

Lower bounds for the parameterized complexity of Minimum Fill-In and other completion problems

Ivan Bliznets* Marek Cygan† Paweł Komosa‡ Lukáš Mach§ Michał Pilipczuk¶

Abstract

In this work, we focus on several completion problems for subclasses of chordal graphs: MINIMUM FILL-IN, INTERVAL COMPLETION, PROPER INTERVAL COMPLETION, THRESHOLD COMPLETION, and TRIVIAALLY PERFECT COMPLETION. In these problems, the task is to add at most k edges to a given graph in order to obtain a chordal, interval, proper interval, threshold, or trivially perfect graph, respectively. We prove the following lower bounds for all these problems, as well as for the related CHAIN COMPLETION problem:

- Assuming the Exponential Time Hypothesis, none of these problems can be solved in time $2^{\mathcal{O}(n^{1/2}/\log^c n)}$ or $2^{\mathcal{O}(k^{1/4}/\log^c k)} \cdot n^{\mathcal{O}(1)}$, for some integer c .
- Assuming the non-existence of a subexponential-time approximation scheme for MIN BISECTION on d -regular graphs, for some constant d , none of these problems can be solved in time $2^{o(n)}$ or $2^{o(\sqrt{k})} \cdot n^{\mathcal{O}(1)}$.

For all the aforementioned completion problems, apart from PROPER INTERVAL COMPLETION, FPT algorithms with running time of the form $2^{\mathcal{O}(\sqrt{k} \log k)} \cdot n^{\mathcal{O}(1)}$ are known. Thus, the second result proves that a significant improvement of any of these algorithms would lead to a surprising breakthrough in the design of approximation algorithms for MIN BISECTION.

To prove our results, we use a reduction methodology based on combining the classic approach of starting with a sparse instance of 3-SAT, prepared using the Sparsification Lemma, with the existence of almost linear-size Probabilistically Checkable Proofs (PCPs). Apart from our main results, we also obtain lower bounds excluding the existence of subexponential algorithms for the OPTIMUM LINEAR ARRANGEMENT problem, as well as improved, yet still not tight, lower bounds for FEEDBACK ARC SET IN TOURNAMENTS.

*St. Petersburg Department of Steklov Institute of Mathematics. E-mail: iabliznets@gmail.com. Partially supported by the Government of the Russian Federation (grant 14.Z50.31.0030), by the Grant of the President of the Russian Federation (MK-6550.2015.1) as well as by Warsaw Center of Mathematics and Computer Science.

†Institute of Informatics, University of Warsaw, Poland. E-mail: cygan@mimuw.edu.pl. Partially supported by the Polish National Science Centre grant DEC-2012/05/D/ST6/03214.

‡Institute of Informatics, University of Warsaw, Poland. E-mail: kompaw01@gmail.com.

§DIMAP and Department of Computer Science, University of Warwick, United Kingdom. E-mail: lukas.mach@gmail.com. Received funding by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013)/ERC grant agreement no. 259385.

¶Institute of Informatics, University of Warsaw, Poland. E-mail: mp248287@mimuw.edu.pl. Supported by the Polish National Science Centre grant DEC-2013/11/D/ST6/03073 and Foundation for Polish Science via the START stipend programme. During the work on these results, Michał Pilipczuk has been holding a post-doc position of Warsaw Center of Mathematics and Computer Science.

1 Introduction

In the MINIMUM FILL-IN problem, also known as CHORDAL COMPLETION, the input is an undirected graph G and an integer k , and the question is whether at most k edges can be added to G in order to turn it into a chordal graph, i.e., a graph without induced cycles of length at least 4 (also known as *holes*). The interest in this problem originates in the study of strategies for Gaussian elimination on sparse matrices, because the optimum number of additional entries of a matrix A that become non-zero during the elimination is tightly connected to the minimum fill-in of the graph G_A obtained by taking A to be its adjacency matrix. See [15, 41] for more information on applications of MINIMUM FILL-IN in the theory of sparse matrices. However, the problems of adding as few edges as possible to obtain a chordal graph, or a graph belonging to some natural subclass of chordal graphs, like interval, proper interval, trivially perfect, or threshold graphs, have numerous other applications ranging from database management, bioinformatics, artificial intelligence, to social networks. We refer to the introductory sections of [6, 7, 18, 19, 21, 25, 32, 43] for a broader discussion and pointers to relevant literature.

MINIMUM FILL-IN is NP-hard, as shown by Yannakakis [44], however the reduction showing this is quite non-obvious; the complexity status of the problem was among the 12 open problems at the end of the first edition of the Garey and Johnson's book [26]. The study of MINIMUM FILL-IN from the point of view of the parameterized complexity, with k being the obvious parameter of interest, started with the pioneering work of Kaplan et al. [32]. They proposed a fixed-parameter (FPT) algorithm with running time $\mathcal{O}(16^k k^6 + k^2 mn)$ that is based on locating holes in the graph and branching on possible ways of adding edges to get rid of them. A similar strategy worked also for PROPER INTERVAL COMPLETION [32], the problem of adding as few edges as possible to obtain a proper interval graph, but not for INTERVAL COMPLETION. The fixed-parameter tractability of the latter has been resolved by Villanger et al. [43] only several years thereafter.

A complete turning point came four years ago, when Fomin and Villanger [25] presented an algorithm for MINIMUM FILL-IN with *subexponential* parameterized complexity, more precisely with running time $\mathcal{O}(2^{\mathcal{O}(\sqrt{k} \log k)} + k^2 nm)$. This was an immense surprise to the parameterized complexity community, since subexponential parameterized algorithms, i.e., with running time $2^{o(k)} \cdot n^{\mathcal{O}(1)}$, were known essentially only in two restricted settings: in topologically constrained graphs via the technique of bidimensionality (see e.g. [16]), and in tournaments, with an important example of FEEDBACK ARC SET IN TOURNAMENTS [3, 23, 33]. For most natural parameterized problems, the existence of such algorithms can be excluded under the Exponential Time Hypothesis, which (when combined with the *Sparsification Lemma* [31]) essentially states that there is no algorithm for 3-SAT with running time $2^{o(n+m)}$; cf. [14, 24].

The work of Fomin and Villanger [25] presented a conceptual breakthrough in the approach to completion problems for subclasses of chordal graphs. The main idea is not to focus on breaking individual obstacles such as holes by single edge additions, as was proposed in the previous works, but to build a structural decomposition of the completed graph by means of a dynamic programming algorithm that minimizes the number of edges missing in the decomposition. The crux is to show that this dynamic programming can be restricted to a space of states that has size $2^{\mathcal{O}(\sqrt{k} \log k)} \cdot n^{\mathcal{O}(1)}$, and can be enumerated efficiently. In case of MINIMUM FILL-IN, the considered decomposition is the clique tree, but this generic approach can be in principle applied to every subclass of chordal graphs whose graphs have a well-defined global structure. And so, following [25], subexponential parameterized algorithms have been designed for THRESHOLD COMPLETION [20], TRIVIAALLY PERFECT COMPLETION [20], PROPER INTERVAL COMPLETION [7],

and even INTERVAL COMPLETION [6]. Apart from PROPER INTERVAL COMPLETION, for all these problems the algorithms have running time $2^{\mathcal{O}(\sqrt{k} \log k)} \cdot n^{\mathcal{O}(1)}$; for the former we are currently stuck at exponent $\mathcal{O}(k^{2/3} \log k)$, but this is conjectured to be an artifact of the technique [7]. Let us remark that in each of these cases the idea of Fomin and Villanger only provides the basic outline of the strategy, while the actual implementation is always class-specific and requires involved technical insight.

Drange et al. [19,20] complemented their results with a number of lower bounds suggesting that even a slight deviation from the setting of adding edges to a subclass of chordal graphs leads to the non-existence of a subexponential parameterized algorithm under the assumption of ETH. Thus, the aforementioned problems for which such algorithms exist are in fact intriguing “singularity points” on the complexity landscape. It is interesting that these singularity points actually correspond to problems that have important practical applications. Hence, the study of this phenomenon is an important direction that naturally belongs to the *optimality programme* (cf. [35,38]): a trend in parameterized complexity that focuses on systematic study of parameterized problems by providing possibly tight upper and lower bounds on their complexity.

From this point of view, the first natural question is exactly how “deep” are the aforementioned singularity points. More precisely, is the running time of the form $2^{\mathcal{O}(\sqrt{k} \log k)} \cdot n^{\mathcal{O}(1)}$ optimal, say under ETH? For subexponential algorithms derived using bidimensionality, the non-existence of $2^{o(\sqrt{k})} \cdot n^{\mathcal{O}(1)}$ -time algorithms under ETH usually follows already from the known NP-hardness reductions [14,24]. However, for the considered completion problems this is not the case: as Fomin and Villanger already observed in [25], the classic NP-hardness reduction for MINIMUM FILL-IN gives only a $2^{o(k^{1/6})} \cdot n^{\mathcal{O}(1)}$ lower bound. A similar situation holds for FEEDBACK ARC SET IN TOURNAMENTS, for which the fastest known algorithms work in time $2^{\mathcal{O}(\sqrt{k})} \cdot n^{\mathcal{O}(1)}$ [23,33].

For this reason, the question of establishing tight upper and lower bounds for MINIMUM FILL-IN was already asked explicitly by Fomin and Villanger [25], repeated by Marx in his survey on the optimality programme [38], and then reiterated for respective subclasses in all the works [6,7,19]. The goal of this paper is to remedy this situation by providing complexity foundations for proving that the square root in the exponent of the running time is hard to improve.

Our results. First, we investigate how strong lower bounds for completion problems can be obtained when we assume only ETH.

Theorem 1.1. *Unless ETH fails, there is an integer $c \geq 1$ such that there are no $2^{\mathcal{O}(\sqrt{n}/\log^c n)}$, and consequently no $2^{\mathcal{O}(k^{1/4}/\log^c k)} \cdot n^{\mathcal{O}(1)}$ algorithms for the following problems: MINIMUM FILL-IN, INTERVAL COMPLETION, PROPER INTERVAL COMPLETION, TRIVIALY PERFECT COMPLETION, THRESHOLD COMPLETION, CHAIN COMPLETION.*

Here, a graph G is a *chain graph* if it is a bipartite graph with a fixed bipartition $A \uplus B$, where the vertices of A can be ordered by a linear ordering \preceq such that $N(u) \subseteq N(v)$ whenever $u \preceq v$, for all $u, v \in A$. In the CHAIN COMPLETION problem we are given a bipartite graph with a fixed bipartition $A \uplus B$, and we want to add at most k edges between A and B to obtain a chain graph. CHAIN COMPLETION also admits an algorithm with running time $\mathcal{O}(2^{\mathcal{O}(\sqrt{k} \log k)} + k^2 nm)$ [25].

Unfortunately, Theorem 1.1 does not provide a tight result for completion problems we are interested in: we still have a gap between $k^{1/4}$ and $k^{1/2}$. We in fact prove a higher lower bound, but we need to ground it on a stronger complexity assumption. More precisely, we consider the approximability of the MIN BISECTION problem on d -regular graphs: given a d -regular graph G

with an even number of vertices, find a partition of $V(G)$ into *equal-sized* parts that minimizes the number of edges between the parts. For $0 \leq \alpha < \beta \leq 1$, problem $\text{GAP MIN BISECTION}(d)_{[\alpha, \beta]}$ is defined as the gap problem of deciding whether the optimum solution is at most αm , or at least βm .

Hypothesis 1.2. *There exist $0 \leq \alpha < \beta \leq 1$, and an integer $d > \frac{4}{\beta - \alpha}$, such that there is no $2^{o(n)}$ -time algorithm for $\text{GAP MIN BISECTION}(d)_{[\alpha, \beta]}$.*

Theorem 1.3. *Unless Hypothesis 1.2 fails, there is no $2^{o(n+m)}$ -time algorithm for CHAIN COMPLETION, and no $2^{o(n)}$ -time algorithms for MINIMUM FILL-IN, INTERVAL COMPLETION, PROPER INTERVAL COMPLETION, TRIVIALY PERFECT COMPLETION, and THRESHOLD COMPLETION. Consequently, none of these problems can be solved in time $2^{o(\sqrt{k})} \cdot n^{\mathcal{O}(1)}$.*

Thus, Theorem 1.3 asserts that the existence of a substantially faster algorithm for any of the considered completion problems would lead to a breakthrough in the design of approximation algorithms for MIN BISECTION. We do not intend to take a stance on whether Hypothesis 1.2 is true or false. However, the existence of an algorithm refuted by Theorem 1.2 is far from the current knowledge: the best polynomial-time approximation algorithms for MIN BISECTION achieve approximation factor $\mathcal{O}(\log OPT)$ [40], whereas to refute Hypothesis 1.2 one would need to obtain approximation factor arbitrarily close to one. This could be, however, possible due to the assumption of d -regularity, for a constant d , as well as the access to subexponential-time computations. We are not aware of any approximation algorithms for MIN BISECTION that would significantly use any of these assumptions. In fact, Berman and Karpinski [5] have shown that approximating MIN BISECTION in graphs of maximum degree 3 is as hard as in general graphs from the point of view of polynomial time approximation.

Finally, using the methodology of Theorem 1.1 we can give an improved lower bound for FEEDBACK ARC SET IN TOURNAMENTS. For this problem, the first subexponential FPT algorithm with running time $2^{\mathcal{O}(\sqrt{k} \log k)} \cdot n^{\mathcal{O}(1)}$ was given by Alon et al. [3], which was later improved to $2^{\mathcal{O}(\sqrt{k})} \cdot n^{\mathcal{O}(1)}$ independently by Feige [23] and by Karpinski and Schudy [33].

Theorem 1.4. ¹ *Unless ETH fails, there is an integer $c \geq 1$ such that there is no $2^{\mathcal{O}(\sqrt{n}/\log^c n)}$, and consequently no $2^{\mathcal{O}(k^{1/4}/\log^c k)} \cdot n^{\mathcal{O}(1)}$ -time algorithm for FEEDBACK ARC SET IN TOURNAMENTS.*

A direct inspection of the classic NP-hardness proofs [1,2,13] gives only a $2^{\mathcal{O}(n^{1/4})}$, and consequently also a $2^{\mathcal{O}(k^{1/8})} \cdot n^{\mathcal{O}(1)}$ lower bound, but Theorem 1.4 is still not tight: we have a gap between $k^{1/2}$ and $k^{1/4}$, similarly as in Theorem 1.1. Unfortunately, despite efforts we are unable to close this gap using Hypothesis 1.2, or a similar assumption.

Our techniques. The main idea of this work is to carefully combine the standard approach of starting with a sparse instance of 3-SAT, prepared using the Sparsification Lemma [31], with the existence of almost linear PCPs. More precisely, an application of the Sparsification Lemma reduces solving an instance φ of 3-SAT to solving a number of instances $\phi_1, \phi_2, \dots, \phi_\ell$ of 3-SAT on the same variable sets, but in each ϕ_i we have that the number of clauses m is linear in the number of variables n . Then, to each ϕ_i we apply the result of Dinur [17] on the existence of almost linear-size PCPs. This transforms each ϕ_i into an instance ϕ'_i of 3-SAT where the number

¹As the reduction in the proof of Theorem 1.4 is randomized, here we assume the randomized version of the ETH.

of variables and clauses is $\mathcal{O}(n \log^c n)$, but the instance is either satisfiable, or at most a ρ -fraction of clauses can be satisfied, for some $\rho < 1$.

