one node in $\text{wit}_1(v)$. We have $k_v \leq |\text{wit}_1(v)| \cdot \log^7 n$ since $\#_W(u) \leq \log^7 n$ for all $u \in V_1$. Below, we give an example of such an auxiliary graph:

We now define our split auxiliary (bipartite) graph G . We start with its set of nodes $V'_1 \cup V'_2$:

1. $V'_1 = V_1 \setminus \text{wit}_1(\mathcal{E}_s)$; and
2. $V'_2 = \cup_{v \in V_2} C(v)$, where $C(v) = \{v^{(0)}\}$ if $v \notin \text{wit}_2(\mathcal{E}_s)$ and $C(v) = \{v^{(0)}, v^{(1)}, \dots, v^{(k_v)}\}$ if $v \in \text{wit}_2(\mathcal{E}_s)$.

So we deleted nodes $\text{wit}_1(\mathcal{E}_s)$ from V_1 and replaced each node $v \in \text{wit}_2(\mathcal{E}_s)$ by $k_v + 1$ new nodes. Next we define the edge set $E(G)$ of G . Every move $W_i = \{u, v\}$ in W that corresponds to an edge (u, v) in H^* with $u \in V_1 \setminus \text{wit}_1(\mathcal{E}_s)$ and $v \in V_2$ will add an edge in G as follows:

1. If $v \notin \text{wit}_2(\mathcal{E}_s)$, then we add $(u, v^{(0)})$ to G ; and
2. Otherwise ($v \in \text{wit}_2(\mathcal{E}_s)$), letting $\mu_i \in [0 : k_v]$ be the number of moves before W_i that contain at least one node in $\text{wit}_1(v)$ (note that W_i does not contain $\text{wit}_1(v)$; actually W_i cannot contain $\text{wit}_1(\mathcal{E}_s)$), we add $(u, v^{(\mu_i)})$ to G .

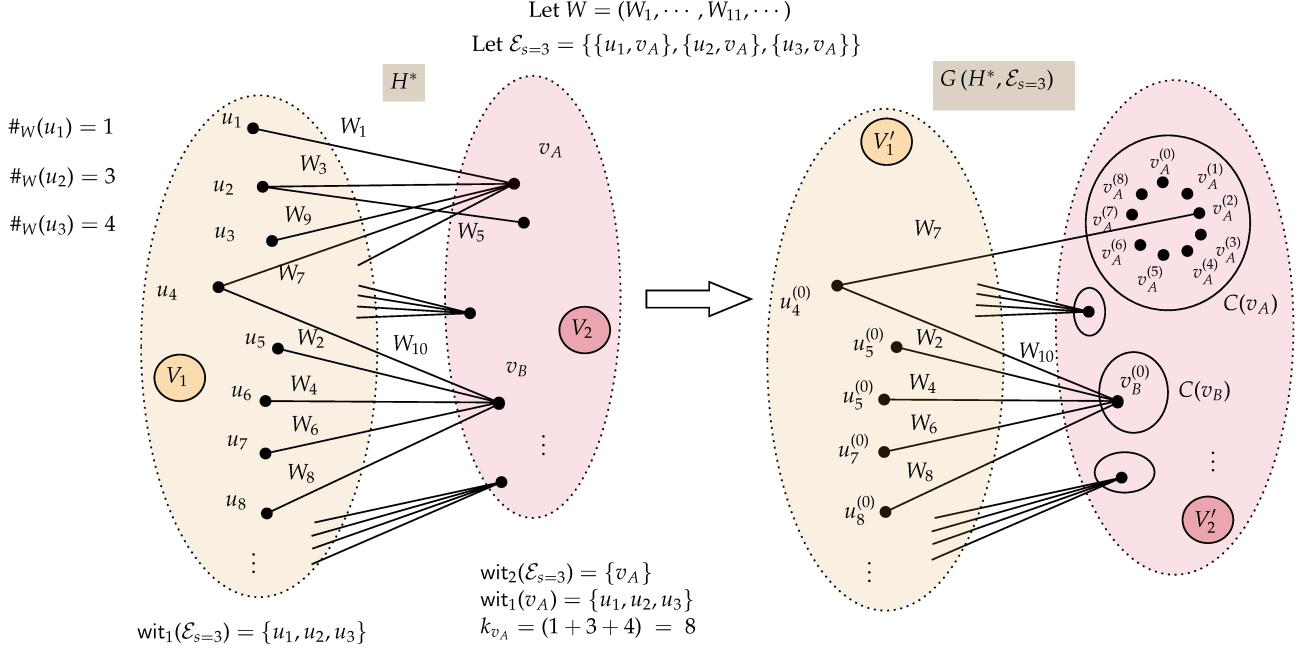


Figure 4: An exemplifying case of an auxiliary graph H^* and splitting graph $G(H^*, \mathcal{E}_s)$

Therefore, every edge in G corresponds to a move in W which corresponds to an edge in H^* that does not contain a node in $\text{wit}_1(\mathcal{E}_s)$. It is clear that each simple cycle of G corresponds to a cycle of H , which in turn corresponds to a dependent cycle of W (Since we assume w.l.o.g that all edges of auxiliary graph, and its split one, correspond to moves of the same sign (See Corollary 5.4)). So $\text{imprv}_{\tau_0, W}(C)$ is well defined for simple cycles C of G . Our motivation for constructing and working on G is because of the following lemma:

Lemma 5.7. *Let C be a simple cycle of G . Then none of the edges in \mathcal{E}_s is a witness of C .*

Proof. Let $e = \{u, v\} \in \mathcal{E}_s$ and $u \in \text{wit}_1(\mathcal{E}_s)$. By the definition of G , u has no copy in G . So u does not appear on the cycle. Let C be a cycle in G with nodes $(w_1^{(i_1)}, w_2^{(i_2)}, \dots, w_t^{(i_t)}, w_1^{(i_1)})$ (we use $w_j^{(0)}$ to denote w_j for $w_j \in V_1'$). If the cycle does not contain any vertex in $C(v)$, then $\text{imprv}_{\tau_0, W}(C)_e = 0$. Now suppose the cycle contains nodes in $C(v)$, specifically, $w_{j_1} = \dots = w_{j_m} = v$. Let the corresponding cycle C on W be c_1, \dots, c_t where $W_{c_i} = \{w_i, w_{i+1}\}$ if $i < t$ and $W_{c_t} = \{w_t, w_1\}$, and let b be the cancellation vector of C . We can write down the value of the improvement vector on edge e .

$$\begin{aligned} \text{imprv}_{\tau_0, W}(C)_e &= \sum_{k=1}^m (b_{j_k-1} \text{imprv}_{\tau_0, W}(c_{j_k-1})_e + b_{j_k} \text{imprv}_{\tau_0, W}(c_{j_k})_e) \\ &= - \sum_{k=1}^m (b_{j_k-1} \tau_{c_{j_k-1}}(v) \tau_{c_{j_k-1}-1}(u) + b_{j_k} \tau_{c_{j_k}}(v) \tau_{c_{j_k}-1}(u)) \end{aligned} \quad (22)$$

By the construction of G , u doesn't appear in any move between c_{j_k} and c_{j_k-1} (otherwise, in G , the edge corresponding to $W_{c_{j_k}}$ and edge corresponding to $W_{c_{j_k-1}}$ wouldn't be connected to $w_{j_k}^{(i_{j_k})}$ with the same i_{j_k} .(For an illustrative explanation, see Figure 5 & 6)) So $\tau_{c_{j_k-1}-1}(u) = \tau_{c_{j_k}-1}(u)$. By definition of

cancellation vector,

$$b_{j_k-1}\tau_{c_{j_k-1}}(v) + b_{j_k}\tau_{c_{j_k}}(v) = 0.$$

So each term in (22) is 0, and $\text{imprv}_{\tau_0, W}(C)_e = 0$.

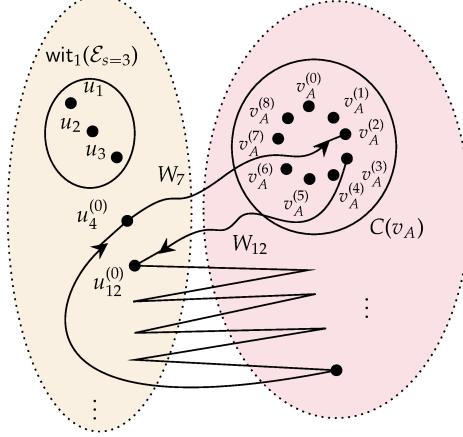


Figure 5: Following the example of Figure 7, $u_4^{(0)}$ and $u_{11}^{(0)}$ can not close cycle, since $u_3 \in W_9$ and $u_3 \in \text{wit}_1(v_A)$ and appears between W_7 and W_{12} . Thus, $u_4^{(0)}$ and $u_{11}^{(0)}$ fail to make a cycle due to the intervention of a node in $u_3 \in \text{wit}_1(v_A)$.

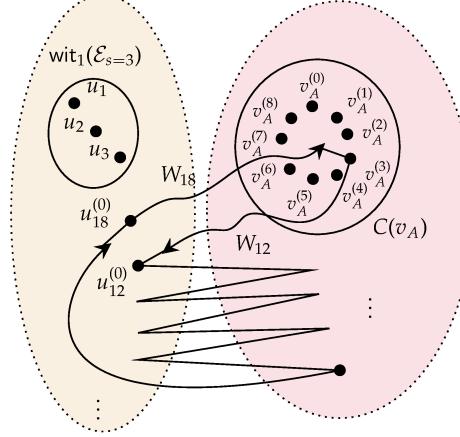


Figure 6: Again by the example of Figure 7, since $u_{12}^{(0)}$ and $u_{18}^{(0)}$ achieved to form a cycle with moves $W_{12} = \{u_{12}, v_A\}$ and $W_{18} = \{u_{18}, v_A\}$, we know from our construction that none of $\text{wit}_1(v_A)$ appear in the sub-window $W_{13} \dots W_{17}$. \square

To finish the proof, it suffices now to find a simple cycle C of G that has a witness $(u, v) \in E(W)$ with one of its vertices $u \in V'_1$. We start by checking that all conditions for H^* still hold for G . It is clear that G is a bipartite graph with no parallel edges. By the definition of $\text{wit}_1(v)$ for each $v \in \text{wit}_2(\mathcal{E}_s)$, we have $\sum_{v \in \text{wit}_2(\mathcal{E}_s)} |\text{wit}_1(v)| \leq 2s$, also $|\text{wit}_1(\mathcal{E}_s)| \leq 2s$. The number of nodes $|V'_1 \cup V'_2|$ in G is at most

$$O(L/\log^3 n) + \sum_{v \in \text{wit}_2(\mathcal{E}_s)} k_v \leq O(L/\log^3 n) + \sum_{v \in \text{wit}_2(\mathcal{E}_s)} |\text{wit}_1(v)| \cdot \log^7 n = O(L/\log^3 n).$$

where the last equality used that $s \leq L/\log^{10} n$. The number of edges in G is at least

$$\Omega(L/\log n) - |\text{wit}_1(\mathcal{E}_s)| \cdot \log^7 n = \Omega(L/\log n).$$

Let's work on another preprocessing of G to simplify the proof. Note that the average degree of nodes in G is at least $\Omega((L/\log n)/(L/\log^3 n)) = \Omega(\log^2 n)$. The following simple lemma shows that one can clean up G to get a bipartite graph G^* such that every node has degree at least $100 \log n$ and the number of edges in G^* remains to be $\Omega(L/\log n)$:

Lemma 5.8. *There is a bipartite graph $G^* = (V_1^* \cup V_2^*, E(G^*))$ with $V_1^* \subseteq V'_1$, $V_2^* \subseteq V'_2$ and $E(G^*) \subseteq E(G)$ such that every node in G^* has degree at least $100 \log n$ and $|E(G^*)| = \Omega(L/\log n)$.*

Proof. Keep deleting nodes in G with degree less than $100 \log n$ one by one (and its adjacent edges) until

no such nodes exist. The number of edges we delete during the whole process is no more than

$$(|V'_1| + |V'_2|) \cdot 100 \log n \leq O(L/\log^2 n).$$

So the remaining graph (which trivially has minimum degree at least $100 \log n$) has at least $\Omega(L/\log n)$ many edges. \square

Let us list the properties of $G^* = (V_1^* \cup V_2^*, E(G^*))$ we will use in the rest of the proof:

1. $V_1^* \subseteq V'_1 = V_1 \setminus \text{wit}_1(\mathcal{E}_s)$ and $V_2^* \subseteq V'_2$ so each node in V_2^* is in $C(v)$ for some $v \in V_2$.
2. The degree of any node is at least $100 \log n$; and
3. For any $u \in V_1^*$ and $v \in V_2$, the number of neighbors of u in $V_2^* \cap C(v)$ is at most one.
4. $E(G^*)$ has no parallel edges, $|E(G^*)| \geq \Omega(L/\log n)$ and w.l.o.g. each edge in $E(G^*)$ correspond to a move of same sign.

We prove the following lemma to finish the proof:

Lemma 5.9. *Let $u \in V_1^*$ and $v \neq v' \in V_2^*$ such that $(u, v^{(j)}), (u, v'^{(j')}) \in E(G^*)$ for some j and j' , and the corresponding moves $W_i = \{u, v\}$ and $W_{i'} = \{u, v'\}$ in W are not consecutive⁶. Then, the graph G^* has a simple cycle C such that C has a witness $e = \{u, w\} \in E(W)$ with $w \in V_1^*$.*

Proof. We begin with a simple sufficient condition for a simple cycle of G^* to satisfy the above condition.

First, let $u \in V_1^*$ and $v \neq v' \in V_2^*$ such that $(u, v^{(j)}), (u, v'^{(j')}) \in E(G^*)$ for some j and j' , and the corresponding moves $W_i = \{u, v\}$ and $W_{i'} = \{u, v'\}$ in W are not consecutive. Assume that $i < i'$ without loss of generality; then $i + 1 < i'$. The following claim shows that there must be a node $w \notin \{u, v, v'\}$ that moves an odd number times in $W_{i+1}, \dots, W_{i'-1}$:

Claim 5.10. *There is a node $w \notin \{u, v, v'\}$ that appears in an odd number of moves in $W_{i+1}, \dots, W_{i'-1}$.*

Proof. This follows from the fact that W is a valid move sequence. We distinguish two cases.

If v' appears an even number of times in $W_{i+1}, \dots, W_{i'-1}$, then use the condition of validity on the subsequence $W_i, \dots, W_{i'-1}$: there is at least one node $w \notin W_i = \{u, v\}$ that appears an odd number of times in $W_i, \dots, W_{i'-1}$. Since $w \notin W_i$, node w appears an odd number of times in $W_{i+1}, \dots, W_{i'-1}$. Hence $w \neq v'$ and thus $w \notin \{u, v, v'\}$ and the claim follows.

If v' appears an odd number of times in $W_{i+1}, \dots, W_{i'-1}$, then v' appears an even number of times in $W_i, \dots, W_{i'}$. Use the condition of validity on the subsequence $W_i, \dots, W_{i'}$: there is at least one node $w \notin W_i = \{u, v\}$ that appears an odd number of times in $W_i, \dots, W_{i'}$. Then $w \neq v'$, and since also $w \notin W_i = \{u, v\}$, it follows that w appears an odd number of times in $W_{i+1}, \dots, W_{i'-1}$ and the claim follows again. \square

⁶It is worth mentioning, that V_2^* always includes at least two vertices which are copies from different initial nodes v, v' . Indeed, if G^* was actually a star graph around $V_2^* = \{v^*\}$, then $\Omega(L/\log n) = E(G^*) = \Theta(V(G^*)) = O(L/\log^3 n)$, which leads to a contradiction. Additionally, notice that $v^{(j)}$ and $v'^{(j')}$ correspond to different nodes in the initial graph, otherwise the initial auxiliary graph H^* would have parallel edges.

We remark that Claim 5.10 holds even when W is a mixture of 1-moves and 2-moves. This will be important when we deal with the general case in Section 6.

We write $w^*(u, v^{(j)}, v'^{(j')}) \in V(W)$ to denote such a node w promised in the above claim (if more than one exist pick one arbitrarily). The next claim gives us a sufficient condition for a simple cycle C of G^* to satisfy the condition of the lemma:

Claim 5.11. *Let*

$$C = u_1 v_1^{(j_1)} u_2 v_2^{(j_2)} \cdots u_k v_k^{(j_k)} u_1$$

be a simple cycle of G^* for some nonnegative integers j_1, \dots, j_k . Suppose for some $i \in [k]$ we have that $w := w^*(u_i, v_{i-1}^{(j_{i-1})}, v_i^{(j_i)}) \in V(W)$ does not appear in C (where $v_{i-1}^{(j_{i-1})}$ denotes $v_k^{(j_k)}$ if $i = 1$), i.e., $w \notin \{u_1, \dots, u_k, v_1, \dots, v_k\}$, then $(u_i, w) \in E(W)$ must be a witness of C .

Proof. Let the corresponding cycle in W be (c_1, \dots, c_{2k}) , edge $\{u_l, v_l^{(j_l)}\}$ corresponds to move c_{2l-1} and edge $(v_l^{(j_l)}, u_{l+1})$ corresponds to c_{2l} (when $l = k$, u_{l+1} denotes u_1). Let b be its cancellation vector. Recall

$$(\text{imprv}_{\tau_0, W}(C))_{\{w, u_i\}} = \sum_{l=1}^{2k} b_l (\text{imprv}_{\tau_0, W}(c_l))_{\{w, u_i\}}.$$

Since w does not appear in C , $(\text{imprv}_{\tau_0, W}(c_l))_{\{w, u_i\}} \neq 0$ only when $u_i \in W_{c_l}$, i.e., when $l = 2i - 2$ or $l = 2i - 1$ (if $i = 1$, it is $l = 2k$ or $l = 1$). So

$$(\text{imprv}_{\tau_0, W}(C))_{\{w, u_i\}} = -b_{2i-2} \tau_{c_{2i-2}}(u_i) \tau_{c_{2i-2}}(w) - b_{2i-1} \tau_{c_{2i-1}}(u_i) \tau_{c_{2i-1}}(w).$$

By the definition of w^* , w moved odd number of times between move c_{2i-2} and c_{2i-1} . So $\tau_{c_{2i-2}}(w) = -\tau_{c_{2i-1}}(w)$. Also, by property 4 of G^* , Corollary 5.4 and the definition of a dependent cycle, $b_{2i-2} \tau_{c_{2i-2}}(u_i) + b_{2i-1} \tau_{c_{2i-1}}(u_i) = 0$. So

$$(\text{imprv}_{\tau_0, W}(C))_{\{w, u_i\}} = -2b_{2i-2} \tau_{c_{2i-2}}(u_i) \tau_{c_{2i-2}}(w) \neq 0.$$

This finishes the proof of the claim. \square

Finally we prove the existence of a simple cycle C of G^* that satisfies the condition of the above claim. To this end, we first review a simple argument which shows that any bipartite graph with n nodes and minimum degree at least $100 \log n$ must have a simple cycle. Later we modify it to our needs.

The argument goes by picking an arbitrary node in the graph as the root and growing a binary tree of $\log n$ levels as follows:

1. In the first round we just add two distinct neighbors of the root in the graph as its children.
2. Then for each round, we grow the tree by one level by going through its current leaves one by one to add two children for each leaf. For each leaf u of the current tree we just pick two of its neighbors in the graph that do not appear in ancestors of u and add them as children of u . Such neighbors always exist since the tree will have no more than $\log n$ levels and each node has degree at least $100 \log n$ in the graph.

Given that there are only n nodes in the graph, there must be a node that appears more than once in the tree at the end. Let's consider the first moment when we grow a leaf by adding one child (labelled by u) and u already appeared in the tree. Note that the two nodes labelled by u are not related in the tree since we maintain the invariant that the label of a node does not appear in its ancestors. Combining paths from these two nodes to the first node at which the two paths diverge, we get a simple cycle of the graph.

We now adapt the above argument to prove the existence of a simple cycle C of G^* that satisfies the condition of Claim 5.11 by building a binary tree of $2 \log n$ levels as follows. We start with an arbitrary node $u_{\text{root}} \in V_1^*$ as the root of the tree and expand the tree level by level, leaf by leaf, as follows:

1. Case 1: The leaf we would like to grow is labelled a node $u \in V_1^*$. In this case we add two children as follows. Let $u_1 v_1^{(j_1)} \dots u_{k-1} v_{k-1}^{(j_{k-1})} u$ be the path from the root (u_1) to u in the tree, where $u_1, \dots, u_{k-1}, u \in V_1^*$ and $v_1^{(j_1)}, \dots, v_{k-1}^{(j_{k-1})} \in V_2^*$. We pick two neighbors $v^{(j)}, v'^{(j')}$ of u in G^* with distinct $v, v' \in V_2$ as its children in the tree. We would like v and v' to satisfy the following two properties: (1) v and v' do not lie in $\{v_1, v_2, \dots, v_{k-1}\}$, (2) v and v' are different from $w^*(u_i, v_i^{(j_i)}, v_i'^{(j'_i)})$ for every $i = 1, \dots, k-1$, where $v_i'^{(j'_i)}$ denotes the other child of v_i in the tree and (3) the move corresponding to $\{u, v^{(j)}\}$ in W and the move corresponding to $\{u, v'^{(j')}\}$ in W are not consecutive moves. The existence of $v^{(j)}$ and $v'^{(j')}$ that satisfy (1), (2) and (3) follows trivially from the fact that every node (in particular, u here) has degree at least $100 \log n$ in G^* . Indeed, to satisfy (2) and (3), for each time we may reject at most 2 possible leafs. Given that the tree will only grow for $2 \log n$ levels –the half of times with V_1^* leafs and the rest half with V_2^* –, we have $k \leq \log n$ and there are at most $2k \leq 2 \log n$ edges of u that need to be avoided. Moreover, no two edges from u go to the same $C(v)$ for some $v \in V_2$ (Because we don't allow parallel edges).
2. Case 2: The leaf we would like to grow is labelled a node $v^{(j)} \in V_2^*$. In this case we just add one neighbor $u \in V_1^*$ of $v^{(j)}$ as its only child. Let $u_1 v_1^{(j_1)} \dots v_{k-1}^{(j_{k-1})} u v^{(j)}$ be the path from the root to $v^{(j)}$. We pick a neighbor $u \in V_1^*$ of $v^{(j)}$ in G^* that satisfies (1) $u \notin \{u_1, \dots, u_k\}$ and (2) u is different from $w^*(u_i, v_i^{(j_i)}, v_i'^{(j'_i)})$ for every $i = 1, \dots, k-1$ and u is different from $w^*(u_k, v^{(j)}, v'^{(j')})$, where $v_i'^{(j'_i)}$ denotes the other child of u_i and $v_k'^{(j')}$ denotes the other child of u_k in the tree. The existence of such u follows from the same argument as Case 1.