Having such an instance at hand, we can proceed with the chain of reductions proposed by Garey et al. [27] and Yannakakis [44]. Namely, we first use simple manipulations to obtain an equivalent instance of MAX CUT with a gap, and then reduce it to the OPTIMUM LINEAR ARRANGEMENT (OLA) problem: Given a graph G on n vertices, find a linear ordering $\pi: V(G) \rightarrow \{1, 2, \dots, n\}$ that minimizes the cost defined as $\sum_{uv \in E(G)} |\pi(u) - \pi(v)|$; value $|\pi(u) - \pi(v)|$ is also called the *cost* of uv . This is precisely the moment where we exploit that we are working with an instance of MAX CUT with a gap. Namely, the construction introduces a huge clique to the complement of the MAX CUT instance; this clique is supposed to separate in the ordering the sides of an optimum max-cut solution. In Garey et al. [27], its size must be large enough so that the cost of any edge jumping over the clique dwarfs the “noise” contribution that is given by internal ordering of parts on the left and on the right. Hence, the clique is chosen to be of size $\Theta(n^4)$, which explodes the size of the instance. However, we observe that by starting with an instance with a gap, we can accommodate the noise in the gap, and therefore we only need the clique to be of linear size. Thus, we obtain the following theorem that can be of independent interest.

Theorem 1.5. *Unless ETH fails, for some $c > 1$ there is no algorithm solving OPTIMUM LINEAR ARRANGEMENT in time $2^{\mathcal{O}(n/\log^c n)}$.*

Note that OLA can be solved by a simple Held-Karp dynamic programming on subsets in time $2^n \cdot n^{\mathcal{O}(1)}$, so Theorem 1.5 gives an almost tight result.

To obtain Theorem 1.1 one just applies slightly modified elegant reduction from (the decision variant of) OLA to MINIMUM FILL-IN, proposed by Yannakakis [44]. The reduction firstly reduce OLA to CHAIN COMPLETION. Because of this fact later we get a very specific MINIMUM FILL-IN output instance. And this immediately gives us a reduction to PROPER INTERVAL COMPLETION and INTERVAL COMPLETION problems. With slight changes similar result is obtained for THRESHOLD COMPLETION and TRIVIALY PERFECT COMPLETION. Theorem 1.4 is derived by applying the same methodology to the randomized hardness reduction for FEEDBACK ARC SET IN TOURNAMENTS of Ailon et al. [1].

The moment when we lose tightness in our chain of reductions is precisely the last reduction from OLA to MINIMUM FILL-IN. This is because in the construction we introduce a new vertex for every *edge* of the original graph, which blows up the vertex set quadratically. Therefore, the reduction would give a tight lower bound if we started with a sparse instance of OPTIMUM LINEAR ARRANGEMENT. This is precisely our idea behind the proof of Theorem 1.3.

Namely, we observe that by starting from the hardness of GAP MIN BISECTION(d) $_{[\alpha, \beta]}$, instead of MAX CUT, we can give an alternative hardness reduction to OPTIMUM LINEAR ARRANGEMENT in graphs of bounded degree by a constant d (OLA $_{\leq}(d)$ for short). This reduction is much more intricate. Essentially, we replace the usage of a huge clique (which would blow up the degree) by a careful construction using several layers of expander graphs that mimics the same behaviour. This proves the following result.

Theorem 1.6. *Unless Hypothesis 1.2 fails, there exists an integer $d \in \mathbb{N}$ such that there is no $2^{\mathcal{O}(n)}$ -time algorithm solving OLA $_{\leq}(d)$ in multigraphs.*

From this, Theorem 1.3 follows similarly as Theorem 1.1 followed from Theorem 1.5.

Related work. The combination of ETH, Sparsification Lemma, and PCP tools was already used in the recent line of work on developing the theory of hardness of FPT approximation; see e.g. [8–11, 28, 34]. There, the main goal is to provide lower bounds on how the access to FPT-time computations can help in the design of approximation algorithms for fundamental inapproximable problems, like CLIQUE or SET COVER.

The existence of almost linear PCPs combined with ETH was also used by Marx [37] to give almost tight lower bounds for the running times of polynomial time approximation schemes (PTASes) for several fundamental problems on planar graphs and in the Euclidean plane. The basic principle behind the approach of Marx [37] is very similar to ours. Namely, it is observed that using ETH and PCP tools one can prepare a sparse instance of the maximization variant of 3-SAT where it is hard to distinguish between a fully satisfiable instance and an instance where only a $(1 - \epsilon)$ -fraction of the clauses can be satisfied in time $2^{\mathcal{O}(n/\log^c n)}$, for some constant c (see also Theorem 3.1 in this work). This observation is used as a base for further reductions refuting the existence of certain approximation algorithms for geometric and planar problems.

Thus, this work provides another example where the said combination appears to be useful. This time we use it to prove improved lower bounds on the complexity of FPT algorithms solving certain graph modification problems exactly. The new idea in this work is that the gap property of the considered instances can be used not only to exclude the existence of approximation algorithms, but also to limit the instance size explosion in a chain of NP-hardness reductions by using more thrifty constructions.

Hypothesis 1.2 can possibly have links with the hypothesis put forward by Feige [22]. Essentially, Feige conjectures the hardness of distinguishing a “typical” 3-SAT instance from a satisfiable one in polynomial time. From this, he derives as a corollary a variant of Hypothesis 1.2, but without the assumption of d -regularity, for a constant d , and with subexponential time replaced by polynomial. It is conceivable that Hypothesis 1.2 can be also implied by some stronger variant of Feige’s conjecture, but we refrain from formalizing this link due to many technical problems that arise when attempting to do this.

Outline of the paper. In Section 2, we give all the needed definitions and list the tools used in our reductions, such as the Exponential Time Hypothesis and the PCP theorem. Next, in Section 3, we prove Theorem 1.5. The main technical contribution of the paper, that is the proof of Theorem 1.6, is contained in Section 4. Theorem 1.6 is the main ingredient in the proof of Theorem 1.3, that is presented in Section 5. Section 6 is devoted to the proof of Theorem 1.4.

2 Preliminaries

2.1 Parameterized complexity

A *parameterized problem* Q is a subset of $\Sigma^* \times \mathbb{N}$, for a fixed finite alphabet Σ . An instance of the problem Q is an element $(x, k) \in \Sigma^* \times \mathbb{N}$ with the integer k called the parameter. A parameterized problem is *fixed parameter tractable* if there exists an algorithm deciding whether $(x, k) \in Q$, and working in time $f(k) \cdot \text{poly}(|x|)$ for every instance (x, k) , where f is a computable function. For denoting the time of an algorithm we will use the \mathcal{O}^* notation that suppresses factors polynomial in the size of the input, e.g. $\mathcal{O}^*(f(k))$.

We denote a polynomial deterministic linear reduction from a problem X to Y by $X \leq_{\text{P}}^{\text{lin}} Y$, i.e., $X \leq_{\text{P}}^{\text{lin}} Y$ means that there is a polynomial time deterministic algorithm, which given an instance

I of the problem X produces an instance I' of the problem Y of size $\mathcal{O}(|I|)$, such that I is a yes-instance if and only if I' is a yes-instance.

2.2 Graph notations

A standard graph theoretical notation is used throughout the paper. A *graph* G is a tuple $(V(G), E(G))$, where $V(G)$ is the set of *vertices* and $E(G) \subseteq \binom{V(G)}{2}$ is the set of *edges*. When it is clear from the context what graph we refer to, we will denote by V and E the set of vertices and edges of G . A graph H is a *subgraph* of G if $V(H) \subseteq V(G)$ and $E(H) \subseteq E(G)$. A graph H is an *induced subgraph* of G if $V(H) \subseteq V(G)$ and $E(H) = E(G) \cap \binom{V(H)}{2}$. An induced subgraph of G with vertex set X is denoted by $G[X]$. A *complement* of G is the graph with the vertices V and the edge $\binom{V}{2} - E$ and we denote it by \overline{G} . For $X \subseteq V$, we denote by $\delta_G(X)$ the set of edges with exactly one endpoint in X . We define a *cut*, as the set of edges $E_G(A, B)$, for a partition (A, B) of V . We denote the size of the cut by $|E_G(A, B)|$.

We often use n , and m to denote the size of V and E , respectively. For a vertex v , $\deg_G(v)$ denotes the degree (the number of incident edges) of the vertex v . We say that a graph is *d-regular* if the degree of each vertex is equal to d . We denote by Δ_G the maximum degree of G . The set $N_G(v) = \{w : (v, w) \in E(G)\}$ is the neighbourhood of v . We extend this notation to subsets of vertices X , i.e. $N_G(X) = \bigcup_{v \in X} N_G(v) \setminus X$. We will also omit subscripts, i.e. $\deg(v), \Delta, N(X), \delta(X), E(U, V)$, when it is clear from the context, which graph we refer to.

If $X, Y \subseteq V$ are disjoint, then $E_G(X, Y)$ is the set of edges between X and Y . We use $G[X, Y]$ to denote the *induced bipartite subgraph of G with parts X and Y* . That is, $G[X, Y]$ is the graph with vertex set $X \cup Y$ that contains precisely the edges $E_G(X, Y)$.

For $X \subseteq V$, we denote by $\delta_G(X)$ the set of edges with exactly one endpoint in X . We define a *cut*, as the set of edges $E_G(A, B)$, for a partition (A, B) of V . We denote the size of the cut by $|E_G(A, B)|$.

2.3 Expanders

The *Cheeger number* $h(G)$ of a graph G is defined as

$$h(G) := \min \left\{ \frac{|\delta(X)|}{|X|} : X \subseteq V(G), |X| \leq \frac{|V(G)|}{2} \right\},$$

We say that a graph G is a *(d, e)-expander* if it is d -regular and has Cheeger number of at least e . When discussing expanders, it is convenient to allow parallel edges and self-loops, which naturally appear in most of the expanders constructions. It is important to note that a self-loop contributes 1 to the degree of a vertex.

There are many efficient constructions of expanders available in the literature. The following theorem produces graphs with provably optimal Cheeger numbers:

Theorem 2.1. [36, 39] *Let $d = p^k + 1$, where p is a prime and $k \in \mathbb{N}$, and let q be a prime congruent to 1 mod 4. Assume $p \neq q$. There exists a $(d, \frac{d}{2} - \sqrt{d-1})$ -expander on $q+1$ vertices. Furthermore, such an expander can be constructed in polynomial time.*

Throughout this paper, we use expanders of various sizes, hence the following theorem appears to be useful in our setting, despite providing weaker bound on the Cheeger number.

Theorem 2.2. [Theorem 21.19 of [4]] *Let $p > 0$ be a real. Then there exists a positive integer d , such that for every positive integer n there exists a d -regular multigraph $G_{n,d}$ on n vertices with $h(G_{n,d}) \geq p$. Moreover, graph $G_{n,d}$ can be constructed in time polynomial in n .*

2.4 Linear arrangements

A *linear arrangement* of a graph $G = (V, E)$ is a function (a vertex ordering) $\pi : V \rightarrow \{1, \dots, n\}$. The *cost* of a linear arrangement π is defined by $\sum_{uv \in E} |\pi(u) - \pi(v)|$. We call π the optimum linear arrangement if its cost is minimized over all vertex orderings of G and we denote this cost by $\text{OLA}(G)$.

2.5 Satisfiability

We employ a standard notation related to SAT problems. We use symbols x_1, \dots, x_n for the variables of an instance, and C_1, \dots, C_m for the clauses. An l -CNF formula is *El-CNF*, if it has exactly l literals. We say that an assignment of x_1, \dots, x_n *NAE-satisfies* an l -CNF formula if every clause contains a satisfied and an unsatisfied literal. An l -AND is a conjunction of clauses, where each clause is a conjunction of at most l literals.

2.6 Exponential Time Hypothesis

The Exponential Time Hypothesis (ETH), introduced by Impagliazzo, Paturi and Zane [30, 31] is now an established tool used for proving conditional lower bounds in the parameterized complexity area (see [35] for a survey on ETH-based lower bounds). Intuitively, ETH states that 3-SAT cannot be solved in time subexponential in the number of variables.

Hypothesis 2.3 (Exponential Time Hypothesis (ETH) [30, 31]). *There is no $2^{o(n)}$ time randomized algorithm for 3-SAT.*

Lemma 2.4. (Sparsification Lemma, [31]) *For every $\epsilon > 0$, there is an algorithm that takes a 3-CNF formula ϕ and returns l 3-CNF formulas ϕ_1, \dots, ϕ_l , such that: i) $l = \mathcal{O}(2^{\epsilon n})$, ii) for every i , ϕ_i has n variables, and every such variable appears in at most c_ϵ clauses of ϕ_i , for some constant $c_\epsilon \geq 0$, iii) ϕ is satisfiable if and only if at least one of ϕ_i is satisfiable. The running time of the algorithm is $\mathcal{O}^*(2^{\epsilon n})$.*

Consequently, based on ETH we have subexponential hardness of 3-SAT in terms of both the number of variables and clauses.

Theorem 2.5 ([31]). *Unless ETH fails, there is no $2^{o(m+n)}$ time algorithm for 3-SAT.*

2.7 Gap problems and PCPs

In a gap version of a problem, the input instance is promised to belong to one of two languages specifying the allowed input, and the goal is to decide which case (language) a given instance belongs to. Gap problems are associated with two parameters α, β . For example, in $\text{GAP MAX CUT}_{[\alpha, \beta]}$, we are to distinguish between the case when a given graph admits a cut of size at least βm and the case when a given graph does not admit a cut of size larger than αm . Similarly for satisfiability problems with gap $[\alpha, \beta]$ we are to distinguish between a formula, for which at least βm clauses can be satisfied, and the case where it is impossible to satisfy more than αm clauses.

To introduce gaps in our reductions, we will use the following fundamental result of Dinur.

Theorem 2.6. (Almost Linear Size PCP, [17]) $3\text{-SAT} \in \mathbf{PCP}_{1, \frac{1}{2}}(\log(n) + \mathcal{O}(\log \log(n)), \mathcal{O}(1))$

3 Combining known reductions for OLA with PCPs and ETH

In this section we discuss the proof of Theorem 1.5, i.e., explain how to show that unless ETH fails, there is no $2^{\mathcal{O}(\frac{n}{\log^c(n)})}$ -time algorithm for OPTIMUM LINEAR ARRANGEMENT, for some constant $c \in \mathbb{N}$. We combine several well known concepts from the complexity theory: the ETH hypothesis, the PCP theorem, the Sparsification Lemma, gap reductions, and gap amplification.

First, we show that the combination of the ETH hypothesis, together with the almost linear size PCP theorem, and the Sparsification Lemma, gives subexponential hardness of GAP E3-SAT. This fact was already observed and used by Marx [37].

Theorem 3.1 (see also Lemma 2.2 of [37]). *Unless ETH fails, there exist $c \in \mathbb{N}$, and $r \in (0, 1)$ such that there is no $2^{\mathcal{O}(\frac{m}{\log^c(m)})}$ -time algorithm for GAP E3-SAT $_{[r, 1]}$.*

Proof. From Theorem 2.6, it follows that there is a $\mathbf{PCP}_{1, \frac{1}{2}}(\log(n) + A \log \log(n), B)$ verifier V for 3-SAT, for some constants $A, B > 0$.

Assume we are given the input E3-CNF formula ϕ . Fix any $\epsilon > 0$ and apply Lemma 2.4 to ϕ , obtaining $l = \mathcal{O}(2^{\epsilon n})$ instances ϕ_i , $i = 1, \dots, l$, each with size bounded by cn , for some constant $c > 0$, depending on ϵ . Let $R = \log(cn) + A \log \log(cn)$.

If we take ϕ_i as the input to the verifier V , then for each random binary string of length at most R , it reads at most B bits of the proof (we will assume without loss of generality that V always reads exactly B bits). Thus, V has access to N bits of the proof in total, where $N \leq B \cdot 2^R = B \cdot cn \log^A(cn)$. We create variables x_1, \dots, x_N corresponding to each of these bits. For a boolean string r of length R , we define a boolean function $f_r : \{0, 1\}^B \rightarrow \{0, 1\}$, with its arguments corresponding to relevant variables x_{i_1}, \dots, x_{i_B} that can be read from the proof. f_r is defined as follows, for every $S \in \{0, 1\}^B$ it evaluates to true on S if and only if V accepts ϕ_i using S as the bits from the proof. Every f_r can be written as an equivalent E3-CNF formula F_r (possibly adding some constant number of variables). Let C be an upper bound (depending on B) for the number of clauses for each such formula. Let ϕ'_i be a formula combined by AND of all F_r . It has the number of clauses bounded by $C \cdot 2^R$ and the following holds.