Given that the tree has $2 \log n$ levels and there are only n nodes, there must be a node that appears more than once in the tree at the end, and let's consider the first moment when we grow a leaf by adding a child and the same node already appeared in the tree. Similarly we trace the two paths and let $u \in V_1^*$ be the node where the two paths diverge; note that given the construction of the tree, this node must be a node in V_1^* , given that nodes in V_2^* only have one child in the tree. On the one hand, the way we construct the tree makes sure that combining the two paths leads to a simple cycle C of G^* . On the other hand, let $v^{(j)}, v'^{(j')} \in V_2^*$ be the two children of u (which are next to u on the cycle). Then it is easy to verify that $w^*(u, v^{(j)}, v'^{(j')})$ does not appear on the cycle we just found.

This ends the proof of the lemma. □

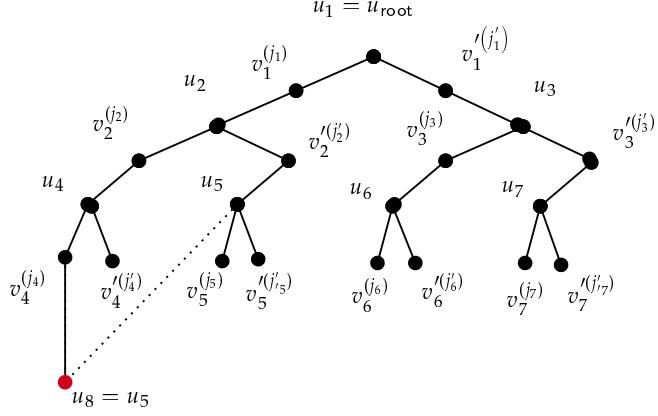


Figure 7: Example of Finding-Cycle Process.

6 General Case

We prove Lemma 3.3 for the general case. Let $\mathcal{S} = (\mathcal{S}_1, \dots, \mathcal{S}_N)$ be a valid move sequence of length $N = n \log^{10} n$ that consists of both 1-moves and 2-moves. We will consider two cases and deal with them separately: (1) the number of 1-moves in \mathcal{S} is at least $N/\log^5 n$; and (2) the number of 1-moves is at most $N/\log^5 n$.

6.1 Case 1

We consider the case when there are at least $N/\log^5 n$ many 1-moves. In this case we show that there is a window W of \mathcal{S} such that $\text{rank}_{\text{arcs}}(W)$ is large. The arguments used in this case are similar to those used in [27, 24, 13]. Given a window W of \mathcal{S} , we write $V_2(W)$ to denote the set of nodes $u \in V(W)$ such that at least two 1-moves in W are $\{u\}$.

Lemma 6.1. *There is a window W of \mathcal{S} such that*

$$|V_2(W)| = \Omega\left(\frac{\text{len}(W)}{\log^6 n}\right).$$

Proof. Any 1-move that is not the first 1-move of the vertex generates a new arc of \mathcal{S} , so the total number of arcs is at least $|\text{arcs}(\mathcal{S})| \geq N/\log^5 n - n$. Define the length of an arc α (i, j) , $\text{len}(\alpha)$, to be $j - i$. Partition all arcs based on their length, for any integer i that $0 \leq i \leq \lfloor \log_2 N \rfloor$, define

$$\text{arcs}_i(\mathcal{S}) := \left\{ \alpha : \alpha \in \text{arcs}(\mathcal{S}), \text{len}(\alpha) \in [2^i, 2^{i+1}) \right\}.$$

Since $\sum_{i=0}^{\lfloor \log_2 N \rfloor} |\text{arcs}_i(\mathcal{S})| \geq N/\log^5 n - n$, there exists i^* such that

$$|\text{arcs}_{i^*}(\mathcal{S})| \geq \frac{N/\log^5 n - n}{\log_2 N + 1} \geq \frac{N}{10 \log^6 n}.$$

Let W'_r be a window of length 2^{i^*+2} starting at a uniformly random position in $\{-2^{i^*+2} + 1, \dots, N\}$, and $W_r = W'_r \cap [N]$. For any arc $\alpha \in \text{arcs}_{i^*}(\mathcal{S})$, there are $\text{len}(W'_r) - \text{len}(\alpha)$ possible starting points for W_r

to contain α . So

$$\Pr [\alpha \in \text{arcs}(W_r)] \geq \frac{\text{len}(W_r) - \text{len}(\alpha)}{N + 2^{i^*+2} - 1} \geq \frac{2^{i^*+1}}{N + 2^{i^*+2}}.$$

From linearity of expectation, and $2^{i^*+2} \leq 4N$,

$$\mathbb{E}[\text{arcs}_{i^*}(W_r)] \geq |\text{arcs}_{i^*}(\mathcal{S})| \cdot \frac{2^{i^*+1}}{N + 2^{i^*+2}} \geq \frac{2^{i^*+1}}{50 \log^6 n} \geq \frac{\text{len}(W_r)}{100 \log^6 n}.$$

We can pick W so that $|\text{arcs}_{i^*}(W)| \geq \text{len}(W)/100 \log^6 n$. By definition of an arc, any vertex in one of the arcs in $\text{arcs}_{i^*}(W)$ must be in $V_2(W)$. On the other hand, any arc $\alpha \in \text{arcs}_{i^*}(W)$ has length at least $2^{i^*} \geq \text{len}(W)/4$. So any vertex can have at most 4 arcs in $\text{arcs}_{i^*}(W)$. We have

$$V_2(W) \geq \# \text{ vertices in } \text{arcs}_{i^*}(W) \geq |\text{arcs}_{i^*}(W)|/4 \geq \frac{\text{len}(W)}{400 \log^6 n}$$

This finishes the proof of the lemma. \square

Lemma 6.2. *We have $\text{rank}_{\text{arcs}}(W) \geq \Omega(|V_2(W)|)$.*

Proof. Let u_1, u_2, \dots, u_k be the vertices in $V_2(W)$, and $\alpha_j = (s_j, e_j)$ be the arc of u_j formed by its first and second 1-move. Since sequence W is a valid move sequence, there exists a vertex $v_j \neq u_j$ that moved odd number of times between s_j and e_j , i.e., $\tau_{s_j-1}(v_j) = -\tau_{e_j-1}(v_j)$. Pick an arbitrary such v_j for each j . Take a subset U of $V_2(W)$ by the following process:

- $V \leftarrow V_2(W)$
- $U \leftarrow \emptyset$
- For j from 1 to k , if $u_j \in V$, $V \leftarrow V \setminus \{u_j, v_j\}$, $U \leftarrow U \cup \{u_j\}$.

In each step we delete at most two element from V and add one element to U , so $|U| \geq |V_2(W)|/2$. Let $U = \{u_{i_1}, u_{i_2}, \dots, u_{i_m}\}$, ordered by the sequence they are added. By the process, for any $u_{i_j} \in U$, and any $j' > j$, $v_{i_{j'}} \neq u_{i_j}$.

Recall

$$\text{imprv}_{\tau_0, W}(\alpha_{i_j})(u_{i_j}, v_{i_j}) = -\tau_{s_{i_j}-1}(v_{i_j}) + \tau_{e_{i_j}-1}(v_{i_j}) = 2\tau_{e_{i_j}-1}(v_{i_j}) \neq 0.$$

And for any $j' > j$, $u_{i_j} \neq u_{i_{j'}}$, $v_{i_{j'}} \neq u_{i_j}$, so $\text{imprv}_{\tau_0, W}(\alpha_{i_{j'}})(u_{i_{j'}}, v_{i_{j'}}) = 0$. Consider the matrix formed by taking the j -th column to be $\text{imprv}_{\tau_0, W}(\alpha_{i_j})$. The row indexed by (u_{i_j}, v_{i_j}) would be of the form

$$\left(\underbrace{*, *, \dots, *}_{j-1 \text{ unknown numbers}}, \quad 2\tau_{e_{i_j}-1}(v_{i_j}) \neq 0, 0, \dots, 0 \right).$$

This means the matrix has a lower triangular square submatrix of size at least $m \geq |V_2(W)|/2$. So we have $\text{rank}_{\text{arcs}}(W) \geq |V_2(W)|/2$. \square

6.2 Case 2

Let \mathcal{S} be a valid move sequence of length N with no more than $N/\log^5 n$ many 1-moves. Let W be a window of \mathcal{S} . We write $\#_W(u)$ to denote the number of moves (including both 1-moves and 2-moves)

that u appears in W , and write $\#_W^2(u)$ to denote the number of 2-moves that u appears in W .

We start by showing a lemma similar to Lemma 4.6 in Section 4.

Lemma 6.3. *Let \mathcal{S} be a valid move sequence of length $N = n \log^{10} n$ with no more than $N/\log^5 n$ many 1-moves. Then there exists a window W of \mathcal{S} such that at least $\Omega(\text{len}(W)/\log n)$ many moves of W are 2-moves $W_i = \{u, v\}$ that satisfy*

$$\log^3 n \leq \#_W^2(u) \leq \#_W(u) \leq 2\log^7 n \quad \text{and} \quad \#_W^2(v) \geq \log^3 n \quad (23)$$

Proof. We would like to apply Lemma 4.6 (which works on move sequences that consist of 2-moves only). To this end, we let \mathcal{S}' be the move sequence obtained from \mathcal{S} by removing all its 1-moves. Let $N' := \text{len}(\mathcal{S}') \geq (1 - 1/\log^5 n)N$. Applying Lemma 4.6 on \mathcal{S}' ,⁷ there must exist a positive integer L and among the $N' - L + 1$ windows W' of \mathcal{S}' of length L , at least $\Omega(1/\log n)$ -fraction of them satisfy that $\Omega(L/\log n)$ many 2-moves $\{u, v\}$ in it satisfy

$$\log^3 n \leq \#_{W'}^2(u) \leq \log^7 n \quad \text{and} \quad \#_{W'}^2(v) \geq \log^3 n. \quad (24)$$

Let's denote these windows of \mathcal{S}' by W'_1, \dots, W'_s for some $s = \Omega((N' - L + 1)/\log n)$. For each W'_i we let \mathcal{S}_{k_i} (or \mathcal{S}_{ℓ_i}) to denote the move in \mathcal{S} that corresponds to the first (or last, respectively) move in W'_i , and let W_i denote the window $(\mathcal{S}_{k_i}, \dots, \mathcal{S}_{\ell_i})$ of \mathcal{S} .

If $L \geq N'/2$, we can just take W to be W_1 . We note that the number of 2-moves in W_1 that satisfy (24) is at least $\Omega(L/\log n) = \Omega(\text{len}(W_1)/\log n)$ given that $L \geq N'/2$. On the other hand, the number of $u \in V(W_1)$ that appears in at least $\log^7 n$ 1-moves of W_1 is at most $(N/\log^5 n)/\log^7 n = O(N/\log^{12} n)$. Thus, the number of 2-moves $\{u, v\}$ in W_1 that satisfy (24) but not $\#_{W_1}(u) \leq 2\log^7 n$ is at most

$$\log^7 n \cdot O\left(\frac{N}{\log^{12} n}\right) = o\left(\frac{L}{\log n}\right).$$

So we assume below that $L < N'/2$.

We claim that W_i can satisfy the condition of the lemma if $\text{len}(W_i) \leq (1 + 1/\log^2 n)L$. To see this is the case, we note that the number of nodes $u \in W_i$ that appears in at least $\log^7 n$ many 1-moves is at most $(L/\log^2 n)/\log^7 n = O(L/\log^9 n)$. Thus, the number of 2-moves $\{u, v\}$ in W_i that satisfy (24) but not $\#_{W_i}(u) \leq 2\log^7 n$ is at most

$$\log^7 n \cdot O\left(\frac{N}{\log^9 n}\right) = o\left(\frac{L}{\log n}\right).$$

So it suffices to show that $\text{len}(W_i) \leq (1 + 1/\log^2 n) \cdot \text{len}(W'_i)$ for some window W_i . Assume this is not the case. Then the total number of 1-moves in W_1, \dots, W_s is at least

$$\frac{1}{\log^2 n} \cdot \sum_{i \in [s]} \text{len}(W'_i) \geq \Omega\left(\frac{(N' - L + 1)L}{\log^3 n}\right) = \Omega\left(\frac{NL}{\log^3 n}\right)$$

using $L < N'/2$. However, each 1-move can only appear in no more than L many windows of length L .

⁷Note that \mathcal{S}' has length not exactly N but $(1 - 1/\log^5 n)N$ but the statement of Lemma 4.6 still holds.

Given that there are only $N/\log^5 n$ many 1-moves, the same number can be upper bounded by

$$\frac{N}{\log^5 n} \cdot L,$$

a contradiction. This finishes the proof of the lemma. \square

So we now have a valid move sequence W of length L , as a window of the original valid sequence S , such that the number of 2-moves in W that satisfy (23) is at least $\Omega(L/\log n)$. The rest of the proof follows the same arguments used in Section 5. We give a sketch below.

First we define the same auxiliary graph $H = (V(W), E)$ such that there is a one-to-one correspondence between E and 2-moves in W . Note that the degree of a node u in H is the same as $\#_W^2(u)$.

We then show that there are disjoint sets of nodes $V_1, V_2 \subset V(W)$ and a subset of edges $E' \subseteq E$ that satisfy conditions similar to those of Lemma 5.1:

Lemma 6.4. *There are two disjoint sets of nodes $V_1, V_2 \subset V(W)$ and a subset of edges $E' \subseteq E$ such that*

1. *Every edge in E' has one node in V_1 and the other node in V_2 ;*
2. *$|V_1 \cup V_2| = O(L/\log^3 n)$ and $|E'| = \Omega(L/\log n)$;*
3. *$\#_W(u) \leq 2\log^7 n$ for every node $u \in V_1$.*

The proof is exactly the same as that of Lemma 5.1, except that we define V to be the set of nodes v with $\#_W^2(v) \geq \log^3 n$ and V_h to be the set of nodes v with $\#_W(v) \geq 2\log^7 n$.

Next we focus on E'' , which contains all edges in E' of the same sign (or edges in E' of different signs, whichever contains more edges). Similarly every cycle in E'' corresponds to a dependent cycle of W . The case when E'' contains many parallel edges can be handled exactly the same way as in Lemma 5.5. So we may delete all parallel edges from E'' , and finish the proof using Lemma 5.6.

The proof of Lemma 5.6 for the general case is very similar, with the following changes:

1. In the definition of k_v for each $v \in \text{wit}_2(\mathcal{E}_s)$, we need it to be the number of moves (including both 1-moves and 2-moves) in W that involve at least one node in $\text{wit}_1(v)$. This can still be bounded from above by $|\text{wit}_1(v)| \cdot 2\log^7 n$ since we have $\#_W(u) \leq \log^7 n$ for all $u \in V_1$ as promised in Lemma 6.4 above.
2. As we commented earlier, Claim 5.10 works even when W consists of both 1-moves and 2-moves.

This finishes the proof of Lemma 3.3 for the general case.

7 Binary Max-CSP and Function Optimization Problems

We recall the definition of binary maximum constraint satisfaction problems, and more generally function optimization problems.

Definition 7.1. *An instance of Binary Max-CSP (Constraint Satisfaction Problem), or MAX 2-CSP, consists of a set $V = \{x_1, \dots, x_n\}$ of variables that can take values over $\{0, 1\}$ and a set $C = \{c_1, \dots, c_m\}$ of constraints with given respective weights w_1, \dots, w_m , where each constraint is a predicate on a pair of variables. The MAX 2-CSP problem is: given an instance, find an assignment that maximizes the sum of the weights of the satisfied constraints.*

Several problems can be viewed as special cases of Binary Max-CSP where the predicates of the constraints are restricted to belong to a fixed family \mathcal{P} of predicates; this restricted version is denoted $\text{Max-CSP}(\mathcal{P})$. For example, the Max Cut problem in graphs is equivalent to $\text{Max-CSP}(\mathcal{P})$ where \mathcal{P} contains only the “not-equal” predicate ($x \neq y$, where x, y are the two variables). The Max Directed Cut problem, where the input graph is directed and we seek a partition of the nodes into two parts N_1, N_2 that maximizes the total weight of the edges directed from N_1 to N_2 , corresponds to the case that \mathcal{P} contains only the $<$ predicate (i.e. $x < y$). MAX 2SAT corresponds to the case that \mathcal{P} consists of all 4 possible clauses on two variables.

A generalization of MAX 2-CSP is the class of *Binary function optimization problems* (BFOP) where instead of constraints (predicates) we have functions on two arguments that take values in $\{0, 1, \dots, d\}$ instead of $\{0, 1\}$, where d is a fixed constant (or even is polynomially bounded). For convenience and consistency with the notation of configurations in the Max Cut problem, we will use in the following $\{-1, 1\}$ as the domain of the variables instead of $\{0, 1\}$. That is, the problem is: Given a set $V = \{x_1, \dots, x_n\}$ of variables with domain $D = \{-1, 1\}$, a set $F = \{f_1, \dots, f_m\}$ of functions, where each f_i is a function of a pair (x_{i_1}, x_{i_2}) of variables, and given respective weights w_1, \dots, w_m , find an assignment $\tau : V \rightarrow D$ to the variables that maximizes $\sum_{i=1}^m w_i \cdot f_i(\tau(x_{i_1}), \tau(x_{i_2}))$.

Even though a function in BFOP (or a constraint in Max-2CSP) has two arguments, its value may depend on only one of them, i.e. it may be essentially a unary function (or constraint). More generally, it may be that the two arguments of the function can be decoupled and the function can be separated into two unary functions. We say that a binary function $f(x, y)$ is *separable* if there are unary functions f_1, f_2 such that $f(x, y) = f_1(x) + f_2(y)$ for all values of x, y ; otherwise f is *nonseparable*. For binary domains there is a simple criterion for separability: a function $f(x, y)$ is separable if and only if $f(-1, -1) + f(1, 1) = f(-1, 1) + f(1, -1)$ [25]. If in a given BFOP instance some binary functions are separable, then we can decompose them into the equivalent unary functions. Thus, we may assume, without loss of generality, that a given BFOP instance has unary and binary functions, where all the binary functions are nonseparable. We say that an instance is *complete*, if every pair of variables appear as the arguments of a (nonseparable) binary function in the instance.

The 2-FLIP local search algorithm can be applied to a MAX 2-CSP or BFOP problem to compute a locally optimal assignment that cannot be improved by flipping the value of any one or two variables. We will show that the smoothed complexity of 2-FLIP for any complete MAX 2-CSP or BFOP instance is (at most) quasipolynomial.

Theorem 1.3. *Let I be an arbitrary complete instance of a MAX 2-CSP (or BFOP) problem with n variables*

and m constraints (functions) with independent random weights in $[-1, 1]$ with density at most $\phi > 0$. Then, with probability at least $1 - o_n(1)$ over the draw of the weights, any implementation of 2-FLIP takes at most $m\phi n^{O(\log^{10} n)}$ steps to terminate.

Proof. Consider a (complete) instance I of a BFOP problem with n variables and m functions, and a sequence \mathcal{S} of moves of 2-FLIP starting from an initial configuration. The proof follows the same structure as the proof for Max Cut. The only thing that changes is the improvement vector in each step, which depends on the specific functions of the instance: the vector has one coordinate for each function f_i in the instance and the entry is equal to the change in the value of the function resulting from the move. Arcs and cycles of \mathcal{S} are defined in the same way as in Max Cut, and the improvement vectors of arcs and cycles are defined in an analogous way from the improvement vectors of the moves.

The heart of the proof for Max Cut is Lemma 3.3 which showed that there is a window \mathcal{W} and a set of arcs or a set of cycles of \mathcal{W} whose improvement vectors have rank $\Omega(\frac{\log(\mathcal{W})}{\log^{10} n})$. We will show that the lemma holds for any BFOP problem.

We associate with the BFOP instance I the graph G where the nodes correspond to the variables of I and the edges correspond to the binary functions of I ; since I is a complete instance, the graph G is the complete graph, possibly with multiple edges connecting the same pair of nodes (if there are multiple functions with the same pair of arguments). We will identify the variables of I with the nodes of G and the functions of I with the edges of G .