If ϕ_i is satisfiable, then there is a proof for which V always accepts. By taking the corresponding values from this proof and setting them to x_1, \dots, x_N , we obtain a satisfying assignment of ϕ'_i .

For the proof in the other direction, if ϕ_i is not satisfiable, then for any proof for at least half of all binary strings r the verifier V rejects ϕ_i . Consider an arbitrary proof and set the variables x_i to the corresponding values from this proof. Observe, that for the binary strings r for which V rejects ϕ_i at least one clause in F_r at least one clause is false. Consequently, at least $\frac{1}{2} \cdot 2^R$ clauses of ϕ'_i are false, i.e., a fraction of at least $\frac{1}{2C}$ of the total number of clauses of ϕ'_i .

Finally, let us put $c = A + 1$, and $r = 1 - \frac{1}{2C}$ and suppose there is a $\mathcal{O}^*(2^{\mathcal{O}(\frac{m}{\log^{A+1}(m)})})$ -time algorithm for GAP E3-SAT $_{[1 - \frac{1}{2C}, 1]}$. Using this algorithm, we can check whether ϕ_i is satisfiable in time $2^{\mathcal{O}(n)}$ and therefore check whether the initial formula ϕ is satisfiable in time $\mathcal{O}^*(2^{\epsilon n} + 2^{\epsilon n} \cdot 2^{\mathcal{O}(n)}) = \mathcal{O}^*(2^{\epsilon n})$. As ϵ is an arbitrary positive number, we obtain a contradiction with the ETH. \square

Next, we inspect the chain of three textbook NP-hardness reductions [42], that starts with GAP E3-SAT and ends with GAP MAX CUT. All of them produce an output instance of linear size in terms of the size of the input instance and at the same time preserve gaps. Combining them with the previous theorem, we will establish (almost) subexponential hardness of GAP MAX CUT.

Theorem 3.2.

$$\begin{aligned} \text{GAP E3-SAT}_{[\alpha,\beta]} &\leq_P^{\text{lin}} \text{GAP E4-NAE-SAT}_{[\alpha,\beta]} \leq_P^{\text{lin}} \\ \text{GAP E3-NAE-SAT}_{[\frac{1+\alpha}{2}, \frac{1+\beta}{2}]} &\leq_P^{\text{lin}} \text{GAP MAX CUT}_{[\frac{16+\alpha}{18}, \frac{16+\beta}{18}]} \end{aligned}$$

Proof. For simplicity, the proof of Theorem 3.2 is split into the following four separate lemmas.

Lemma 3.3.

$$\text{GAP E3-SAT}_{[\alpha,\beta]} \leq_P^{\text{lin}} \text{GAP E4-NAE-SAT}_{[\alpha,\beta]}$$

Proof. Given an E3-CNF formula $\phi = C_1 \wedge \dots \wedge C_m$, we create a new E4-CNF formula $\phi' = C'_1 \wedge \dots \wedge C'_m$ by adding a new variable z to every clause. That is, if $C_i = l_1 \vee l_2 \vee l_3$, then we set $C'_i = l_1 \vee l_2 \vee l_3 \vee z$.

If some assignment satisfies k clauses of ϕ , then by additionally setting $z = 0$ the corresponding k clauses of ϕ' are NAE-satisfied. In the other direction, if some assignment φ NAE-satisfies k clauses of ϕ' , then its negation $\tilde{\varphi}$ also NAE-satisfies k clauses of ϕ' . W.l.o.g. assume that φ sets z to false, which means that φ restricted to the variables of ϕ satisfies k clauses of ϕ . Consequently we obtain a gap preserving reduction and the theorem follows. \square

Lemma 3.4.

$$\text{GAP E4-NAE-SAT}_{[\alpha,\beta]} \leq_P^{\text{lin}} \text{GAP E3-NAE-SAT}_{[\frac{1+\alpha}{2}, \frac{1+\beta}{2}]}$$

Proof. Given an E4-CNF formula $\phi = C_1 \wedge \dots \wedge C_m$, we add m new variables z_1, \dots, z_m , and replace every clause $C_i = l_1 \vee l_2 \vee l_3 \vee l_4$ with the following two clauses: $C'_i = l_1 \vee l_2 \vee z_i$, $C''_i = l_3 \vee l_4 \vee \neg z_i$, obtaining an E3-CNF formula ϕ' . In the following we show that one can NAE-satisfy at least k clauses of ϕ if and only if one can NAE-satisfy at least $m + k$ clauses of ϕ' , which suffices to prove the theorem.

First, observe that if an assignment NAE-satisfies C_i , then by setting z_i appropriately both C'_i and C''_i become NAE-satisfied. On the other hand if an assignment does not NAE-satisfy C_i , then setting z_i to an arbitrary value NAE-satisfies exactly one clause out of C'_i and C''_i . Consequently if there is an assignment, which NAE-satisfies k clauses of ϕ , then it can be extended to an assignment of variables of ϕ' , which NAE-satisfies $m + k$ variables of ϕ' .

Let us assume that there is an assignment φ , which NAE-satisfies at least $m + k$ clauses of ϕ' . Note that there is a set I of at least k indices i , such that φ NAE-satisfies both C'_i and C''_i . Consider a fixed $i \in I$ and w.l.o.g. assume that φ sets z_i to false. As φ NAE-satisfies both C'_i and C''_i , we infer that φ sets at least one of the literals l_1, l_2 to true, and at least one of the literals l_3, l_4 to false. Consequently φ restricted to the variables of ϕ NAE-satisfies all the clauses C_i for $i \in I$, which finishes the proof of the theorem as $|I| \geq k$. \square

Lemma 3.5.

$$\text{GAP E3-NAE-SAT}_{[\alpha,\beta]} \leq_P^{\text{lin}} \text{GAP MULTIGRAPH MAX CUT}_{[\frac{3+2\alpha}{6}, \frac{3+2\beta}{6}]}$$

Proof. Given an E3-CNF formula $\phi = C_1 \wedge \dots \wedge C_m$, let n_i be the number of occurrences of x_i in ϕ , both in the positive and negative form. We construct a multigraph G with $2n$ vertices, and $6m$ edges as follows. For each variable x_i , we create two vertices corresponding to literals x_i , $\neg x_i$, and add exactly n_i edges between them. Finally, for every clause C_i we add a triangle connecting the vertices corresponding to its literals. We will prove that there is an assignment which NAE-satisfies at least k clauses of ϕ if and only if G admits a cut of size at least $3m + 2k$.

Let us assume that there is an assignment which NAE-satisfies k clauses of ϕ . We put vertices of G , that correspond to literals evaluated to true, on one side of the cut, and remaining literals to the other side of the cut. Every edge connecting x_i and $\neg x_i$ is clearly in the cut, and all of them contribute $\sum_i n_i = 3m$ to the cut. Furthermore, every triangle corresponding to a NAE-satisfied clause has exactly two edges cut, all of them contribute $2k$ to the cut. Thus, we obtain the cut of size at least $3m + 2k$.

Let us assume that G has a cut of size at least $3m + 2k$. First, suppose that x_i and $\neg x_i$ are on the same side of the cut for some variable x_i . They contribute at most n_i edges to the cut, and we can move one of the vertices x_i or $\neg x_i$ to the other side of the cut without decreasing the number of edges in the cut. Thus, we can assume that variables are separated from their negations. The edges connecting x_i and $\neg x_i$ contribute exactly $3m$ to the cut. If a clause triangle is cut 2 times, then it corresponds to a NAE-satisfied clause, otherwise the clause is not NAE-satisfied as all its literals have the same value. Thus, we deduce that the considered cut corresponds to an assignment which NAE-satisfies k clauses of ϕ . \square

Lemma 3.6.

$$\text{GAP MULTIGRAPH MAX CUT}_{[\alpha, \beta]} \leq_P^{lin} \text{GAP MAX CUT}_{[\frac{2+\alpha}{3}, \frac{2+\beta}{3}]}$$

Proof. Given a multigraph G , for every edge $e = uv$ of G , we create two auxiliary vertices w_e, z_e , add three edges $uw_e, w_e z_e, z_e v$, and remove the edge uv , obtaining a resulting graph G' . We claim that G has a cut of size at least k if and only if G' has a cut of size at least $2m + k$.

Let C be a cut of G of size k . If $e = uv \notin C$, we put w_e and z_e to the side of the cut opposite to u . This way two of the new edges, i.e., uw_e and $z_e v$, are in the cut. If $e = uv \in C$, then putting w_e, z_e to the opposite sides of u, v , respectively, gives three edges in the cut. All together we obtain a cut of G' of size $2m + k$.

Let C be a cut of G' of size $2m + k$. Clearly, there are at least k edges uv of the original graph G , such that all the three edges $uw_e, w_e z_e$, and $z_e v$ are in the cut C . Note that for such an edge uv the vertex w_e is on the side opposite to u , z_e in on the side opposite to w_e and v is on side opposite to z_e , which means that v in on the side of the cut opposite to u . Consequently a restriction of C to vertices of G gives a cut of size k . \square

Lemmas 3.3, 3.4, 3.5, and 3.6 together prove Theorem 3.2. \square

Thus, we infer:

Theorem 3.7. *Unless ETH fails, there exist $c \in \mathbb{N}$, and $0 \leq \alpha < \beta \leq 1$ such that there is no $2^{\mathcal{O}(\frac{m}{\log^c(m)})}$ -time algorithm for $\text{GAP MAX CUT}_{[\alpha, \beta]}$.*

Finally, we modify the reduction by Garey et al. [27] from MAX CUT to OLA. The construction from [27] introduces a huge clique to the complement of the MAX CUT instance; this clique is supposed to separate in the ordering the sides of an optimum max-cut solution. The clique has to

be large enough so that the cost of any edge going over the clique eclipses the “noise” contribution that is given by internal ordering of parts on the left and on the right. For this reason the clique is chosen to be of size $\Theta(n^4)$, which imposes a significant blow-up in the instance size and immediately prevents from obtaining a desired subexponential hardness result.

This is precisely the moment where we exploit that we are working with an instance of MAX CUT with a gap, as starting with a gap instance, we can accommodate the noise in the gap, and therefore we only need the clique to be of linear size. Note that in the following theorem the number of vertices produced is linear, however the produced instance might be dense, hence we cannot use the notation $\leq_{\text{P}}^{\text{lin}}$ and formulate the reduction properties explicitly.

Theorem 3.8. *There is a polynomial time algorithm, which given an instance I of GAP MAX CUT $_{[\alpha, \beta]}$ with n vertices produces an instance I' of OLA with $\mathcal{O}(n)$ vertices, such that if I admits a cut with at least βm edges, then I' is a yes-instance, and if I does not have a cut with at least αm edges, then I' is a no-instance.*

Proof. Given the input graph G , we create a corresponding instance of OLA, G' as follows. We take a complement of G , i.e., \overline{G} and add a disjoint clique C of size Mn , for $M = \lceil \frac{2}{\beta - \alpha} \rceil$ and fully connect it to \overline{G} .

We will show that G has a cut of size at least βm if and only if G' has a linear arrangement of cost at most $\binom{(M+1)n+1}{3} - \beta m \cdot Mn$, that will prove the equivalence of the instances.

Let (A, B) be a cut of size at least βm . We define a linear arrangement π as follows. First, we list the vertices of A , then we list the vertices of the clique C and then the vertices of B . The order of vertices inside A , B and C can be arbitrary. Then by counting the costs of all the edges in the clique on $(M+1)n$ vertices we obtain:

$$\sum_{uv \in E(G')} |\pi(u) - \pi(v)| + \sum_{uv \in E(G)} |\pi(u) - \pi(v)| = \binom{(M+1)n+1}{3} (*)$$

Now, the costs of each edge going across the clique, i.e., each edge in $E_{G'}(A, B)$, is at least Mn , thus $\sum_{uv \in E(G)} |\pi(u) - \pi(v)| \geq \beta m \cdot Mn$, and it implies the desired inequality.

In the other direction, let π be a linear arrangement of G' of cost at most $\binom{(M+1)n+1}{3} - \beta m \cdot Mn$. First, we will prove that there exists an optimum linear arrangement π , such that the vertices of C has to be in the same consecutive block of π . Before that, we will introduce notation regarding to linear arrangements, also used in other proofs.

We assume the vertices of G' are ordered from left to right according to a linear arrangement π . When we speak about the i -th vertex (from the left), we mean the vertex mapped to the number i by π . A set of vertices U is *consecutive in π* , if $\pi(U) = \{p, p+1, \dots, q-1, q\}$ for $p, q \in \mathbb{N}$. The set of all vertices that are to the left of every vertex from some set U is called *vertices to the left of U* and denoted by $L(U)$. Similarly, we define *vertices to the right of U* and denote them by $R(U)$. A *block of U* is any inclusion-wise maximal non-empty subset of U that is consecutive in π . The *left-most block of U* is the block of U whose vertices are mapped to the smallest values by π . The *second left-most block of U* is the first block of U to the right of the left-most block of U . The *inner block of U* is the set of all vertices from $V(G') \setminus U$ located simultaneously to the right the left-most block of U and to the left of the second left-most block of U (in the case when U forms a single block, the inner block does not exist).

Claim 3.9. *There exists an optimum linear arrangement π of G' such that the vertices of C are consecutive in π .*

Proof. Let us choose π to be an optimum linear arrangement of G' that minimize the number of vertices of G that lie between the vertices of C , i.e. $f(\pi) = |\{v \in V(G) : \pi(u) < \pi(v) < \pi(w), u, w \in V(C)\}|$ is the smallest possible. We claim that the vertices of C are consecutive in π , i.e. $f(\pi) = 0$.

Assume that this is not the case. Let X be the inner block of C and we consider the following two cases:

- $|E_G(L(X), X)| \leq |E_G(X, R(X))|$
- $|E_G(L(X), X)| > |E_G(X, R(X))|$

In each of them, we will create another linear arrangement that will contradict the assumptions we made.

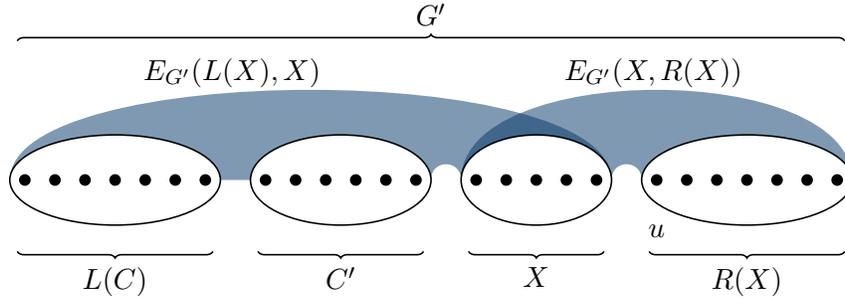


Figure 1: The situation in the proof of Claim 3.9.

If the first case occurs, let C' be the left-most block of C . We swap places of C' and X . Then, we consider how it affects the cost of π . First, we note that by (*), the optimality of π can be equivalently restated as maximizing the sum of cost of edges of G , i.e. π is optimal if $\sum_{uv \in E(G)} |\pi(u) - \pi(v)|$ is the largest possible. We claim that the value of this sum has not decreased. Clearly, the cost of edges with both endpoints in $V(G) \setminus X$ or X does not change. Hence, it is enough to inspect the edges with exactly one endpoint in X . Every such edge goes from X to $L(X)$ or from X to $R(X)$. An edge going from X to $L(X)$ has decreased its length by $|C'|$, and an edge going from X to $R(X)$ has increased its length by $|C'|$. Thus, as $|E_G(L(X), X)| \leq |E_G(X, R(X))|$, the overall contribution of those edges has not decreased, but the number of vertices between C has dropped by $|X|$, leading to the contradiction.