In the general case of the Max Cut problem, in Case 1 where there is a large number of 1-moves, we identified a window \mathcal{W} and a large set A' of arcs in the window whose set of improvement vectors are linearly independent. The argument relied only on the zero-nonzero structure of the improvement vectors: it showed that the matrix M formed by these vectors and a set of rows corresponding to a certain set E' of witness edges is a lower triangular matrix with nonzero diagonal. Take a set F' of functions of I that contains for each edge $\{u, v\} \in E'$ a function $f_k(u, v)$ with this pair as arguments (it exists because the instance I is complete), and form the matrix M' with the set F' as rows and the set A' of arcs as columns. We will show that the matrix M' has the same zero-nonzero structure as M , thus it also has full rank.

Consider an arc of the move sequence \mathcal{S} corresponding to two moves $\mathcal{S}_i = \{u\}$, $\mathcal{S}_j = \{u\}$, $i < j$, and a function f_k of I . If u is not one of the arguments of the function, then the corresponding entry of the improvement vector of the arc is obviously 0. If u is one of the arguments, i.e. the k -th function is $f_k(u, v)$ (similarly if it is $f_k(v, u)$), then the corresponding entry of the improvement vector of the arc is $\gamma_i(u)[f_k(\gamma_i(u), \gamma_i(v)) - f_k(-\gamma_i(u), \gamma_i(v))] - \gamma_j(u)[f_k(\gamma_j(u), \gamma_j(v)) - f_k(-\gamma_j(u), \gamma_j(v))]$. If v moves an even number of times between \mathcal{S}_i and \mathcal{S}_j , then $\gamma_i(v) = \gamma_j(v)$ and it follows that the entry is 0, both in the case that $\gamma_i(u) = \gamma_j(u)$ and in the case that $\gamma_i(u) = -\gamma_j(u)$. On the other hand, if v moves an odd number of times between \mathcal{S}_i and \mathcal{S}_j , then $\gamma_i(v) = -\gamma_j(v)$ and it follows that the k -th entry of the improvement vector is $\gamma_i(u)[f_k(\gamma_i(u), \gamma_i(v)) - f_k(-\gamma_i(u), \gamma_i(v))] - \gamma_j(u)[f_k(\gamma_j(u), -\gamma_i(v)) - f_k(-\gamma_j(u), -\gamma_i(v))]$. Letting $\gamma_i(u) = a$, $\gamma_i(v) = b$, the entry is $a[f_k(a, b) + f_k(-a, -b) - f_k(-a, b) - f_k(a, -b)]$ (both when $\gamma_i(u) = \gamma_j(u)$ and when $\gamma_i(u) = -\gamma_j(u)$); this quantity is nonzero because f_k is nonseparable. Thus, the entry for $f_k(u, v)$ of the improvement vector of the arc is nonzero exactly when the entry of the arc in the Max Cut problem for the edge (u, v) is nonzero. It follows that the matrix M' has the same zero-nonzero structure as M , thus it also has full rank.

In Case 2 of the Max Cut problem, where the number of 2-moves is very large, there were two

subcases. In the first subcase, where there are many parallel edges in the graph that we associated with the window of the move sequence, we found a large set of 2-cycles whose improvement vectors were linearly independent. In the other case, where there are "few" parallel edges, we constructed a large set of cycles (of length $O(\log n)$), again with linearly independent improvement vectors. In both cases, the proof of linear independence relied again only on the zero-nonzero structure of the vectors, and not on the precise value of the entries. We will argue that in both cases, the corresponding vectors of these cycles in the BFOP instance I have the same zero-nonzero structure.

In the first subcase we found many 2-cycles $(u_1, v_1), \dots, (u_k, v_k)$, and corresponding "witness" edges $(u_1, z_1), \dots, (u_k, z_k)$ such that the matrix M with rows corresponding to the witness edges and columns corresponding to the 2-cycles in the Max Cut problem is lower triangular with nonzero diagonal. The nodes u_i are distinct (the v_i and the z_i may not be distinct) and $z_i \neq u_i, v_i$ for all i . For each witness pair (u_i, z_i) pick a function f_{r_i} of instance I with this pair of variables as arguments, in either order, say wlog the function is $f_{r_i}(u_i, z_i)$. Consider the matrix M' with rows corresponding to the functions $f_{r_i}(u_i, z_i)$ and columns corresponding to the 2-cycles (u_i, v_i) . Note that the entry $M(j, i)$ is nonzero if one of the nodes u_j, z_j is in $\{u_i, v_i\}$ and the other node appears an odd number of times between the two moves $\{u_i, v_i\}$, and it is 0 otherwise, i.e. if $\{u_j, z_j\} \cap \{u_i, v_i\} = \emptyset$, or if one of u_j, z_j is in $\{u_i, v_i\}$ and the other node appears an even number of times between the two moves $\{u_i, v_i\}$. Importantly it cannot be that $\{u_j, z_j\} = \{u_i, v_i\}$ because $u_j \neq u_i, v_i$. Examining the value $M'(j, i)$ in the same way as in the case of arcs above, we observe that if $M(j, i) = 0$ then also $M'(j, i) = 0$, and if $M(j, i) \neq 0$ then also $M'(j, i) \neq 0$. Thus, M' has the same zero-nonzero structure as M and hence it has also full rank.

In the second subcase of Case 2, we found many cycles C_1, \dots, C_k and corresponding witness edges $\{u_i, v_i\}$ such that for every i , (1) C_i does not contain any u_j for $j < i$, nor v_i , (2) C_i has exactly two edges incident to u_i and node v_i appears an odd number of times between the two moves corresponding to these two edges, (3) if C_i contains v_j for some $j < i$ (the cycle C_i may go more than once through v_j), then u_j does not appear between any pair of moves that correspond to consecutive edges of the cycle C_i incident to v_i . We used these properties in Max Cut to show that the matrix M whose rows correspond to the witness edges and the columns correspond to the cycles C_i is lower triangular with nonzero diagonal. As before, for each witness pair (u_i, v_i) pick a function f_{r_i} of instance I with this pair of variables as arguments, and let M' be the matrix with these functions as rows and the cycles C_i as columns. We can use the above properties to show that the matrix M' is also lower triangular with nonzero diagonal. Property (2) and the fact that $v_i \notin C_i$ (from property (1)) imply that $M'(i, i) \neq 0$ for all i . Properties (1) and (3) can be used to show that $M'(j, i) = 0$ for all $j < i$. Therefore, M' has full rank.

Once we have Lemma 3.3 for the BFOP instance I , the rest of the proof is the same as for Max Cut. The only difference is that, if the maximum value of a function in I is d (a constant, or even polynomial in n), then the maximum absolute value of the objective function is md instead of n^2 that it was in Max Cut. \square

8 Conclusions

We analyzed the smoothed complexity of the SWAP algorithm for Graph Partitioning and the 2-FLIP algorithm for Max Cut and showed that with high probability the algorithms terminate in quasi-polynomial time for any pivoting rule. The same result holds more generally for the class of maximum binary constraint satisfaction problems (like Max-2SAT, Max Directed Cut, and others). We have not made any attempt currently to optimize the exponent of $\log n$ in the bound, but we believe that with a more careful analysis the true exponent will be low. There are several interesting open questions raised by this work. We list some of them below.

1. Can our bounds be improved to polynomial? In the case of the 1-FLIP algorithm in the full perturbation model (i.e. when all edges of K_n are perturbed) a polynomial bound was proved in [23]. Can a similar result be shown for 2-FLIP and SWAP?
2. Can our results be extended to the structured smoothed model, i.e., when we are given a graph G and only the edges of G are perturbed? In the case of 1-FLIP we know that this holds [22, 25], but 2-FLIP is much more challenging.
3. We saw in this paper how to analyze local search when one move flips simultaneously two nodes. This is a qualitative step up from the case of single flips, that creates a number of obstacles which had to be addressed. This involved the introduction of nontrivial new techniques in the analysis of the sequence of moves, going from sets to graphs. Dealing with local search that flips 3 or more nodes will require extending the methods further to deal with hypergraphs. We hope that our techniques will form the basis for handling local search algorithms that flip multiple nodes in one move, e.g. k -FLIP for higher k , and even more ambitiously powerful methods like Kernighan-Lin that perform a deep search in each iteration and flip/swap an unbounded number of nodes.
4. Can our results be extended to Max k -Cut or k -Graph Partitioning where the graph is partitioned into $k > 2$ parts? In the case of 1-FLIP for Max k -Cut quasi-polynomial bounds were shown in [24].
5. Can similar results be shown for Max-CSP with constraints of higher arities, for example Max 3SAT? No bounds are known even for 1-FLIP. In fact, analyzing 1-FLIP for Max 3SAT seems to present challenges that have similarities with those encountered in the analysis of 2-FLIP for Max 2SAT and Max Cut, so it is possible that the techniques developed in this paper will be useful also in addressing this problem.

A Missing Proofs from Section 2

Lemma A.1. For any configuration $\gamma_0 \in \{\pm 1\}^n$ that is an extension of τ_0 , letting $\gamma_0, \gamma_1, \dots, \gamma_\ell \in \{\pm 1\}^n$ be the sequence of configurations induced by \mathcal{S} and letting $w[u, i, j] := \gamma_i(u) \cdot \text{imprv}_{\gamma_0, \mathcal{S}}(i) - \gamma_j(u) \cdot \text{imprv}_{\gamma_0, \mathcal{S}}(j)$, we have that

$$(w[u, i, j]_e)_{e \in E_n} = \begin{cases} (\tau_i(u) \cdot \text{imprv}_{\tau_0, \mathcal{S}}(i) - \tau_j(u) \cdot \text{imprv}_{\tau_0, \mathcal{S}}(j))_e & \text{for every entry } e \in E(\mathcal{S}), \\ 0 & \text{otherwise.} \end{cases}$$

for any arbitrary choice of $u \in V(\mathcal{S})$.

Proof. Note that $\tau_i(u) = \gamma_i(u)$, $\tau_j(u) = \gamma_j(u)$, and by definition, $\text{imprv}_{\tau_0, \mathcal{S}}(i) = \text{imprv}_{\gamma_0, \mathcal{S}}(i)^*$ is the projection of $\text{imprv}_{\gamma_0, \mathcal{S}}(i)$ on $E(\mathcal{S})$.

Thus, for any edge $e \in E(\mathcal{S})$, w_e is the same as $(\tau_i(u) \cdot \text{imprv}_{\tau_0, \mathcal{S}}(i) - \tau_j(u) \cdot \text{imprv}_{\tau_0, \mathcal{S}}(j))_e$. For any edge $e \notin E(\mathcal{S})$, if $e = \{v_1, v_2\}$ doesn't contain u , the improvement vectors are 0 on e and correspondingly $w[u, i, j]_e = 0$. For the last case, let us assume $e = \{u, v\}$ where v is inactive. We have that

$$\text{imprv}_{\gamma_0, \mathcal{S}}(i)_e = \gamma_{i-1}(v)\gamma_{i-1}(u) = -\gamma_{i-1}(v)\gamma_i(u) \quad \& \quad \text{imprv}_{\gamma_0, \mathcal{S}}(j)_e = -\gamma_{j-1}(v)\gamma_j(u).$$

Since v is not active, $\gamma_i(v) = \gamma_0(v)$ for any $i \in [\ell]$. So, we get that

$$(\gamma_i(u) \cdot \text{imprv}_{\gamma_0, \mathcal{S}}(i) - \gamma_j(u) \cdot \text{imprv}_{\gamma_0, \mathcal{S}}(j))_e = -\gamma_{i-1}(v) + \gamma_{j-1}(v) = 0.$$

This finishes the proof of the lemma. \square

Lemma A.2. Let $C = (c_1, \dots, c_t)$ be a dependent cycle of \mathcal{S} and let b be its cancellation vector. Then for any configurations $\tau_0 \in \{\pm 1\}^{V(\mathcal{S})}$ and $\gamma_0 \in \{\pm 1\}^n$ such that γ_0 is an extension of τ_0 , letting $w[C] := \sum_{j \in [t]} b_j \cdot \text{imprv}_{\gamma_0, \mathcal{S}}(c_j)$, we have that

$$(w[C]_e)_{e \in E_n} = \begin{cases} (\sum_{j \in [t]} b_j \cdot \text{imprv}_{\tau_0, \mathcal{S}}(c_j))_e & \text{for every entry } e \in E(\mathcal{S}), \\ 0 & \text{otherwise.} \end{cases}$$

Proof. Again recall that $\tau_i(u) = \gamma_i(u)$, $\tau_j(u) = \gamma_j(u)$, and by definition, $\text{imprv}_{\tau_0, \mathcal{S}}(i) = \text{imprv}_{\gamma_0, \mathcal{S}}(i)^*$ is the projection of $\text{imprv}_{\gamma_0, \mathcal{S}}(i)$ on $E(\mathcal{S})$.

Thus, for edge $e \in E(\mathcal{S})$, $w[C]_e$ is the same as $\sum_{j \in [t]} b_j \cdot \text{imprv}_{\tau_0, \mathcal{S}}(c_j)_e$. For edge $e \notin E(\mathcal{S})$, if e doesn't contain u_i for any $i \in [t]$, the improvement vectors are 0 on e and correspondingly $w[C]_e = 0$. For the last case, let us assume $e = (u, v)$ where v is inactive. Then we have (where index 0 corresponds to t)

$$\begin{aligned} \sum_{j \in [t]} b_j \cdot \text{imprv}_{\gamma_0, \mathcal{S}}(c_j)_e &= \sum_{i \in [t]: u_i = u} (b_{i-1} \text{imprv}_{\gamma_0, \mathcal{S}}(c_{i-1})_e + b_i \text{imprv}_{\gamma_0, \mathcal{S}}(c_i)_e) \\ &= \sum_{i \in [t]: u_i = u} (-b_{i-1} \gamma_{c_{i-1}-1}(v) \gamma_{c_{i-1}}(u_i) - b_i \gamma_{c_i-1}(v) \gamma_{c_i}(u_i)). \end{aligned}$$

Since v is inactive, $\gamma_{c_{i-1}-1}(v) = \gamma_{c_i-1}(v) = \gamma_0(v)$. Each term above is equal to

$$-\gamma_0(v)(b_{i-1}\gamma_{c_{i-1}}(u_i) + b_i\gamma_{c_i}(u_i)) = 0.$$

This finishes the proof of the lemma. \square

B Rank Invariance of Improving vectors over Initial Configuration

In this section we prove that the rank of the improvement vectors for the set of $1\text{-move}(\mathcal{S})$, $2\text{-move}(\mathcal{S})$, $\text{arcs}(\mathcal{S})$ and $\text{cycles}(\mathcal{S})$ is independent of the initial configuration γ_0 of the vertices.

Proof of Lemmas 2.4, 2.8 & 2.10. We start by recalling the following useful facts:

Fact B.1. Let $\gamma_0, \gamma'_0 \in \{\pm 1\}^n$ be two arbitrary initial configurations. Then, it holds that

$$\gamma_0(v)\gamma'_0(v) = \gamma_i(v)\gamma'_i(v) \quad \text{for any } i \in [\ell],$$

where γ_i, γ'_i is obtained from $\gamma_{i-1}, \gamma'_{i-1}$ by flipping nodes in \mathcal{S}_i , for a move sequence $\mathcal{S} = (\mathcal{S}_1, \dots, \mathcal{S}_\ell)$.

Fact B.2. Let A be a $(k_1 \times k_2)$ real-valued matrix A and B, C are full-rank $(k_1 \times k_1)$ and $(k_2 \times k_2)$ squared matrices correspondingly. Then, it holds that $\text{rank}(A) = \text{rank}(A^\top)$ and $\text{rank}(A) = \text{rank}(BAC)$.

Let $\mathcal{M}_{1\text{-move}}(\gamma_0, \mathcal{S}), \mathcal{M}_{2\text{-move}}(\gamma_0, \mathcal{S})$ be the matrices whose columns are the improvement vectors of 1-move for a given initial configuration γ_0 and let $\mathcal{M}_{1\text{-move}}(\tau_0, \mathcal{S}), \mathcal{M}_{2\text{-move}}(\tau_0, \mathcal{S})$ be the submatrices of $\mathcal{M}_{1\text{-move}}(\tau_0, \mathcal{S}), \mathcal{M}_{2\text{-move}}(\tau_0, \mathcal{S})$ including only the rows which correspond to $E(\mathcal{S})$.

Schematically, we have that for the 1-move case:

$$\mathcal{M}_{1\text{-move}}(\tau_0, \mathcal{S}) = \begin{bmatrix} & & \vdots & & \\ \cdots & (\text{imprv}_{\tau_0, \mathcal{S}}(\mathcal{S}_k = \{u\})_{e=\{u,v\}} = \tau_{k-1}(u)\tau_{k-1}(v) & \cdots & & \\ & & \vdots & & \\ \cdots & (\text{imprv}_{\tau_0, \mathcal{S}}(\mathcal{S}_k = \{u\})_{e=\{z,w\}} = 0 & \cdots & & \\ & & \vdots & & \\ & & & & \end{bmatrix}_{|E(\mathcal{S})|, |1\text{-move}(\mathcal{S})|}$$

and for the 2-move case:

$$\mathcal{M}_{2\text{-move}}(\tau_0, \mathcal{S}) = \begin{bmatrix} & & \vdots & & \\ \cdots & (\text{imprv}_{\tau_0, \mathcal{S}}(\mathcal{S}_{k'} = \{u, v\})_{e=\{u,v\}} = 0 & \cdots & & \\ \cdots & (\text{imprv}_{\tau_0, \mathcal{S}}(\mathcal{S}_{k'} = \{u, v\})_{e=\{u,z\}} = \tau_{k'-1}(u)\tau_{k'-1}(z) & \cdots & & \\ \cdots & (\text{imprv}_{\tau_0, \mathcal{S}}(\mathcal{S}_{k'} = \{u, v\})_{e=\{v,w\}} = \tau_{k'-1}(v)\tau_{k'-1}(w) & \cdots & & \\ & & \vdots & & \\ \cdots & (\text{imprv}_{\tau_0, \mathcal{S}}(\mathcal{S}_{k'} = \{u, v\})_{e=\{z,w\}} = 0 & \cdots & & \\ & & \vdots & & \\ & & & & \end{bmatrix}_{|E(\mathcal{S})|, |2\text{-move}(\mathcal{S})|}$$

Notice that we can derive $\mathcal{M}_{1\text{-move}}(\tau'_0, \mathcal{S}), \mathcal{M}_{2\text{-move}}(\tau'_0, \mathcal{S})$ by multiplying $\mathcal{M}_{1\text{-move}}(\tau_0, \mathcal{S}), \mathcal{M}_{2\text{-move}}(\tau_0, \mathcal{S})$ from left by the squared diagonal $|E_n| \times |E_n|$ matrix $\mathcal{D}[\tau_0, \tau'_0]$ such that

$$\mathcal{D}[\tau_0, \tau'_0]_{(e,e)=((u,v),(u,v))} = \tau_0(u)\tau'_0(u)\tau_0(v)\tau'_0(v) \text{ and } \mathcal{D}[\tau_0, \tau'_0]_{(e,e')} = 0 \text{ for } e \neq e'.$$

Indeed, we have that for an entry (e, k) representing an edge $e = (u, v)$ and the k -th μ -move

$$\begin{aligned} (\mathcal{D}[\tau_0, \tau'_0]\mathcal{M}_{\mu\text{-move}}(\tau_0, \mathcal{S}))_{e,k} &= \begin{cases} \tau_0(u)\tau'_0(u)\tau_0(v)\tau'_0(v)\tau_k(u)\tau_k(v) & \text{if } \mathcal{M}_{\mu\text{-move}}(\tau_0, \mathcal{S})_{e,k} = \tau_k(u)\tau_k(v) \\ 0 & \text{if } \mathcal{M}_{\mu\text{-move}}(\tau_0, \mathcal{S})_{e,k} = 0 \end{cases} \\ \stackrel{\text{Fact B.1}}{=} & \begin{cases} \tau_k(u)\tau'_k(u)\tau_k(v)\tau'_k(v)\tau_k(u)\tau_k(v) & \text{if } \mathcal{M}_{\mu\text{-move}}(\tau_0, \mathcal{S})_{e,k} = \tau_k(u)\tau_k(v) \\ 0 & \text{if } \mathcal{M}_{\mu\text{-move}}(\tau_0, \mathcal{S})_{e,k} = 0 \end{cases} \end{aligned}$$

Since $\tau_i(v)^2$ equals 1 for any $v \in V_n$ for every $i \in [\ell]$, we get that

$$(\mathcal{D}[\tau_0, \tau'_0]\mathcal{M}_{\mu\text{-move}}(\tau_0, \mathcal{S}))_{e,k} = \begin{cases} \tau'_k(u)\tau'_k(v) & \text{if } \mathcal{M}_{\mu\text{-move}}(\tau_0, \mathcal{S})_{e,k} = \tau_k(u)\tau_k(v) \\ 0 & \text{if } \mathcal{M}_{\mu\text{-move}}(\tau_0, \mathcal{S})_{e,k} = 0 \end{cases} = (\mathcal{M}_{\mu\text{-move}}(\tau'_0, \mathcal{S}))_{e,k}$$

for any $\mu \in \{1, 2\}$. Since $\mathcal{D}[\tau_0, \tau'_0]$ is a full-rank matrix, leveraging Fact B.2, the above argument proves that $\text{rank}(\mathcal{M}_{\mu\text{-move}}(\tau_0, \mathcal{S}))$ is independent of the initial configuration.