If the second case occurs, let C' be the second-left most block of C . Similarly we swap places of C' and X . Again, only affected edges are the ones that have exactly one endpoint in X . The inequality $|E_G(L(X), X)| > |E_G(X, R(X))|$ implies that $\sum_{uv \in E(G)} |\pi(u) - \pi(v)|$ has increased at least by $|X|$, contradicting the optimality of π .

Thus, we have proved that C has to be consecutive in π . □

Let π be the optimum linear arrangement from the claim above. We define a cut (A, B) , by taking A to be $L(C)$ and B to be $R(C)$. Then by (*) we know:

$$\beta m \cdot Mn \leq \sum_{uv \in E(G)} |\pi(u) - \pi(v)| =$$

$$\sum_{\substack{uv \in E(G), \\ u, v \in A}} |\pi(u) - \pi(v)| + \sum_{\substack{uv \in E(G), \\ u, v \in B}} |\pi(u) - \pi(v)| + \sum_{\substack{uv \in E(G), \\ u \in A, v \in B}} |\pi(u) - \pi(v)| < \\ nm + (Mn + n) \cdot |E_G(A, B)| \leq 2nm + Mn \cdot |E_G(A, B)|$$

Thus:

$$|E_G(A, B)| > \beta m - \frac{2m}{M} \geq \beta m - (\beta - \alpha) \cdot m = \alpha m$$

By the promise given by the gap problem we know that if G admits a cut greater than αm , then it actually admits a cut of size at least βm . \square

The proof of Theorem 1.5 follows immediately from Theorems 3.7 and 3.8.

4 Sparse reduction

We now introduce a polynomial-time Turing reduction from a gap version of the MIN BISECTION problem on d -regular graphs to OPTIMUM LINEAR ARRANGEMENT. Its key property is that the instances of the former problem result in instances of OPTIMUM LINEAR ARRANGEMENT with linear number of vertices and bounded degree. Even though the created instance of OPTIMUM LINEAR ARRANGEMENT is a multigraph, it does not cause any additional difficulties in further reductions described in Section 5. The reduction allows one to distinguish between instances of MIN BISECTION with at most αm edges and at least βm edges in the optimum cut for some fixed choice of $0 \leq \alpha < \beta \leq 1$ by solving the resulting instance of OPTIMUM LINEAR ARRANGEMENT. This relates OPTIMUM LINEAR ARRANGEMENT to the following hypothesis:

Hypothesis 1.2. *There exist $0 \leq \alpha < \beta \leq 1$, and an integer $d > \frac{4}{\beta - \alpha}$, such that there is no $2^{o(n)}$ -time algorithm for GAP MIN BISECTION(d) $_{[\alpha, \beta]}$.*

The main result of this section is:

Theorem 1.6. *Unless Hypothesis 1.2 fails, there exists an integer $d \in \mathbb{N}$ such that there is no $2^{o(n)}$ -time algorithm solving OLA $_{\leq}(d)$ in multigraphs.*

We first describe a transformation $T(\cdot)$ from an instance G of MIN BISECTION to an instance of OPTIMUM LINEAR ARRANGEMENT that forms the key component of our reduction. Then, we introduce several technical claims about its properties. Finally, we prove Theorem 1.3 by showing how to decide the instances of GAP MIN BISECTION based on the cost of the optimum arrangement of $T(G)$. The fact that our reduction exhibits only a linear increase in the size of the instance is crucial in achieving the $2^{\Omega(n)}$ bound.

The result of transformation $T(\cdot)$ is influenced by several parameters. The choice of their values is deferred to the proof of Theorem 1.3. Consider an instance G of the MIN BISECTION problem, where G is a d_G -regular graph. Assume $V(G) = \{v_1, \dots, v_n\}$. The transformation produces a graph $G' := T(G)$ with the vertex set $\{v_1, \dots, v_n, x_1, \dots, x_{Z \cdot \lceil \varphi n \rceil}\}$, where $Z \in \mathbb{N}$ and $\varphi \in (0, 1)$ are constants chosen later. Note that G' contains the vertices of G . Indeed, we are going to construct the edge-set in such a way that G is an induced subgraph of G' . It is actually convenient to introduce notation for some of the induced subgraphs of G' . The subgraph with the vertex set $\{x_1, \dots, x_{Z \cdot \lceil \varphi n \rceil}\}$ is denoted by H . The graph H is (arbitrarily) divided into Z disjoint induced subgraphs H_i of size $\lceil \varphi n \rceil$ each, for some constants $Z \in \mathbb{N}$ and $\varphi \in (0, 1)$.

The result of the transformation is illustrated in Figure 2. The edge-set of G' is constructed as follows:

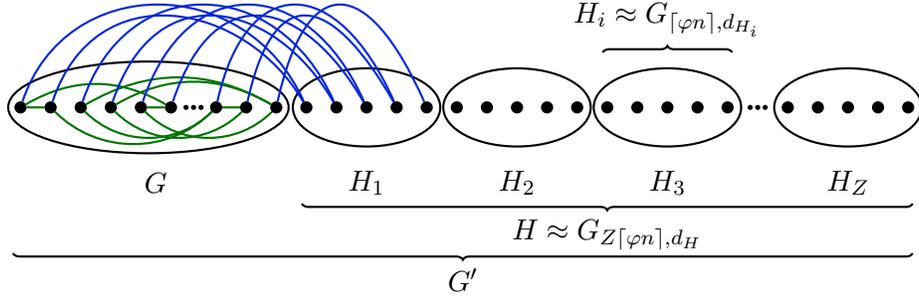


Figure 2: The resulting instance $G' = T(G)$ after applying the reduction. The original graph G is an induced subgraph of G' with its edges shown in green. The edges of the bipartite graph added between $V(G)$ and $V(H_1)$ are shown in blue.

- The induced subgraph of G' on $\{v_1, \dots, v_n\}$ is G .
- We construct a d_H -regular expander $G_{|H|, d_H}$ using Theorem 2.2 satisfying $h(G_{|H|, d_H}) \geq p_H$, and add its edges on the vertices of H (the value p_H will be determined later).
- For each $i \in \{1, \dots, Z\}$ we construct a d_{H_i} -regular expander $G_{|H_i|, d_{H_i}}$ using Theorem 2.2 satisfying $h(G_{|H_i|, d_{H_i}}) \geq p_{H_i}$, and add its edges on the vertices of H_i (the value p_{H_i} will be determined later).
- For each $i \in \{1, \dots, Z\}$ we add a bipartite graph on parts $V(G)$ and $V(H_i)$ such that all vertices of $V(G)$ have degree 1 in this bipartite graph and the degrees of vertices from $V(H_i)$ differ by at most 1. We denote the maximum degree of the $V(H_i)$ part of this added bipartite graph by $\Delta_{H, G}$. It is at most $\lceil \frac{1}{\varphi} \rceil$.

Note that we are constructing a multigraph, that is when an edge is to be added several times in the construction process, we keep all its copies.

In the following part, we give the proof of correctness of the transformation, as well as determine the parameters driving the reduction. In the proof, we first show that the vertices of H have to be consecutive in an optimum linear arrangement. Next, we show that actually vertices of each small expander H_i are consecutive in such an ordering. This is crucial when we analyze the change of the cost of the ordering when moving a vertex of G from one side of H to the other, and in turn prove that in an optimum ordering the parts of G to the left and to the right of H are almost of the same size. Interestingly, in our reduction we have to use the hypothetical oracle solving the decision version of OPTIMUM LINEAR ARRANGEMENT to find the cost of an optimum ordering of an auxiliary graph by using binary search.

The constructed G' is influenced (apart from the input graph G) by our choice of parameters Z, φ, p_H and p_{H_i} , which in turn influence d_H and d_{H_i} by Theorem 2.2. The lemmas below impose a particular structure on the optimum linear arrangement of G' , provided certain inequalities between these parameters are satisfied. Eventually, the lemmas are employed in the proof of the main theorem of this section.

The following technical *Swapping Lemma* establishes a condition on degrees in two consecutive sets of G' under which the swapping of the two sets results in a decreased cost of the ordering.

Lemma 4.1 (Swapping Lemma). *Consider an ordering π of any finite graph G' . Assume that the sets $X, Y \subseteq V(G')$ are consecutive and X immediately precedes Y . Let $L := L(X)$ and $R := R(Y)$. Assume*

- the value P_X upper bounds the degree of vertices from X in the induced bipartite subgraph $G'[L, X]$,
- P_C is an upper bound on the maximum degree of $G'[X, Y]$, and
- P_Y is an upper bound on the degree of a vertex from Y in $G'[Y, R]$.

Finally, let p be a lower bound on the average degree of a vertex from X in $G'[X, R]$. Then the inequality $p > P_X + 2P_C + P_Y$ implies that swapping the vertices of X with the vertices of Y in the order specified by π results in a decrease in the cost of the ordering.

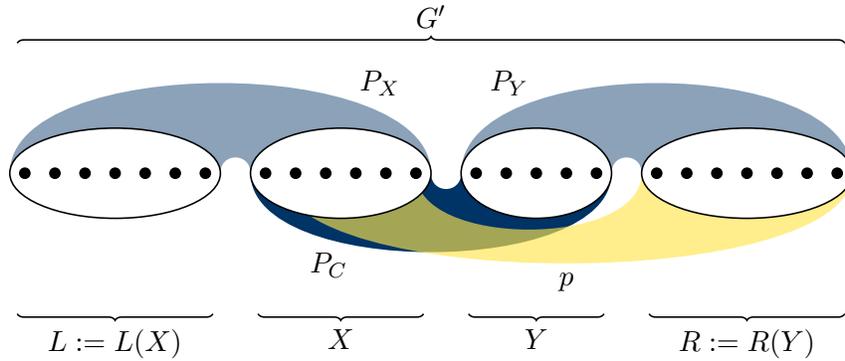


Figure 3: The vertex set of G' is partitioned into four sets, L, X, Y , and R , in Lemma 4.1. The bounds P_X, P_Y, P_C (upper bounds), and p (a lower bound) on the degrees of the induced bipartite subgraphs are also shown.

Proof. The situation is illustrated in Figure 3. The length of all edges connecting a pair of vertices from one of the sets L, X, Y, R remains unchanged after swapping X and Y in the ordering. The same holds for edges connecting L with R . The length of each edge connecting X and Y increases by at most $|X| + |Y| \leq 2 \max\{|X|, |Y|\}$. The cost of each edge connecting X and L increases by at most $|Y|$. Similarly, the cost of each edge connecting Y and R increases by at most $|X|$. On the other hand, the edges connecting X and R are shortened, each by $|Y|$. The upper bounds on maximum degrees and the lower bound on average degree from the statement of the lemma allow us to lower bound the decrease in the cost of the ordering after the swap is performed. For example, the decrease in total cost of the edges connecting X with R is at least $p|X||Y|$. The decrease in cost after swapping is at least

$$p|X||Y| - 2 \min\{|X|, |Y|\} P_C \max\{|X|, |Y|\} - |X|P_X|Y| - |Y|P_Y|X|,$$

which is equal to

$$|X||Y|(p - 2P_C - P_X - P_Y).$$

Assuming the inequality from the lemma, this is strictly larger than zero. \square

We now make several claims about the optimum ordering π of $G' := T(G)$, where G is a d_G -regular graph. Recall that $Z, \varphi, p_H, p_{H_i}, d_H$, and d_{H_i} are the parameters of the transformation $T(\cdot)$ still to be determined.

Lemma 4.2. *If $p_H > 3\Delta_{H,G} + 3Z + d_G$ and π is an optimum linear arrangement of G' , then $V(H)$ is consecutive in π .*

Proof. Suppose $V(H)$ is not consecutive in π . Consider the left-most block of $V(H)$ and denote its elements by X . We can assume that $|X| \leq \frac{|H|}{2}$ – otherwise we take the right-most block of $V(H)$ and proceed with a mirrored version of the following argument. Denote by Y the inner block of $V(H)$ and set $L := L(X), R := R(Y)$. The following choice of values satisfies the assumptions on degree upper-bounds of the Lemma 4.1:

$$P_X := \Delta_{H,G}, \quad P_Y := d_G + Z, \quad P_C := \Delta_{H,G} + Z.$$

Since H is an expander, $|X| \leq |H|/2$ and $H \setminus X \subseteq R$, we take $p = p_H$. It remains to show the inequality from the statement of the Swapping Lemma. We have:

$$p = p_H > 3\Delta_{H,G} + 3Z + d_G = P_X + 2P_C + P_Y.$$

Thus, we can swap X and Y and decrease the cost of the ordering. This contradicts the optimality of π . \square

Lemma 4.3. *Let $i \in \{1, \dots, Z\}$. If $p_{H_i} > d_{H_{i+1}} + 4d_H + 2\Delta_{H,G}$, π is an optimum linear arrangement of G' , and for each $i' < i$ the vertices of $H_{i'}$ are consecutive in π , then the vertices of H_i are consecutive in π . (For the purposes of this lemma, we set $d_{H_{Z+1}} := 0$.)*

Proof. Assume H_i not to be consecutive and denote by X the left-most block of $V(H_i)$ in π . Similarly to the situation in the proof of Lemma 4.2, we can assume $|X| \leq |H_i|/2$. (Otherwise we consider the right-most block instead and proceed with a mirrored version of the argument.) We show that moving X to the right decreases the cost of the arrangement.

Denote by u the vertex positioned by π immediately to the right of X . Due to Lemma 4.2, we know $u \notin V(G)$. Therefore, $u \in V(H_j)$ for $j \neq i$. We distinguish two cases: either $j < i$ or $j > i$.

Suppose that $u \in V(H_j)$ for $j < i$. Note that H_j is consecutive in π . We set $Y := V(H_j), L := L(X)$, and $R := R(Y)$. Again, we employ the Swapping Lemma. The following degree upper-bounds satisfy its assumptions:

$$P_X := \Delta_{H,G} + d_H, \quad P_Y := \Delta_{H,G} + d_H, \quad P_C := d_H.$$

Since H_i is an expander, $|X| \leq |H_i|/2$ and $V(H_i) \setminus X \subseteq R$, we can set the average degree lower-bound p to p_{H_i} . By the inequality from the statement of this lemma, we have

$$p = p_{H_i} > 4d_H + 2\Delta_{H,G} = P_X + 2P_C + P_Y.$$

Thus, the inequality from the Swapping Lemma holds and we can use it to decrease the cost of ordering, contradicting the optimality of π .

Suppose therefore that $u \in V(H_j)$ for $j > i$. We now use Lemma 4.1 again to move the block X one position to the right, effectively swapping X and $Y := \{u\}$. We set $L := L(X), R := R(Y)$. This time, we set:

$$P_X := \Delta_{H,G} + d_H, \quad P_Y := \Delta_{H,G} + d_H + d_{H_{i+1}}, \quad P_C := d_H.$$

Similarly to the previous cases, we set $p := p_{H_i}$. The inequality from the Swapping Lemma is again satisfied:

$$p = p_{H_i} > 4d_H + 2\Delta_{H,G} + d_{H_{i+1}} = P_X + 2P_C + P_Y.$$

Once more, we get a contradiction with the optimality of π . □

Due to Lemma 4.2 we know that an optimum linear arrangement of G' places vertices of H consecutively, assuming the inequalities from its statement are satisfied. Furthermore, by iterating Lemma 4.3 we get that within H , the vertices of each H_i are grouped together in the arrangement. To be precise, in the optimum ordering the subgraphs H_i are placed in the order $H_{\ell_1}, H_{\ell_2}, \dots, H_{\ell_Z}$, where $(\ell_1, \ell_2, \dots, \ell_Z)$ is a permutation of $\{1, 2, \dots, Z\}$. The vertices of G can thus be only to the left of H or to its right. The next lemma shows that H divides the graph G into two roughly equal parts.