More interestingly, in order to prove Lemma 2.4, it suffices to prove that the rank of the column matrix with the improvement vectors of arcs is independent of the initial configuration. In fact, let

$$M_{\text{arcs}(\mathcal{S})}(\tau_0, \mathcal{S}) = \begin{bmatrix} \vdots & & \vdots & & \vdots \\ \text{imprv}_{\tau_0, \mathcal{S}}(\alpha_1) & \cdots & \text{imprv}_{\tau_0, \mathcal{S}}(\alpha_k) & \cdots & \text{imprv}_{\tau_0, \mathcal{S}}(\alpha_{|\text{arcs}(\mathcal{S})|}) \\ \vdots & & \vdots & & \vdots \end{bmatrix}$$

By definition, we have that for any $\alpha \in \text{arcs}(\mathcal{S})$:

$$\text{imprv}_{\tau_0, \mathcal{S}}(\alpha) = \tau_i(u) \cdot \text{imprv}_{\tau_0, \mathcal{S}}(i) - \tau_j(u) \cdot \text{imprv}_{\tau_0, \mathcal{S}}(j) \in \mathbb{Z}^{E(\mathcal{S})},$$

or equivalently $M_{\text{arcs}(\mathcal{S})}(\tau_0, \mathcal{S}) = \mathcal{M}_{1\text{-move}}(\tau_0, \mathcal{S}) \cdot \mathcal{T}(\tau_0, \mathcal{S})$, where $\mathcal{T}(\tau_0, \mathcal{S})$ is a sparse $(|\text{1-move}(\mathcal{S})| \times |\text{arcs}(\mathcal{S})|)$ rectangular matrix such that

$$\mathcal{T}(\tau_0, \mathcal{S})_{k\text{-th 1-move, } \alpha=(i,j)} = \begin{cases} (-1)^{\frac{(k-i)}{(j-i)}} \tau_k(u) & k \in \{i, j\} \\ 0 & \text{otherwise} \end{cases}, \text{ where } u \text{ is the corresponding node of arc } \alpha.$$

Schematically, we have that matrix $\mathcal{T}(\tau_0, \mathcal{S})$ includes the $(\tau_i(\text{node}(\alpha)), -\tau_j(\text{node}(\alpha)))_{\alpha=(i,j) \in \text{arcs}(\mathcal{S})}$ pairs expanded to $\{0, \pm 1\}^{|\text{1-move}(\mathcal{S})|}$.

$$\mathcal{T}(\tau_0, \mathcal{S}) = \begin{bmatrix} 0 & \tau_i(\text{node}(\alpha_1 = (i, j))) & 0 & \cdots & -\tau_j(\text{node}(\alpha_1 = (i, j))) & 0 & \cdots & 0 \\ \vdots & 0 & \vdots & \vdots & 0 & \vdots & \vdots & \vdots \\ \vdots & \vdots & 0 & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \tau_{i'}(\text{node}(\alpha_{\text{arcs}(\mathcal{S})} = (i', j'))) & 0 & \cdots & \tau_{i'}(\text{node}(\alpha_{\text{arcs}(\mathcal{S})} = (i', j'))) & 0 & 0 \end{bmatrix}^T$$

Again, notice that we can derive $\mathcal{T}(\tau'_0)$ by multiplying $\mathcal{T}(\tau_0)$ from right by the squared diagonal $|\text{arcs}(S)| \times |\text{arcs}(S)|$ matrix \mathcal{D}' such that

$$\mathcal{D}'[\tau_0, \tau'_0]_{(\alpha, \alpha) = ((i, j), (i, j))} = \tau_0(u)\tau'_0(u) \text{ where } u \text{ is the node of } \alpha \text{ and } \mathcal{D}'[\tau_0, \tau'_0]_{(\alpha, \alpha')} = 0 \text{ for } \alpha \neq \alpha'.$$

Indeed, we have that for an entry (k, α) representing an arc $\alpha = (i, j)$ and the k -th 1-move

$$(\mathcal{T}(\tau_0)\mathcal{D}'[\tau_0, \tau'_0])_{k, \alpha} = \begin{cases} \tau_0(u)\tau'_0(u)\tau_i(u) = \tau_i(u)\tau'_i(u)\tau_i(u) = \tau'_i(u) & \text{if } (\mathcal{T}(\tau_0))_{k, \alpha} = \tau_i(u) \\ -\tau_0(u)\tau'_0(u)\tau_j(u) = \tau_j(u)\tau'_j(u)\tau_j(u) = -\tau'_j(u) & \text{if } (\mathcal{T}(\tau_0))_{k, \alpha} = -\tau_j(u) \\ 0 & \text{if } (\mathcal{T}(\tau_0))_{k, \alpha} = 0 \end{cases} = (\mathcal{T}(\tau'_0))_{k, \alpha}$$

where first equality leverages Fact B.1 and the second one uses the fact $\tau_i(v)^2$ equals 1 for any $v \in V_n$ for every $i \in [\ell]$. To sum-up, it holds that

$$\begin{aligned} M_{\text{arcs}(S)}(\tau'_0, \mathcal{S}) &= \mathcal{M}_{1\text{-move}}(\tau'_0, \mathcal{S}) \cdot \mathcal{T}(\tau'_0, \mathcal{S}) = (\mathcal{D}[\tau_0, \tau'_0]\mathcal{M}_{1\text{-move}}(\tau_0, \mathcal{S})) \cdot (\mathcal{T}(\tau_0, \mathcal{S})\mathcal{D}'[\tau_0, \tau'_0]) \\ &= \mathcal{D}[\tau_0, \tau'_0](\mathcal{M}_{1\text{-move}}(\tau_0, \mathcal{S})) \cdot \mathcal{T}(\tau_0, \mathcal{S})\mathcal{D}'[\tau_0, \tau'_0] = \mathcal{D}[\tau_0, \tau'_0]M_{\text{arcs}(S)}(\tau'_0, \mathcal{S})\mathcal{D}'[\tau_0, \tau'_0] \end{aligned}$$

Given that $\mathcal{D}'[\tau_0, \tau'_0]$ and $\mathcal{D}[\tau_0, \tau'_0]$ are full-rank matrices, leveraging Fact B.2, we can prove that the above argument proves that $\text{rank}(\mathcal{M}_{\text{arcs}(S)}(\tau_0, \mathcal{S}))$ is independent of the initial configuration.

Now, in order to prove Lemma 2.8, we recall the definition of a dependent cycle C of size t , namely let $C = (c_1 = \{u_1, u_2\}, c_2 = \{u_2, u_3\}, \dots, c_t = \{u_t, u_1\})$ in a move sequence \mathcal{S} . For an initial configuration $\tau_0 \in \{\pm 1\}^{V(\mathcal{S})}$, we say that cycle is *dependent* with respect to τ_0 if there exists $b \in \{\pm 1\}^t$ such that

$$\begin{bmatrix} \tau_1(u_1) & 0 & \dots & \dots & 0 & \tau_t(u_1) \\ \tau_1(u_2) & \tau_2(u_2) & \dots & \dots & \dots & 0 \\ 0 & \tau_2(u_3) & \tau_3(u_3) & \dots & \dots & 0 \\ \vdots & \vdots & \ddots & \ddots & \ddots & 0 \\ \vdots & \vdots & \vdots & \ddots & \ddots & 0 \\ 0 & 0 & 0 & 0 & \tau_{t-1}(u_t) & \tau_t(u_t) \end{bmatrix} \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ \vdots \\ \vdots \\ b_t \end{bmatrix} = \begin{bmatrix} 0 \\ \vdots \\ \vdots \\ \vdots \\ \vdots \\ 0 \end{bmatrix} \equiv \Delta(\tau_0) \cdot b = 0_t$$

Again, notice that we can derive $\Delta(\tau'_0)$ by multiplying $\Delta(\tau_0)$ from left by the squared diagonal $|C| \times |C|$ matrix \mathcal{D}'' such that

$$\mathcal{D}''[\tau_0, \tau'_0]_{(u_k, u'_k)} = \tau_0(u)\tau'_0(u) \text{ for } k = k' \text{ and } \mathcal{D}''[\tau_0, \tau'_0]_{(u_k, u'_k)} = 0 \text{ for } k \neq k'.$$

Indeed, we have that for an entry (k, c_i) representing a 2-move $c_i = \{u_{i \bmod t}, u_{i+1 \bmod t}\}$ and the u_k -th node of the cycle C :

$$(\mathcal{D}''[\tau_0, \tau'_0]\Delta(\tau_0))_{k, c_i} = \begin{cases} \tau_0(u_k)\tau'_0(u_k)\tau_i(u_k) = \tau_i(u_k)\tau'_i(u_k)\tau_i(u_k) = \tau'_i(u_k) & \text{if } (\Delta(\tau_0))_{k, c_i} = \tau_i(u_k) \\ 0 & \text{if } (\Delta(\tau_0))_{k, c_i} = 0 \end{cases} = (\Delta(\tau'_0))_{k, c_i}$$

Since $\mathcal{D}''[\tau_0, \tau'_0]$ is full-rank matrix, we get that $(\Delta(\tau_0)) = (\Delta(\tau'_0))$. Therefore if there exists a non-zero vector b such that $\Delta(\tau_0) \cdot b = 0$ then $\Delta(\tau'_0) \cdot b = 0$ as well, completing the proof for Lemma 2.8.

For the last case of Lemma 2.10, we start by the following observations:

1. If $\Delta(\tau_0) \cdot b = 0$, then for any $b' = \lambda b$, it also holds that $\Delta(\tau_0) \cdot b' = 0$ for any non-zero λ constant.
2. If $\Delta(\tau_0) \cdot b = 0$ and $b_k = 0$ for some $k \in [t]$, then b is the zero vector, $b = 0$.
3. More precisely, the vector that belong to the (right) null space of $\Delta(\tau_0)$, i.e., all the vectors b such that $\Delta(\tau_0) \cdot b = 0$, are of the following form:

$$\Delta(\tau_0) \cdot b = 0 \Leftrightarrow b = b_1 \cdot \left(1, -\frac{\tau_{c_1}(u_2)}{\tau_{c_2}(u_2)}, \dots, (-1)^{k-1} \frac{\prod_{i \in [2:k]} \tau_{c_{i-1}}(u_i)}{\prod_{i \in [2:k]} \tau_{c_i}(u_i)}, \dots, (-1)^{t-1} \frac{\prod_{i \in [2:t]} \tau_{c_{i-1}}(u_i)}{\prod_{i \in [2:t]} \tau_{c_i}(u_i)} \right)^T$$

4. The term $\frac{\prod_{i \in [2:k]} \tau_{c_{i-1}}(u_i)}{\prod_{i \in [2:k]} \tau_{c_i}(u_i)}$ is independent of the initial configuration.

Items (1)-(3) are simple linear algebras derivation. Item (4) holds since

$$\frac{\prod_{i \in [2:k]} \tau_{c_{i-1}}(u_i)}{\prod_{i \in [2:k]} \tau_{c_i}(u_i)} = \frac{\prod_{i \in [2:k]} \tau_{c_{i-1}}(u_i)}{\prod_{i \in [2:k]} \tau_{c_i}(u_i)} \times \underbrace{\frac{\prod_{i \in [2:k]} \tau'_{c_{i-1}}(u_i) \tau_{c_{i-1}}(u_i)}{\prod_{i \in [2:k]} \tau'_{c_i}(u_i) \tau_{c_i}(u_i)}}_{=1 \text{ from Fact B.1}} = \frac{\prod_{i \in [2:k]} \tau'_{c_{i-1}}(u_i)}{\prod_{i \in [2:k]} \tau'_{c_i}(u_i)}$$

Hence, if b is a cancelling vector with $b_1 = 1$, then from items (1),(2),(3),(4) we have that b is unique for every cycle C and independent initial configuration τ_0 .

More interestingly, in order to prove Lemma 2.10, it suffices to prove that the rank of the column matrix with the improvement vectors of cycles is independent of the initial configuration. In fact, let

$$M_{\text{cycles}(\mathcal{S})}(\tau_0, \mathcal{S}) = \begin{bmatrix} \vdots & & \vdots & & \vdots \\ \text{imprv}_{\tau_0, \mathcal{S}}(C_1) & \cdots & \text{imprv}_{\tau_0, \mathcal{S}}(C_k) & \cdots & \text{imprv}_{\tau_0, \mathcal{S}}(C_{|\text{cycles}(\mathcal{S})|}) \\ \vdots & & \vdots & & \vdots \end{bmatrix}$$

By definition, we have that for any $C \in \text{cycles}(\mathcal{S})$:

$$\text{imprv}_{\tau_0, \mathcal{S}}(C = (c_1, \dots, c_t)) = \sum_{j \in [t]} b_j(C) \cdot \text{imprv}_{\tau_0, \mathcal{S}}(c_j).$$

or equivalently $M_{\text{cycles}(\mathcal{S})}(\tau_0, \mathcal{S}) = \mathcal{M}_{\text{2-move}}(\tau_0, \mathcal{S}) \cdot \mathcal{B}(\mathcal{S})$ where $\mathcal{T}(\tau_0, \mathcal{S})$ is a sparse $(|\text{2-move}(\mathcal{S})| \times |\text{cycles}(\mathcal{S})|)$ rectangular matrix such that

$$\mathcal{B}(\mathcal{S})_{\rho-\text{th 2-move}, C} = \begin{cases} b_\rho(C) & c_\rho \in C \\ 0 & \text{otherwise} \end{cases}, \text{ where } u \text{ is the corresponding node of arc } \alpha.$$

Schematically, we have that matrix $\mathcal{B}(\mathcal{S})$ includes the $(b(C))_{C \in \text{cycles}(\mathcal{S})}$ vectors expanded to $\{0, \pm 1\}^{|\text{2-move}(\mathcal{S})|}$

$$\mathcal{B}(\mathcal{S}) = \begin{bmatrix} 0 & b_1(C_1) & 0 & b_p(C_1) & 0 & b_{|C_1|}(C_1) & 0 & \cdots & 0 \\ b_1(C_2) & 0 & \cdots & b_{p'}(C_2) & \vdots & 0 & \cdots & b_{|C_2|}(C_2) & 0 \\ 0 & \cdots & \vdots & \cdots & \vdots & \cdots & \cdots & \cdots & 0 \\ b_1(C_{|\text{cycles}(\mathcal{S})|}) & b_{p''}(C_{|\text{cycles}(\mathcal{S})|}) & 0 & \cdots & \cdots & 0 & \cdots & b_{|C_{|\text{cycles}(\mathcal{S})|}}(C_{|\text{cycles}(\mathcal{S})|}) & 0 \end{bmatrix}^\top$$

Having noticed the above ones, it is easy to see that $\text{rank}(\mathcal{M}_{\text{arcs}(\mathcal{S})}(\tau_0, \mathcal{S}))$ is independent of the initial configuration, since

$$\begin{aligned} \mathcal{M}_{\text{arcs}(\mathcal{S})}(\tau'_0, \mathcal{S}) &= \mathcal{M}_{2\text{-move}}(\tau'_0, \mathcal{S}) \cdot \mathcal{B}(\mathcal{S}) = (\mathcal{D}[\tau_0, \tau'_0] \cdot \mathcal{M}_{2\text{-move}}(\tau_0, \mathcal{S})) \mathcal{B}(\mathcal{S}) \\ &= \mathcal{D}[\tau_0, \tau'_0] \cdot (\mathcal{M}_{2\text{-move}}(\tau_0, \mathcal{S}) \mathcal{B}(\mathcal{S})) = \mathcal{D}[\tau_0, \tau'_0] \cdot \mathcal{M}_{\text{arcs}(\mathcal{S})}(\tau_0, \mathcal{S}). \end{aligned}$$

Thus, by Fact B.2 we get that $\text{rank}(\mathcal{M}_{\text{arcs}(\mathcal{S})}(\tau_0, \mathcal{S})) = \text{rank}(\mathcal{M}_{\text{arcs}(\mathcal{S})}(\tau'_0, \mathcal{S}))$, concluding also the proof of Lemma 2.10. \square

References

- [1] Daniel A Spielman and Shang-Hua Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *Journal of the ACM (JACM)*, 51(3):385–463, 2004.
- [2] René Beier and Berthold Vöcking. Random knapsack in expected polynomial time. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 232–241, 2003.
- [3] Matthias Englert, Heiko Röglin, and Berthold Vöcking. Smoothed analysis of the 2-opt algorithm for the general tsp. *ACM Transactions on Algorithms (TALG)*, 13(1):1–15, 2016.
- [4] Heiko Röglin and Berthold Vöcking. Smoothed analysis of integer programming. *Mathematical programming*, 110(1):21–56, 2007.
- [5] Daniel Dadush and Sophie Huisbergs. A friendly smoothed analysis of the simplex method. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 390–403, 2018.
- [6] Aditya Bhaskara, Moses Charikar, Ankur Moitra, and Aravindan Vijayaraghavan. Smoothed analysis of tensor decompositions. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 594–603, 2014.
- [7] Brendan Farrell and Roman Vershynin. Smoothed analysis of symmetric random matrices with continuous distributions. *Proceedings of the American Mathematical Society*, 144(5):2257–2261, 2016.
- [8] Avrim Blum and John Dunagan. Smoothed analysis of the perceptron algorithm for linear programming. 2002.
- [9] David Arthur, Bodo Manthey, and Heiko Röglin. Smoothed analysis of the k-means method. *Journal of the ACM (JACM)*, 58(5):1–31, 2011.
- [10] David Arthur and Sergei Vassilvitskii. Worst-case and smoothed analysis of the icp algorithm, with an application to the k-means method. In *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS’06)*, pages 153–164. IEEE, 2006.
- [11] Vidyashankar Sivakumar, Steven Wu, and Arindam Banerjee. Structured linear contextual bandits: A sharp and geometric smoothed analysis. In *International Conference on Machine Learning*, pages 9026–9035. PMLR, 2020.
- [12] Shant Boodaghians, Joshua Brakensiek, Samuel B Hopkins, and Aviad Rubinstein. Smoothed complexity of 2-player nash equilibria. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 271–282. IEEE, 2020.
- [13] Shant Boodaghians, Rucha Kulkarni, and Ruta Mehta. Smoothed efficient algorithms and reductions for network coordination games. In Thomas Vidick, editor, *11th Innovations in Theoretical*

Computer Science Conference, ITCS 2020, January 12-14, 2020, Seattle, Washington, USA, volume 151 of LIPIcs, pages 73:1–73:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.

- [14] René Beier, Heiko Röglin, Clemens Rösner, and Berthold Vöcking. The smoothed number of pareto-optimal solutions in bicriteria integer optimization. *Mathematical Programming*, pages 1–37, 2022.
- [15] Lirong Xia. The smoothed possibility of social choice. *Advances in Neural Information Processing Systems*, 33:11044–11055, 2020.
- [16] David S. Johnson, Cecilia R. Aragon, Lyle A. McGeoch, and Catherine Schevon. Optimization by simulated annealing: An experimental evaluation; part i, graph partitioning. *Oper. Res.*, 37(6):865–892, 1989.
- [17] M.R. Garey, D.S. Johnson, and L. Stockmeyer. Some simplified np-complete graph problems. *Theor. Comput. Sci.*, pages 237–267, 1976.
- [18] Brian W. Kernighan and Shen Lin. An efficient heuristic procedure for partitioning graphs. *Bell Syst. Tech. J.*, 49(2):291–307, 1970.
- [19] David S Johnson, Christos H Papadimitriou, and Mihalis Yannakakis. How easy is local search? *Journal of Computer and System Sciences*, 37(1):79–100, 1988.
- [20] Alejandro A Schäffer and Mihalis Yannakakis. Simple local search problems that are hard to solve. *SIAM journal on Computing*, 20(1):56–87, 1991.
- [21] Robert Elsässer and Tobias Tscheuschner. Settling the complexity of local max-cut (almost) completely. In *International Colloquium on Automata, Languages, and Programming*, pages 171–182. Springer, 2011.
- [22] Michael Etscheid and Heiko Röglin. Smoothed analysis of local search for the maximum-cut problem. *ACM Trans. Algorithms*, 13(2):25:1–25:12, March 2017.
- [23] Omer Angel, Sébastien Bubeck, Yuval Peres, and Fan Wei. Local max-cut in smoothed polynomial time. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 429–437. ACM, 2017.
- [24] Ali Bibak, Charles Carlson, and Karthekeyan Chandrasekaran. Improving the smoothed complexity of flip for max cut problems. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 897–916. SIAM, 2019.
- [25] Xi Chen, Chenghao Guo, Emmanouil V. Vlatakis-Gkaragkounis, Mihalis Yannakakis, and Xinzhi Zhang. Smoothed complexity of local max-cut and binary max-csp. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2020, page 1052–1065, New York, NY, USA, 2020. Association for Computing Machinery.

- [26] Charles M. Fiduccia and Robert M. Mattheyses. A linear-time heuristic for improving network partitions. In *Proceedings of the 19th Design Automation Conference*, pages 175–181. ACM/IEEE, 1982.
- [27] Heiko Röglin. *The complexity of Nash equilibria, local optima, and Pareto optimal solutions*. PhD thesis, Aachen, Techn. Hochsch., Diss., 2008, 2008.