Lemma 4.4. *Assume G' has been constructed by the transformation $T(\cdot)$ with parameters satisfying the inequalities from the statements of Lemmas 4.2 and 4.3 as well as $Z\varphi \geq 2$. Moreover, assume $\gamma = 3\varphi d_G$. Consider an optimum linear arrangement π of G' and set $A := L(H), B := R(H)$. Then $||A| - |B|| \leq \gamma n$.*

Proof. Assume the imbalance $||A| - |B||$ is strictly bigger than γn . Without loss of generality, assume $|A| > |B|$. We consider the vertex u such that $\pi(u) = 1$ (i.e., the one placed on the left side of the arrangement).

Moving u to the right-most position results in the following changes in the cost of the arrangement. The cost associated with the edges of G might be increased by at most $d_G(Z\lceil\varphi n\rceil + n) \leq d_G(Z\varphi n + Z + n)$. Note that for sufficiently large n we have $\frac{Z\varphi n}{2} \geq Z$ and by the assumption of the Lemma we have $\frac{Z\varphi n}{2} \geq n$, therefore the cost associated with the edges of G increases by at most $d_G(Z\varphi n + Z + n) \leq 2d_G Z\varphi n$.

In addition to this, the vertex u is connected to precisely one vertex v_i of each H_{ℓ_i} . Before moving u , the edge cost of uv_i was $|A| + (i-1)\lceil\varphi n\rceil + r_i$, for some $0 \leq r_i \leq \lceil\varphi n\rceil - 1$ and after, it becomes $|B| + ((Z-i+1)\lceil\varphi n\rceil - r_i)$, thus the contribution of all those edges has changed by:

$$\begin{aligned} (|B| - |A|)Z + \sum_{i=1}^Z ((Z-i+1)\lceil\varphi n\rceil - r_i) - \sum_{i=1}^Z ((i-1)\lceil\varphi n\rceil + r_i) &= (|B| - |A|)Z + \sum_{i=1}^Z (\lceil\varphi n\rceil - 2r_i) < \\ -\gamma nZ + Z\lceil\varphi n\rceil &\stackrel{\gamma=3d_G\varphi}{=} -3d_G Z\varphi n + Z\lceil\varphi n\rceil < \stackrel{Z\lceil\varphi n\rceil < d_G Z\varphi n}{=} -2d_G Z\varphi n \end{aligned}$$

Note that in the last inequality we have used the assumption $\lceil\varphi n\rceil < d_G \varphi n$ which holds for $d_G \geq 2$ and sufficiently large n , as φ is a constant.

Combining this cost change with the cost increase of edges of G being at most $2d_G Z\varphi n$, we obtain that moving u to the right-most position causes the decrease in the cost of π , leading to a contradiction. □

Being equipped with all the required tools, we are ready to prove Theorem 1.6.

Proof of Theorem 1.6. We prove the theorem by introducing a Turing reduction from GAP MIN BISECTION(d) $_{[\alpha,\beta]}$ to OLA $_{\leq}(d')$, where

$$d > \frac{4}{\beta - \alpha} \tag{1}$$

and d' is some constant depending on d , α , and β . The reduction proceeds in the following way.

Given an instance G of GAP MIN BISECTION(d) $_{[\alpha, \beta]}$, we create an equivalent instance (G', k) of $\text{OLA}_{\leq}(d')$ based on the transformation $T(\cdot)$ applied to G . The value k will depend on, among other parameters, the cost of the optimal arrangement of the expander H (recall this is an induced subgraph of G'). As mentioned in the beginning, the reduction is a Turing reduction. It is therefore equipped with an oracle solving the decision version of OLA, which in turn is enough to find the cost of an optimal ordering by using binary search (we will use it to compute $\text{OLA}(H)$).

We start by establishing the parameters of the transformation in the following order.

- First, we set $\gamma := \frac{\beta - \alpha}{4}$ and $\varphi := \frac{\gamma}{3d_G}$. Note that $\gamma, \varphi \in (0, 1)$ and γ, φ satisfy the condition $\gamma = 3\varphi d_G$ from Lemma 4.4.
- Next, we set the value of Z to the following integer:

$$Z := \lceil \frac{2(2\alpha + 1)}{(\beta - \alpha)\varphi} \rceil.$$

In particular $Z \geq \frac{2}{\varphi}$ satisfying the condition from Lemma 4.4. Moreover, the following inequality holds:

$$2(2\alpha + 1) \leq (\beta - \alpha)Z\varphi. \quad (2)$$

- By construction $\Delta_{H,G} \leq \lceil \frac{1}{\varphi} \rceil = \lceil \frac{3d_G}{\gamma} \rceil$.
- Next, we define p_H , which in turn determines the value of d_H by Theorem 2.2:

$$p_H := 3\Delta_{H,G} + 3Z + d_G + 1.$$

The additive term $+1$ is just to make sure the inequality from Lemma 4.2 is strict and it could be replaced by any positive constant.

- Finally, we set the values p_{H_i} :

$$p_{H_i} := d_{H_{i+1}} + 4d_H + 2\Delta_{H,G} + 1,$$

in the decreasing order $i = Z, \dots, 1$, where for simplicity we denote $d_{H_{Z+1}} = 0$. Note that the value of p_{H_i} determines the value of d_{H_i} via Theorem 2.2.

Observe that with the above definition of all the constants we satisfy all the requirements of Lemmas 4.2, 4.3, 4.4.

Let G' be the result of the transformation $T(G)$ with the above choice of parameters. The value k is then set as follows:

$$k := \text{OLA}(H) + \alpha m \cdot (Z \lceil \varphi n \rceil + n) + m \cdot \frac{n}{2} + \left(\left(\frac{n}{2} + 1 \right) \frac{n}{2} Z + n \sum_{i=1}^Z i \lceil \varphi n \rceil \right). \quad (3)$$

Note that to compute the value of k we need use the oracle solving the decision version of T in a binary search routine. It remains to show that G has a bisection of size at most αm if and only if G' has an optimum linear arrangement of size at most k .

Let us assume G has a bisection with at most αm edges. We claim that k is an upper bound on the cost of an optimum linear arrangement of G' . This is because it accounts for all costs associated with an ordering of G' constructed from the optimum bisection of G . Denote by A, B the partition of $V(G)$ corresponding to an optimum bisection of G . We create an ordering π that first lists all vertices of A , then the vertices of H in the order of an optimum linear arrangement of H , and finally the vertices of B .

The first term of (3) is the cost of all edges inside H . The second term upper bounds the cost of edges of G between A and B : there are at most αm of them and we upper bound the cost of every such edge by $|V(G')| = (Z\lceil\varphi n\rceil + n)$. In the third term, we account for the cost of edges within A and within B . There are at most m of them and since $|A| = |B| = \frac{n}{2}$ every such edge has cost at most $\frac{n}{2}$. The last term is an upper bound on the cost of edges connecting G and H . Every vertex v of G has an edge to exactly one vertex of each H_{ℓ_i} . If $v \in A$ we may bound its cost by $j(v) + i\lceil\varphi n\rceil$, where $j(v)$ is the length of the part of the ordering from v to the first vertex in H . We first count the contribution of the $j(v)$ -terms in the above expression for all choices of $v \in A$. Since $|A| = n/2$, summing over all $v \in A$ and $i = 1, \dots, Z$ we get $\sum_{j=1}^{|A|} jZ = \frac{(\frac{n}{2}+1)\frac{n}{2}}{2}Z$. The situation is analogous for B . The last term of (3) is obtained by summing the remaining edge costs $i\lceil\varphi n\rceil$ for all $i = 1, \dots, Z$ and $v \in G$. This proves the claimed upper bound.

In the other direction, we start with assuming that the graph G' has an optimum arrangement π of cost at most k . The aim is to prove that G has a bisection of size at most αm . Lemma 4.2, Lemma 4.3, and Lemma 4.4 together impose a certain structure on π . Particularly, the vertices of H are placed together in π . We use this fact to construct a bisection of G . Set $A := L(H)$ and $B := R(H)$, and if $|A| < |B|$, then replace A with B . Note that (A, B) is a partition of $V(G)$ and these two sets might have different sizes with the imbalance bounded by Lemma 4.4. We now bound the number of edges between A and B . To this end we lower bound the cost of π in terms of $|E_{G'}(A, B)|$. Specifically, it is at least:

$$\text{OLA}(H) + |E_{G'}(A, B)| \cdot Z\lceil\varphi n\rceil + \left(\left(\frac{n}{2} + 1\right)\frac{n}{2}Z + n \sum_{i=1}^Z (i-1)\lceil\varphi n\rceil \right). \quad (4)$$

There, the first term of (4) accounts for the cost of the edges of H , as it is lower bounded by $\text{OLA}(H)$. The second term of (4) is a lower bound on the cost of edges of G going across the partition (A, B) . There are $|E_{G'}(A, B)|$ of such edges and each contributes at least $Z\lceil\varphi n\rceil$ to the cost. Recall that vertices of H must have the following order $H_{\ell'_1}, H_{\ell'_2}, \dots, H_{\ell'_Z}$, where $(\ell'_1, \ell'_2, \dots, \ell'_Z)$ is some permutation of $\{1, 2, \dots, Z\}$. The third term lower bounds the cost of edges connecting G to H . Every vertex v of G has an edge to exactly one vertex from $H_{\ell'_i}$. Similarly to the analysis above, we lower bound the cost of such an edge by $j(v) + (i-1)\lceil\varphi n\rceil$. The contribution of $j(v)$'s from the above expression for all choices of $v \in G$ is equal to $\left(\sum_{j=1}^{|A|} j + \sum_{j=1}^{|B|} j \right) Z \geq \left(2 \sum_{j=1}^{\frac{n}{2}} j \right) Z \geq \left(\frac{n}{2} + 1\right)\frac{n}{2}Z$. The remaining part is obtained by summing $(i-1)\lceil\varphi n\rceil$ over all $v \in G$ and $i = 1, \dots, Z$. Comparing (3) with (4) we obtain:

$$\begin{aligned} |E_{G'}(A, B)| \cdot Z\lceil\varphi n\rceil &\leq \alpha m \cdot (Z\lceil\varphi n\rceil + n) + m \cdot \frac{n}{2} + n \sum_{i=1}^Z i\lceil\varphi n\rceil - n \sum_{i=1}^Z (i-1)\lceil\varphi n\rceil = \\ \alpha m \cdot Z\lceil\varphi n\rceil + \frac{2(2\alpha + 1)mn}{4} + n \sum_{i=1}^Z \lceil\varphi n\rceil &\stackrel{\text{by (2)}}{\leq} \alpha m \cdot Z\lceil\varphi n\rceil + \frac{(\beta - \alpha)m \cdot Z\varphi n}{4} + n \cdot Z\lceil\varphi n\rceil \stackrel{\varphi n \leq \lceil\varphi n\rceil}{\leq} \end{aligned}$$

$$\alpha m \cdot Z[\varphi n] + \frac{(\beta - \alpha)m}{4} \cdot Z[\varphi n] + n \cdot Z[\varphi n] \stackrel{n = \frac{2}{d_G} m < \frac{\beta - \alpha}{2} m \text{ by (1)}}{<} \frac{\alpha + 3\beta}{4} m \cdot Z[\varphi n].$$

We infer that $|E_{G'}(A, B)| < \frac{\alpha + 3\beta}{4} m$. Next, we create a bisection (A', B') of G as follows. Let C be a set of $\frac{|A| - |B|}{2}$ arbitrary vertices of A , then we put $A' = A - C$, $B' = B \cup C$. By moving these vertices we get at most $\frac{\gamma m}{2} \cdot d_G = \gamma m$ additional edges in the cut, thus:

$$|E_{G'}(A', B')| \leq |E_{G'}(A, B)| + \gamma m < \frac{\alpha + 3\beta}{4} m + \frac{\beta - \alpha}{4} m = \beta m.$$

Therefore, we have $|E_{G'}(A', B')| < \beta m$. Since the problem considered is a gap problem, we know that G admits either a bisection of size at most αm or there is no bisection of size less than βm . Therefore, we conclude the instance admits a bisection of size at most αm . \square

5 Lower Bounds for Minimum Fill-in and Other Completion Problems

In this section we prove Theorems 1.1 and 1.3, that is prove conditional lower bounds (under ETH and under Hypothesis 1.2) for parameterized completion problems such as MINIMUM FILL-IN, CHAIN COMPLETION, PROPER INTERVAL COMPLETION, INTERVAL COMPLETION, THRESHOLD COMPLETION, TRIVIAL PERFECT COMPLETION. As a starting point we use Theorems 1.5 and 1.6, hence our goal is to transform an instance of OLA into an instance of a graph completion problem. The main reduction of this section, which transforms an instance of OLA to CHAIN COMPLETION is a slight modification of the reduction of Yannakakis [44], with the only difference that on bounded degree instance of OLA we obtain linear number of vertices in the final instance of CHAIN COMPLETION. This fact is crucial to prove Theorem 1.3 while for proof of Theorem 1.1 it is enough to use the original version of Yannakakis reduction.

Definition 5.1. *A bipartite graph (A, B, F) with vertices $A \uplus B$ and edges F is a chain graph if the set of vertices A (called left side) can be ordered v_1, v_2, \dots, v_n (called left order) in such a way that $N(v_1) \subseteq N(v_2) \subseteq \dots \subseteq N(v_n)$.*

In the CHAIN COMPLETION problem given a bipartite graph (A, B, F) one is asked to add a minimum number of edges $F' \subseteq A \times B \setminus F$ such that $(A, B, F \cup F')$ is a chain graph.

Lemma 5.2. *There is a polynomial time algorithm, which given an instance $I = (G = (V, E), k)$ of OPTIMUM LINEAR ARRANGEMENT creates an equivalent instance $I' = (G' = (A, B, F), k')$ of CHAIN COMPLETION, such that the number of vertices of G' is bounded by $\mathcal{O}(\Delta_G \cdot |V|)$, where Δ_G is the maximum degree of G . The reduction works even if G is a multigraph.*

Proof. As the left side of G' we take $A = V$. For each vertex $v \in V$ create a set of Δ_G new vertices $S_v = \{v_e : e \in \delta_G(v)\} \cup \{v_i : \deg(v) < i \leq \Delta_G\}$. We define B as the union of all the sets S_v , thus B contains exactly $\Delta_G \cdot |V|$ vertices. The set of edges F is constructed as follows. For each $w \in S_v$ we add to F an edge vw . Additionally, for each $v_e \in B$, where $e \in E, e = uv$ we add to F an edge uv_e , so that the vertex v_e is of degree exactly two in G' . The described transformation is depicted in Fig. 4. To finish the construction of I' we define $k' = k + \Delta_G \frac{n(n-1)}{2} - 2|E|$.

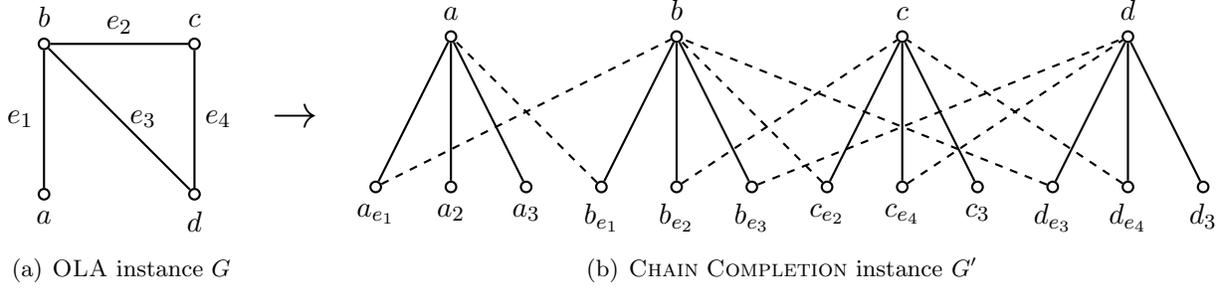


Figure 4: Transformation of OPTIMUM LINEAR ARRANGEMENT to CHAIN COMPLETION

For a given ordering π of the vertices of the graph $G = (V, E)$ denote by $C(G, \pi)$ the cost of arrangement induced by this ordering (i.e., $C(\pi, G) = \sum_{uv \in E} |\pi(u) - \pi(v)|$). For an ordering σ of the left side A of the bipartite graph $G' = (A, B, F)$ denote by $E(G', \sigma)$ the number of edges that we should add to obtain a minimal chain graph in which the left order coincides with σ . We will prove the following claim.

Claim 5.3. *For any ordering π of V (or equivalently A) we have $E(G', \pi) = C(G, \pi) + \Delta_G \frac{n(n-1)}{2} - 2|E|$.*

Proof. Let us inspect what is the number of edges in a minimal chain bipartite graph G'' which has left order $\pi = (v_1, v_2, \dots, v_n)$ and contains $G' = (A, B, F)$ as a subgraph. We know that $N(v_i) \subseteq N(v_j)$ for any $i < j$, which means that each $x \in S_{v_i}$ must be connected to all vertices from the set $\{v_i, v_{i+1}, \dots, v_n\}$. So it means that each vertex from S_{v_i} is connected to at least $(n+1) - i$ vertices from the set A . Moreover, by minimality of G'' the vertex x is connected to exactly this number of the vertices if x does not correspond to any edge in G , i.e., when x is of degree exactly one in G' . If a vertex $w_i \in S_{v_i}$ corresponds to some edge $e \in G$ with endpoints v_i, v_j (note there could be several edges with equal endpoints as we are working with multigraphs) then the vertex w_i is connected to vertices v_i, v_j in the graph G' . Hence, in a minimal chain graph G'' the degree of w_i is either $(n+1) - i$ or $(n+1) - j = (n+1) - i + (i - j)$ depending on whether $i < j$ or $i > j$. Note that there is a second vertex $w_j \in S_{v_j}$ which also corresponds to the edge e . The degrees of w_i and w_j in G'' both are equal to $(n+1) - i$ or $(n+1) - j$, depending whether $i < j$ or $i > j$. In both cases the sum of degrees w_i, w_j in G'' can be written as $((n+1) - i) + ((n+1) - j) + |i - j|$. Hence, for each edge e with endpoints v_i, v_j we have additional cost of $|i - j|$. Summing up, we infer that the number of edges in G'' equals

$$\Delta_G \left(\sum_i^n ((n+1) - i) \right) + \sum_{v_i v_j \in E} |i - j| = \Delta_G \frac{n(n+1)}{2} + C(G, \pi).$$

The number of added edges equals the number of edges in G'' minus the number of edges in G' . So we add exactly

$$\left(\Delta_G \frac{n(n+1)}{2} + C(G, \pi) \right) - (\Delta_G n + 2|E|) = \Delta_G \frac{n(n-1)}{2} + C(G, \pi) - 2|E|$$

edges. □

Graph class name	Forbidden induced subgraphs
Chordal	C_n for $n \geq 4$
Proper Interval	claw, 2-net, 3-tent, C_n for $n \geq 4$
Interval	bipartite claw, umbrella, n -net for $n \geq 2$, n -tent for $n \geq 3$, C_n for $n \geq 4$
Threshold	$2K_2, C_4, P_4$
Trivially perfect	C_4, P_4

Table 1: Characterization of graph classes by forbidden induced subgraphs

Equivalence of the instances I and I' follows from the claim, and it proves the lemma. \square

Having an instance of CHAIN COMPLETION we transform it further to an instance of other completion problems by simply making A a clique, or by making both A and B cliques. By inspecting the forbidden subgraphs definition of each graph class we infer the equivalence of the instances, which is enough to prove Theorems 1.1 and 1.3.

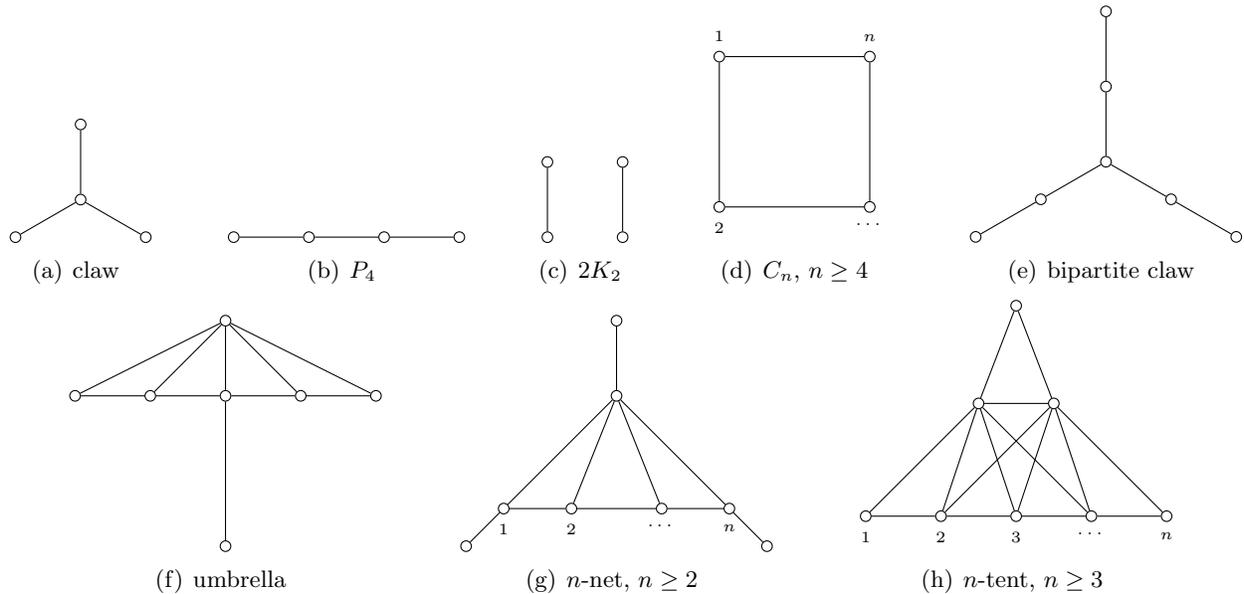


Figure 5: Forbidden induced subgraphs for various graph classes

Classes of chordal, interval, proper interval, threshold, trivially perfect graphs have many characterizations. For our purposes the most convenient one is by the set of forbidden induced subgraphs [12]. The characterization is presented in Table 1.

Lemma 5.4. *There are polynomial time reductions from CHAIN COMPLETION problem to MINIMUM FILL-IN, INTERVAL COMPLETION, PROPER INTERVAL COMPLETION, THRESHOLD COMPLETION, TRIVIAALLY PERFECT COMPLETION problems and these reductions do not change vertex set.*

Proof. For any bipartite graph $H' = (U_1, U_2, F)$ consider a graph $Ch(H') = (U_1, U_2, F \cup \{uv | u, v \in U_1\} \cup \{uv | u, v \in U_2\})$. In [44] it is shown that any bipartite graph H' is a chain graph if and only

if $Ch(H')$ is a chordal graph. It means that in order to reduce an instance of CHAIN COMPLETION $H = (A, B, F)$ to MINIMUM FILL-IN it is enough to construct cliques on sets of vertices A and B . So the constructed graph is a union of two cliques and some edges between cliques. Hence any arbitrarily completion of this graph does not contain a claw, bipartite claw, umbrella, p -net, q -tent for $p \geq 2, q \geq 3$ as these graphs have an independent set of size 3 and a union of two cliques does not. It follows that solutions for MINIMUM FILL-IN, PROPER INTERVAL COMPLETION, INTERVAL COMPLETION problems on such instances coincide and we can look at this reduction to MINIMUM FILL-IN as a reduction to PROPER INTERVAL COMPLETION or INTERVAL COMPLETION problems.

It is left to show a reduction from a CHAIN COMPLETION instance to THRESHOLD COMPLETION and TRIVIALY PERFECT COMPLETION instances. Having a CHAIN COMPLETION instance with a bipartite graph $H = (A, B, F)$ we consider TRIVIALY PERFECT COMPLETION and THRESHOLD COMPLETION problems on the graph $G = (A \cup B, F \cup \{(u, v) | u, v \in A\})$. We just add edges such that A becomes a clique. We show that a minimum chain completion of the graph H corresponds to a completion towards trivially perfect or threshold graph. Let F' be a solution of CHAIN COMPLETION for the graph H . Consider graph $G' = (A \cup B, F \cup \{(u, v) | u, v \in A\} \cup F')$, G' is a union of independent set and clique plus some edges between independent set and clique. So it does not contain induced $2K_2$ or C_4 . If G' contains induced $P_4 = v_1v_2v_3v_4$ then v_2, v_3 belong to clique and v_1, v_4 to independent set. However, this contradict to fact that edges between clique and independent set form a chain completion. As G' does not contain induced $2K_2, P_4, C_4$ it is trivially perfect and threshold graph. Let now F' denote solution of TRIVIALY PERFECT COMPLETION or THRESHOLD COMPLETION on instance G' . To finish the proof of correctness of reduction it is enough to show that $(A, B, (F \cup F') \cap E(A, B))$ is a chain graph. If this graph is not a chain graph then it must contain two independent edges v_1v_2, v_3v_4 [44]. However in such case graph on vertices v_1, v_2, v_3, v_4 will induce a P_4 in graph $G' \cup F'$ which contradict to the fact that $G' \cup F'$ is a threshold or trivially perfect graph.

To sum up all our reductions from CHAIN COMPLETION to MINIMUM FILL-IN, PROPER INTERVAL COMPLETION, INTERVAL COMPLETION, THRESHOLD COMPLETION, TRIVIALY PERFECT COMPLETION add some edges to a graph of CHAIN COMPLETION instance and do not change a vertex set of a graph. \square

At this point we almost proved Theorem 1.1 and Theorem 1.3.

Theorem 1.1. *Unless ETH fails, there is an integer $c \geq 1$ such that there are no $2^{\mathcal{O}(\sqrt{n}/\log^c n)}$, and consequently no $2^{\mathcal{O}(k^{1/4}/\log^c k)} \cdot n^{\mathcal{O}(1)}$ algorithms for the following problems: MINIMUM FILL-IN, INTERVAL COMPLETION, PROPER INTERVAL COMPLETION, TRIVIALY PERFECT COMPLETION, THRESHOLD COMPLETION, CHAIN COMPLETION.*

Proof. If the statement is not true then for some of the problems there is an algorithm running in time $2^{\mathcal{O}(\sqrt{n}/\log^c n)}$. Having instance of OPTIMUM LINEAR ARRANGEMENT on n vertices we can reduce it to problem under consideration with $(\Delta_G + 1)n = \mathcal{O}(n^2)$ vertices by Lemmas 5.2 and 5.4. This gives us an $2^{\mathcal{O}(\sqrt{n^2}/\log^c n^2)} = 2^{\mathcal{O}(n/\log^c n)}$ time algorithm which contradicts Theorem 1.5. As $k \leq n^2$ we also have $2^{\Omega(k^{1/4}/\log^c k)} \cdot n^{\mathcal{O}(1)}$ lower bound on the running time. \square

Theorem 1.3. *Unless Hypothesis 1.2 fails, there is no $2^{o(n+m)}$ -time algorithm for CHAIN COMPLETION, and no $2^{o(n)}$ -time algorithms for MINIMUM FILL-IN, INTERVAL COMPLETION, PROPER INTERVAL COMPLETION, TRIVIALY PERFECT COMPLETION, and THRESHOLD COMPLETION. Consequently, none of these problems can be solved in time $2^{o(\sqrt{k})} \cdot n^{\mathcal{O}(1)}$.*

Proof. In Section 4 we transformed a d -regular MIN BISECTION instance to OPTIMUM LINEAR ARRANGEMENT instance with bounded degree. Pipelined with Lemma 5.2 we get a reduction from d -regular MIN BISECTION to a CHAIN COMPLETION instance with $\mathcal{O}(n)$ vertices and edges. So $2^{o(n+m)}$ -time algorithm for CHAIN COMPLETION contradicts Hypothesis 1.2. By Lemma 5.4 we can reduce CHAIN COMPLETION to MINIMUM FILL-IN, INTERVAL COMPLETION, PROPER INTERVAL COMPLETION, TRIVIALY PERFECT COMPLETION, and THRESHOLD COMPLETION instance without changing the vertex set. Combining all three reductions in one we get reductions from MIN BISECTION to MINIMUM FILL-IN, PROPER INTERVAL COMPLETION, INTERVAL COMPLETION, THRESHOLD COMPLETION, TRIVIALY PERFECT COMPLETION problems which transform an instance with n vertices into an instance with $\mathcal{O}(n)$ vertices. This leads to $2^{\Omega(n)}$ lower bound for all discussed completion problems as well as to $2^{\Omega(\sqrt{k})} \cdot n^{\mathcal{O}(1)}$ lower bound because $k \leq n^2$. \square

6 Hardness of FEEDBACK ARC SET IN TOURNAMENTS

In this section we prove Theorem 1.4, that is, the lower bound on the complexity of FEEDBACK ARC SET IN TOURNAMENTS. We start with preparing an appropriately hard instance of FEEDBACK ARC SET in general digraphs, so that we can apply the reduction of Ailon et al. [1].

6.1 Preparing a hard instance of FAS

By $(E3_{1,1}, E2_{d,d})$ -SAT we denote the version of 3-SAT where every variable has

- exactly 1 positive occurrence in a clause of size 3,
- exactly 1 negative occurrence in a clause of size 3,
- exactly d positive occurrences in clauses of size 2,
- exactly d negative occurrences in clauses of size 2, and
- there are no clauses of size 1.

Similarly as before, GAP $(E3_{1,1}, E2_{d,d})$ -SAT $_{[\alpha, \beta]}$ for $0 \leq \alpha < \beta \leq 1$ is the problem of distinguishing whether the maximum number of clauses that can be satisfied in a given instance of $(E3_{1,1}, E2_{d,d})$ -SAT is at most αm , or at least βm , where m is the total number of clauses. We now give a hardness result for GAP $(E3_{1,1}, E2_{d,d})$ -SAT $_{[\alpha, \beta]}$.

Lemma 6.1. *There exists a positive integer d such that*

$$\text{GAP } E3\text{-NAE-SAT}_{[\alpha, 1]} \leq_P^{\text{lin}} \text{GAP } (E3_{1,1}, E2_{d,d})\text{-SAT}_{[\frac{1+\alpha+3d}{2+3d}, 1]}$$

Proof. Let ϕ be the input E3-CNF formula. For every variable x , let $n(x)$ be the number of occurrences of x in ϕ . Let G_x be a d -regular expander with $n(x)$ vertices, obtained using Theorem 2.2 for $p = 2$; suppose vertices of G_x are numbered $1, 2, \dots, n(x)$. Create ψ as follows:

- Replace x with $n(x)$ new variables $x_1, x_2, \dots, x_{n(x)}$, corresponding to the occurrences of x in ϕ ;
- For every $ij \in E(G_x)$ with $i \neq j$, introduce two clauses $(\neg x_i \vee x_j)$ and $(x_i \vee \neg x_j)$, which (if unbroken) force the evaluation of x_i to be equal to that of x_j ;

- For every $ii \in E(G_x)$, introduce a trivial clause $(\neg x_i \vee x_i)$;
- Perform the same construction for every other variable of ϕ ;
- For every 3-clause C of ϕ , introduce two new clauses C' and C'' to ψ . Clause C' is constructed from C by replacing each occurrence of some variable x with the new variable x_i corresponding to this occurrence; the polarity of the literals in C' is the same as in C . Clause C'' is constructed from C' by reversing the polarity of each literal. For example, if $C = (x \vee \neg y \vee z)$, then $C' = (x_i \vee \neg y_j \vee z_k)$ and $C'' = (\neg x_i \vee y_j \vee \neg z_k)$, where i, j, k are the indices of the occurrences of x, y, z in C , respectively.

It is easy to see that every variable of the new formula ψ appears in exactly two 3-clauses, once positively and once negatively. Also, it has exactly $2d$ occurrences in 2-clauses: d positive and d negative. Since the original formula had only clauses of size 3, the new formula has exactly $3m$ variables and $m(2 + 3d)$ clauses.

If there is a variable evaluation λ for ϕ that NAE-satisfies all clauses of ϕ , then we can construct a variable evaluation λ' for ψ by assigning all the variables x_i that originate in variables x the value $\lambda(x)$. Then it is easy to see that λ' satisfies all the clauses of ψ .

Suppose now that every variable evaluation for ϕ NAE-satisfies at most αm clauses, and for the sake of contradiction suppose that there is a variable evaluation λ' for ψ that satisfies more than $(1 + \alpha + 3d)m$ clauses. Let x be a variable of ϕ , and let us modify λ' as follows: assign all the variables $x_1, x_2, \dots, x_{n(x)}$ the value that is taken by the majority of these variables in the original evaluation λ' (breaking ties arbitrarily). Observe that since $h(G_x) \geq 2$, this step cannot decrease the number of satisfied clauses: if q is the number of variables out of $x_1, x_2, \dots, x_{n(x)}$ that take the minority value, then by replacing their values by the majority value we can unsatisfy at most $2q$ 3-clauses, but we satisfy at least $h(G_x) \cdot q \geq 2q$ 2-clauses that were previously unsatisfied. By performing this operation for every variable of ϕ , we can assume without loss of generality that in λ' all the variables originating in the same variable of ϕ are assigned the same value. This naturally defines a variable evaluation λ for ϕ . Then, provided λ' satisfied more than $(1 + \alpha + 3d)m$ clauses of ψ , we infer that λ NAE-satisfies more than αm clauses of ϕ . This is a contradiction. \square

We now turn our attention to the FEEDBACK VERTEX SET in general directed multigraphs: given a directed multigraph G , find the smallest possible subset of vertices X such that $G - X$ is acyclic. By FVS-BAL(d) we denote the variant of FVS where the input directed multigraph has no loops, and is $2d$ -regular and balanced, i.e., the indegree and the outdegree of every vertex is equal to d . Again, GAP FVS-BAL(d) $_{[\alpha, \beta]}$ is the gap problem where we need to distinguish between the cases when the optimum size of X is at most αn and at least βn , where n is the number of vertices of the multigraph.

Lemma 6.2.

$$\text{GAP (E3}_{1,1}, \text{E2}_{d,d})\text{-SAT}_{[\alpha, 1]} \leq_P^{lin} \text{GAP FVS-BAL}(d + 2)_{[1/2, (4-\alpha)/6]}$$

Proof. Let ϕ be the input instance of (E3_{1,1}, E2_{d,d})-SAT, and let n be the number of variables of ϕ . Then m , the number of clauses of ϕ , is equal to $(2/3 + d)n$.

Construct a directed multigraph G as follows. For every variable x of ϕ , create two vertices u_{\top}^x and u_{\perp}^x , corresponding to setting x to true and false, respectively. Add edges $(u_{\top}^x, u_{\perp}^x)$ and $(u_{\perp}^x, u_{\top}^x)$ to the edge set, for every variable x of ϕ . Moreover, for every 2-clause of ϕ add a 2-cycle between vertices corresponding to its literals (e.g. clause $x \vee \neg y$ gives rise to edges $(u_{\top}^x, u_{\perp}^y)$ and

$(u_{\perp}^y, u_{\top}^x)$), and similarly for every 3-clause of ϕ add a 3-cycle between vertices corresponding to its literals, oriented arbitrarily. Note that in this manner trivial clauses of the form $(\neg x \vee x)$ give rise to additional copies of the 2-cycle $(u_{\top}^x, u_{\perp}^x)(u_{\top}^x, u_{\perp}^x)$. This concludes the construction of G . It is easy to verify using the assumed properties of ϕ that every vertex of G has indegree and outdegree equal to $d + 2$, and moreover in the construction we did not introduce loops. Let $n' = 2n$ be the number of vertices in G .

Suppose first that there exists a variable evaluation λ for ϕ that satisfies all the clauses of ϕ . Define X to be the set of all the vertices $u_{\lambda(x)}^x$ for x being a variable of ϕ ; note that $|X| = n = n'/2$. Since λ satisfies all the clauses, and every variable of ϕ participates in exactly one 3-clause positively and in exactly one 3-clause negatively, then it is easy to see that all the weakly connected components of $G - X$ are either isolated vertices or single edges. Thus, $G - X$ is acyclic.

Assume now that every variable evaluation for ϕ satisfies at most αm clauses, and for the sake of contradiction suppose that there exists a set X with $|X| < \frac{4-\alpha}{6}n' = \frac{4-\alpha}{3}n$ such that $G - X$ is acyclic. Observe that from each pair $\{u_{\top}^x, u_{\perp}^x\}$ at least one vertex has to belong to X . Define a variable evaluation λ for ϕ as follows: if $|\{u_{\top}^x, u_{\perp}^x\} \cap X| = 1$, then $\lambda(x)$ is such that $u_{\lambda(x)}^x \in X$, and otherwise $\lambda(x)$ is chosen arbitrarily. Observe that the first alternative holds for a set of more than $n - \frac{1-\alpha}{3}n$ variables; let us denote them by S . Since $G - X$ is acyclic, each of 2- and 3-cycles constructed for a clause C of ϕ has at least one vertex from X . If all the variables of C belong to S , then it can be easily seen that this implies that λ satisfies C . Hence, the only clauses of ϕ that can be unsatisfied by λ are the ones that contain at least one variable outside S . Every variable of ϕ occurs in at most $2d + 2$ clauses, and there are less than $\frac{1-\alpha}{3} \cdot n$ variables outside S , which means that λ unsatisfies less than $\frac{(1-\alpha)(2d+2)}{3} \cdot n$ clauses. Hence the fraction of unsatisfied clauses is less than

$$\frac{(1-\alpha)(2d+2)n}{3m} = \frac{(1-\alpha)(2d+2)}{(2+3d)} \leq 1-\alpha.$$

This is a contradiction. □

Finally, there is a well-known reduction that reduces FEEDBACK VERTEX SET to FEEDBACK ARC SET in the directed setting. This reduction appears to preserve the gap. In the following, by $\text{GAP FAS}(d)_{[\alpha, \beta]}$ we denote the problem of determining, for a given directed multigraph G without loops whose underlying undirected multigraph is d -regular, whether the minimum number of edges that needs to be removed from G to make it acyclic is at most αm or at least βm , where $0 \leq \alpha < \beta \leq 1$ and m is the number of edges in G .

Lemma 6.3.

$$\text{GAP FVS-BAL}(d)_{[\alpha, \beta]} \stackrel{\text{lin}}{\leq}_P \text{GAP FAS}(d+1)_{[\frac{\alpha}{d+1}, \frac{\beta}{d+1}]}$$

Proof. Let G be the input directed multigraph, and let n and m denote the numbers of edges and vertices of G , respectively; by the assumption that G is $2d$ -regular we know that $m = dn$. Construct a graph G' as follows:

- For every $u \in V(G)$ create two vertices $u^-, u^+ \in V(G')$ and an edge $(u^-, u^+) \in E(G')$;
- For every edge $(u, v) \in E(G)$, create an edge (u^+, v^-) .

This concludes the construction of G' . Let E_1, E_2 be the sets of edges constructed in the first and second bullet point, respectively. Since G was $2d$ -regular and balanced, we infer that every vertex

u^+ has outdegree d and indegree 1, whereas every vertex u^- has outdegree 1 and indegree d . Thus, G' has $n' = 2n$ vertices and $m' = m + n = (d + 1)n$ edges.

Suppose first that X is a subset of vertices of G with size at most αn such that $G - X$ is acyclic. Let $F = \{(u^+, u^-) \mid u \in X\} \subseteq E_1$. Then it can be easily seen that $G' - F$ is acyclic, and $|F| = |X| \leq \alpha n = \frac{\alpha}{d+1} m'$.

Assume now that every subset $X \subseteq V(G)$ for which $G - X$ is acyclic has size at least βn , and for the sake of contradiction suppose that there is a set $F \subseteq E(G')$ such that $G' - F$ is acyclic and $|F| < \frac{\beta}{d+1} m'$. Observe that if F contains some edge $(u^+, v^-) \in E_2$, then we could modify F by removing (u^+, v^-) from F and adding (v^-, v^+) to F (unless it is not already contained in F , in which case we do not add any edge to F). This operation can only decrease the number of edges in F and preserves the property that $G' - F$ is acyclic; this is because after removing (v^-, v^+) , v^- becomes a sink. Thus, without loss of generality we can assume that $F \subseteq E_1$. Let X be the set of vertices $u \in V(G)$ for which $(u^-, u^+) \in F$. Since $G' - F$ is acyclic, it easily follows that $G - X$ is also acyclic. Moreover, $|X| = |F| < \frac{\beta}{d+1} m' = \beta n$. This is a contradiction. \square

Finally, observe that in an instance of FAS without loops one can subdivide every edge once, which doubles the number of edges while not changing the size of the optimum solution. Thus, application of this reduction to the gap problem shrinks the gap twice and makes the directed graph at hand simple: it has no loops, no parallel edges, and moreover if (u, v) is an edge then (v, u) is not. By combining this observation with Theorem 3.1 and Lemmas 3.3, 3.4, 6.1, 6.2, and 6.3, we obtain the following result.

Theorem 6.4. *Unless ETH fails, there exist $0 \leq \alpha < \beta \leq 1$, $c \geq 1$, and $d > 0$ such that there is no $2^{\mathcal{O}(\frac{n}{\log^c(n)})}$ algorithm for GAP FAS $_{[\alpha, \beta]}$ on directed simple graphs of maximum total degree d .*

6.2 Reducing FAS to FAST

Theorem 1.4. *Unless ETH fails, there is an integer $c \geq 1$ such that there is no $2^{\mathcal{O}(\sqrt{n}/\log^c n)}$, and consequently no $2^{\mathcal{O}(k^{1/4}/\log^c k)} \cdot n^{\mathcal{O}(1)}$ algorithm for FEEDBACK ARC SET IN TOURNAMENTS.*

Proof. We provide a randomized reduction that essentially reiterates the argument of Ailon et al. [1]. In the analysis, we use the known fact that for any directed graph H , $\text{fas}(H)$ is equal to the minimum possible number of edges oriented backwards (called *feedback edges*) in an ordering of vertices of $V(H)$, where $\text{fas}(H)$ is the optimum size of a feedback arc set in H . Thus, we may equivalently think of the FAS problem as finding an ordering π of $V(H)$ that minimizes the number of feedback edges. For an ordering π of $V(H)$, by $\text{fas}(H, \pi)$ we denote the number of feedback edges in the ordering π .

Let α, β, c, d be the constants given by Theorem 6.4, and let G be an instance of GAP FAS $_{[\alpha, \beta]}$, where G is a simple directed graph with n vertices and m edges and has maximum total degree d . Note that w.l.o.g. we may assume $m \geq n$, as otherwise there is a vertex in G with no outgoing edges, which can be safely removed. Let us fix an integer k , to be determined later. We consider the k -blow up G_k defined as follows: for every $u \in V(G)$ we create a k vertices u_1, u_2, \dots, u_k in G_k , and for all $1 \leq i, j \leq k$ we put $(u_i, v_j) \in E(G_k)$ if and only if $(u, v) \in E(G)$. Thus, vertices u_i are twins. Ailon et al. [1], based on a communication by Alon, argue that there is an optimum ordering for G_k which may be obtained by taking an optimum ordering for G and replacing every vertex $u \in V(G)$ by a block consisting of vertices $\{u_i\}_{1 \leq i \leq k}$ in any order; hence in particular $\text{fas}(G_k) = k^2 \cdot \text{fas}(G)$. For an ordering σ of $V(G)$, let σ_k be an ordering of $V(G_k)$ constructed in the manner described above.

Construct a tournament T_k from G_k by adding edges between every pair of vertices that are not connected by an edge in G_k , where the orientations of these edges are chosen independently and uniformly at random. Observe that $|E(G_k)| = k^2 \cdot |E(G)| \leq dk^2n/2$. Let $R_k = (V(T_k), E(T_k) \setminus E(G_k))$ be the directed graph consisting only of the edges picked at random. Then for a sufficiently large n we have that $|E(R_k)| = \binom{nk}{2} - |E(G_k)| \geq \frac{(nk)^2}{4}$, and of course $|E(R_k)| \leq \frac{(nk)^2}{2}$.

We now prove that with high probability, FEEDBACK ARC SET(T_k) is closely related to FEEDBACK ARC SET(G_k), because the number of feedback edges that need to be chosen from the edges picked at random is concentrated around the expected value.

Let us fix some ordering π of $V(G_k)$, then $\text{fas}(T_k, \pi) = \text{fas}(G_k, \pi) + \text{fas}(R_k, \pi)$. For $e \in E(R_k)$, let X_e be the indicator random variable having value 1 if e is a feedback edge w.r.t. π , and 0 otherwise. Let also $X = \sum_{e \in E(R_k)} X_e$; then $\mathbb{E}X = \frac{|E(R_k)|}{2}$. Let $\eta = \frac{\beta - \alpha}{3}$. Since X_e -s are independent, from the Chernoff bound we obtain that

$$\Pr(|X - \mathbb{E}X| \geq \eta k^2 n) \leq 2 \exp\left(-\frac{2\eta^2 k^4 n^2}{|E(R_k)|}\right) \leq 2 \exp(-4k^2 \eta^2).$$

Suppose now that there exists an ordering π of $V(G)$ that has at most αm feedback edges. Then, with probability at least $1 - 2 \exp(-4\eta^2 k^2)$ we have that

$$\text{fas}(T_k, \pi_k) \leq \alpha k^2 m + |E(R_k)|/2 + \eta k^2 n \leq \frac{2\alpha + \beta}{3} \cdot k^2 m + |E(R_k)|/2. \quad (5)$$

Hence, if we can set k to be a large enough constant, such that conclusion (5) holds with probability at least $3/4$.

Suppose now that $\text{fas}(G) \geq \beta m$. Then, for a fixed ordering σ of $V(G_k)$ we have that with probability at least $1 - 2 \exp(-4\eta^2 k^2)$ it holds that

$$\begin{aligned} \text{fas}(T_k, \sigma) &\geq \text{fas}(G_k, \sigma) + |E(R_k)|/2 - \eta k^2 n \geq k^2 \text{fas}(G) + |E(R_k)|/2 - \eta k^2 n \\ &\geq k^2 \beta m + |E(R_k)|/2 - \eta k^2 n \geq \frac{\alpha + 2\beta}{3} \cdot k^2 m + |E(R_k)|/2. \end{aligned} \quad (6)$$

We would like to infer that with high probability this conclusion holds for all the possible orderings σ , and for this we will use the union bound. Observe that the number of orderings σ of $V(G_k)$ is $(nk)! = \exp(\mathcal{O}(nk \log(nk)))$, while the probability of failure for each of them is at most $2 \exp(-4\eta^2 k^2)$. Since η is a positive constant, simple computations show that by setting $k = \Theta(n \log n)$, we have that $(nk)! \cdot 2 \exp(-4\eta^2 k^2) \leq 1/4$, and hence conclusion (6) holds simultaneously for all orderings σ with probability at least $3/4$.

Suppose now that FEEDBACK ARC SET IN TOURNAMENTS admitted an algorithm with running time $2^{\mathcal{O}(\frac{n^{1/2}}{\log^{c'} n})}$ for $c' = c + \frac{1}{2}$, where c is as in Theorem 6.4. Apply this algorithm to the constructed tournament T_k to compute $\text{fas}(T_k)$. In case $\text{fas}(G) \leq \alpha m$, then with probability at least $3/4$ we have that $\text{fas}(T_k) \leq \frac{2\alpha + \beta}{3} \cdot k^2 m + |E(R_k)|/2$. In case $\text{fas}(G) \geq \beta m$, then with probability at least $3/4$ we have that $\text{fas}(T_k) \geq \frac{\alpha + 2\beta}{3} \cdot k^2 m + |E(R_k)|/2$. Since $\frac{2\alpha + \beta}{3} < \frac{\alpha + 2\beta}{3}$, these two alternatives are disjoint and the algorithm can, with double-sided error, resolve the input instance of GAP FAS $_{[\alpha, \beta]}$. Since $|V(T_k)| = \Theta(n^2 \log n)$, this procedure runs in time $2^{\mathcal{O}(\frac{n}{\log^c(n)})}$. This is a contradiction with Theorem 6.4. \square

7 Conclusions

In this work we have given evidence that $2^{\mathcal{O}(\sqrt{k}\text{-polylog}(k))} \cdot n^{\mathcal{O}(1)}$ can be the final answer for the running times of parameterized algorithms for MINIMUM FILL-IN, INTERVAL COMPLETION, PROPER INTERVAL COMPLETION, TRIVIALY PERFECT COMPLETION, THRESHOLD COMPLETION, and CHAIN COMPLETION. This evidence is based on a new complexity hypothesis connected to the hardness of approximation for the MIN BISECTION problem. Thus, the answer given by us is not completely satisfactory: the lower bounds that we can give only under the assumption of ETH are much weaker. Rather, our results uncover a surprising link between the parameterized algorithms for MINIMUM FILL-IN and related problems, and the approximability of MIN BISECTION. Thus it seems that the question about the optimality of the former has a much deeper, fundamental nature.

Therefore, we believe that our work strongly motivates further investigation of Hypothesis 1.2. Can this conjecture be linked to ETH and possibly some other strong conjectures like SETH, the existence of linear PCPs, or the conjectures proposed by Feige [22]? Or maybe it can be simply disproved?

Our improved lower bound for FEEDBACK ARC SET IN TOURNAMENTS still has a gap between $k^{1/4}$ and $k^{1/2}$. A closer inspection of the proof uncovers a fundamental obstacle for why we cannot achieve tightness: the Chernoff concentration bound used in the proof of Theorem 1.4 is essentially tight, because in every tournament on n vertices there is a feedback arc set of size $\binom{n}{2}/2 - \Omega(n^{3/2})$ [29]. If this error term was of magnitude $\Theta(n)$ instead of $\Theta(n^{3/2})$, then our approach would give a tight result for FAST. Can this problem be circumvented, or maybe the high anticoncentration of the number of feedback edges in a random ordering of a tournament can be exploited algorithmically to obtain a faster algorithm?

Finally, even assuming Hypothesis 1.2 we do not get tight bounds, due to the (poly)logarithmic factors in the exponent describing the running time of the existing algorithms for completion problems. Bridging this gap can be seen as another forthcoming goal.

Acknowledgements

We would like to thank Per Austrin for valuable discussions about different versions of the PCP theorem and anonymous reviewers for their helpful comments.

References

- [1] N. Ailon, M. Charikar, and A. Newman. Aggregating inconsistent information: Ranking and clustering. *J. ACM*, 55(5), 2008.
- [2] N. Alon. Ranking tournaments. *SIAM J. Discrete Math.*, 20(1):137–142, 2006.
- [3] N. Alon, D. Lokshtanov, and S. Saurabh. Fast FAST. In S. Albers, A. Marchetti-Spaccamela, Y. Matias, S. E. Nikolettseas, and W. Thomas, editors, *Automata, Languages and Programming, 36th International Colloquium, ICALP 2009, Rhodes, Greece, July 5-12, 2009, Proceedings, Part I*, volume 5555 of *Lecture Notes in Computer Science*, pages 49–58. Springer, 2009.
- [4] S. Arora and B. Barak. *Computational Complexity - A Modern Approach*. Cambridge University Press, 2009.

- [5] P. Berman and M. Karpinski. Approximation hardness of bounded degree MIN-CSP and MIN-BISECTION. In P. Widmayer, F. T. Ruiz, R. M. Bueno, M. Hennessy, S. Eidenbenz, and R. Conejo, editors, *Automata, Languages and Programming, 29th International Colloquium, ICALP 2002, Malaga, Spain, July 8-13, 2002, Proceedings*, volume 2380 of *Lecture Notes in Computer Science*, pages 623–632. Springer, 2002.
- [6] I. Bliznets, F. V. Fomin, M. Pilipczuk, and M. Pilipczuk. A subexponential parameterized algorithm for Interval Completion. *CoRR*, abs/1402.3473, 2014. Accepted to SODA 2016 conference.
- [7] I. Bliznets, F. V. Fomin, M. Pilipczuk, and M. Pilipczuk. A subexponential parameterized algorithm for Proper Interval Completion. In A. S. Schulz and D. Wagner, editors, *Algorithms - ESA 2014 - 22th Annual European Symposium, Wroclaw, Poland, September 8-10, 2014. Proceedings*, volume 8737 of *Lecture Notes in Computer Science*, pages 173–184. Springer, 2014.
- [8] E. Bonnet, B. Escoffier, E. J. Kim, and V. T. Paschos. On subexponential and FPT-time inapproximability. *Algorithmica*, 71(3):541–565, 2015.
- [9] E. Bonnet, M. Lampis, and V. T. Paschos. Time-approximation trade-offs for inapproximable problems. *CoRR*, abs/1502.05828, 2015.
- [10] E. Bonnet and V. T. Paschos. Parameterized (in)approximability of subset problems. *Oper. Res. Lett.*, 42(3):222–225, 2014.
- [11] E. Bonnet and V. T. Paschos. Sparsification and subexponential approximation. *CoRR*, abs/1402.2843, 2014.
- [12] A. Brandstädt, J. P. Spinrad, et al. *Graph classes: a survey*, volume 3. Siam, 1999.
- [13] P. Charbit, S. Thomassé, and A. Yeo. The Minimum Feedback Arc Set problem is NP-hard for tournaments. *Combinatorics, Probability & Computing*, 16(1):1–4, 2007.
- [14] M. Cygan, F. V. Fomin, L. Kowalik, D. Loksthanov, D. Marx, M. Pilipczuk, M. Pilipczuk, and S. Saurabh. *Parameterized Algorithms*. Springer, 2015.
- [15] T. A. Davis. *7. Fill-Reducing Orderings*. SIAM.
- [16] E. D. Demaine and M. Hajiaghayi. The bidimensionality theory and its algorithmic applications. *Comput. J.*, 51(3):292–302, 2008.
- [17] I. Dinur. The PCP theorem by gap amplification. *J. ACM*, 54(3):12, 2007.
- [18] P. G. Drange, M. S. Dregi, D. Lokshtanov, and B. D. Sullivan. On the threshold of intractability. In N. Bansal and I. Finocchi, editors, *Algorithms ESA 2015*, volume 9294 of *Lecture Notes in Computer Science*, pages 411–423. Springer Berlin Heidelberg, 2015.
- [19] P. G. Drange, F. V. Fomin, M. Pilipczuk, and Y. Villanger. Exploring subexponential parameterized complexity of completion problems. In E. W. Mayr and N. Portier, editors, *31st International Symposium on Theoretical Aspects of Computer Science (STACS 2014)*, volume 25 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 288–299. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2014.

- [20] P. G. Drange, F. V. Fomin, M. Pilipczuk, and Y. Villanger. Exploring subexponential parameterized complexity of completion problems. In E. W. Mayr and N. Portier, editors, *31st International Symposium on Theoretical Aspects of Computer Science (STACS 2014), STACS 2014, March 5-8, 2014, Lyon, France*, volume 25 of *LIPICs*, pages 288–299. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2014. Journal version accepted to ACM Transactions on Computational Theory (TOCT).
- [21] P. G. Drange and M. Pilipczuk. A polynomial kernel for trivially perfect editing. In N. Bansal and I. Finocchi, editors, *Algorithms ESA 2015*, volume 9294 of *Lecture Notes in Computer Science*, pages 424–436. Springer Berlin Heidelberg, 2015.
- [22] U. Feige. Relations between average case complexity and approximation complexity. In J. H. Reif, editor, *Proceedings on 34th Annual ACM Symposium on Theory of Computing, May 19-21, 2002, Montréal, Québec, Canada*, pages 534–543. ACM, 2002.
- [23] U. Feige. Faster FAST (Feedback Arc Set in Tournaments). *CoRR*, abs/0911.5094, 2009.
- [24] J. Flum and M. Grohe. *Parameterized Complexity Theory*. Texts in Theoretical Computer Science. Springer, 2006.
- [25] F. V. Fomin and Y. Villanger. Subexponential parameterized algorithm for Minimum Fill-in. *SIAM J. Comput.*, 42(6):2197–2216, 2013.
- [26] M. R. Garey and D. S. Johnson. *Computers and intractability*, volume 174. Freeman New York, 1979.
- [27] M. R. Garey, D. S. Johnson, and L. J. Stockmeyer. Some simplified NP-complete graph problems. *Theor. Comput. Sci.*, 1(3):237–267, 1976.
- [28] M. Hajiaghayi, R. Khandekar, and G. Kortsarz. The foundations of fixed parameter inapproximability. *CoRR*, abs/1310.2711, 2013.
- [29] R. Hassin and S. Rubinfeld. Approximations for the maximum acyclic subgraph problem. *Information Processing Letters*, 51:133–140, 1994.
- [30] R. Impagliazzo and R. Paturi. On the complexity of k -SAT. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001.
- [31] R. Impagliazzo, R. Paturi, and F. Zane. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.*, 63(4):512–530, 2001.
- [32] H. Kaplan, R. Shamir, and R. E. Tarjan. Tractability of parameterized completion problems on chordal, strongly chordal, and proper interval graphs. *SIAM J. Comput.*, 28(5):1906–1922, 1999.
- [33] M. Karpinski and W. Schudy. Faster algorithms for Feedback Arc Set Tournament, Kemeny Rank Aggregation and Betweenness Tournament. In O. Cheong, K. Chwa, and K. Park, editors, *Algorithms and Computation - 21st International Symposium, ISAAC 2010, Jeju Island, Korea, December 15-17, 2010, Proceedings, Part I*, volume 6506 of *Lecture Notes in Computer Science*, pages 3–14. Springer, 2010.

- [34] S. Khot and I. Shinkar. On hardness of approximating the parameterized Clique problem. *Electronic Colloquium on Computational Complexity (ECCC)*, 22:13, 2015.
- [35] D. Lokshantov, D. Marx, and S. Saurabh. Lower bounds based on the exponential time hypothesis. *Bulletin of the EATCS*, 105:41–72, 2011.
- [36] A. Lubotzky, R. Phillips, and P. Sarnak. Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988.
- [37] D. Marx. On the optimality of planar and geometric approximation schemes. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2007), October 20-23, 2007, Providence, RI, USA, Proceedings*, pages 338–348. IEEE Computer Society, 2007.
- [38] D. Marx. What’s next? future directions in parameterized complexity. In H. L. Bodlaender, R. Downey, F. V. Fomin, and D. Marx, editors, *The Multivariate Algorithmic Revolution and Beyond - Essays Dedicated to Michael R. Fellows on the Occasion of His 60th Birthday*, volume 7370 of *Lecture Notes in Computer Science*, pages 469–496. Springer, 2012.
- [39] M. Morgenstern. Existence and explicit constructions of $q + 1$ regular ramanujan graphs for every prime power q . *Journal of Combinatorial Theory, Series B*, 62(1):44 – 62, 1994.
- [40] H. Räcke. Optimal hierarchical decompositions for congestion minimization in networks. In C. Dwork, editor, *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*, pages 255–264. ACM, 2008.
- [41] J. D. Rose. *A graph-theoretic study of the numerical solution of sparse positive definite systems of linear equations*. Graph Theory and Computing. Academic Press, New York, 1972.
- [42] T. J. Schaefer. The complexity of satisfiability problems. In *Proceedings of the tenth annual ACM symposium on Theory of computing*, pages 216–226. ACM, 1978.
- [43] Y. Villanger, P. Heggernes, C. Paul, and J. A. Telle. Interval Completion is fixed parameter tractable. *SIAM J. Comput.*, 38(5):2007–2020, 2009.
- [44] M. Yannakakis. Computing the minimum fill-in is NP-complete. *SIAM Journal on Algebraic Discrete Methods*, 2(1):77–79, 1981.

A Problem definitions

OPTIMUM LINEAR ARRANGEMENT (OLA)

Input: A graph $G = (V, E)$, an integer k .

Question: Does there exist a linear arrangement π of G of cost at most k ?

OPTIMUM LINEAR ARRANGEMENT $_{\leq}(d)$ (OLA $_{\leq}(d)$)

Input: A graph $G = (V, E)$ with degree at most d , an integer k .

Question: Does there exist a linear arrangement π of G of cost at most k ?

MAX CUT

Input: A graph $G = (V, E)$, an integer k

Question: Does there exist a cut of size at least k ?

GAP MAX CUT $_{[\alpha, \beta]}$

Input: A graph $G = (V, E)$.

Case 1: G admits a cut of size at least βm .

Case 2: G does not admit a cut of size larger than αm .

MIN BISECTION

Input: A graph $G = (V, E)$ with even number of vertices, an integer k .

Question: Does there exist a cut (A, B) of size at most k , such that $|A| = |B|$?

GAP MIN BISECTION(d) $_{[\alpha, \beta]}$

Input: A d -regular graph $G = (V, E)$ with even number of vertices.

Case 1: G admits a cut (A, B) of size at least βm , such that $|A| = |B|$.

Case 2: G does not admit a cut (A, B) of size larger than αm , such that $|A| = |B|$.

A.1 Satisfiability problems

We consider several variants of the satisfiability problem, in general defined as follows.

X-SAT

Input: An X-SAT formula $\phi = C_1 \wedge \dots \wedge C_m$.

Question: Does there exist an assignment of the variables of ϕ , such that ϕ is satisfiable?

X-SAT $_{[\alpha, \beta]}$

Input: An X-SAT formula $\phi = C_1 \wedge \dots \wedge C_m$.

Case 1: ϕ admits an assignment satisfying at least βm clauses.

Case 2: ϕ does not admit an assignment assignment satisfying more than αm clauses.

Where an X-SAT formula is a formula from the SAT related problem, precisely it is an El -CNF formula for El -SAT, El -NAE-SAT, and an l -AND formula for l -AND-SAT.

We also similarly define a problem X-SAT (d) with the difference that the variables of an input formula occur in at most d clauses, e.g. l -AND-SAT (d).

A.2 Completion problems

The following problems is a generic version of a completion problem to a given graph class X .

X-COMPLETION

Input: An undirected graph G , an integer k .

Question: Is it possible to add at most k edges to G , so that the obtained graph belongs to the graph class X ?

We consider the following list of completion problems: MINIMUM FILL-IN, CHAIN COMPLETION, PROPER INTERVAL COMPLETION, INTERVAL COMPLETION, THRESHOLD COMPLETION, TRIVIAL PERFECT COMPLETION, where MINIMUM FILL-IN is completion to chordal graphs, and other problems have self-descriptive names.