SIAM J. COMPUT.
Vol. 40, No. 2, pp. 567–596

INAPPROXIMABILITY RESULTS FOR MAXIMUM EDGE BICLIQUE, MINIMUM LINEAR ARRANGEMENT, AND SPARSEST CUT*

CHRISTOPH AMBÜHL†, MONALDO MASTROLILLI‡, AND OLA SVENSSON§

Abstract. We consider the Minimum Linear Arrangement problem and the (Uniform) Sparsest Cut problem. So far, these two notorious NP-hard graph problems have resisted all attempts to prove inapproximability results. We show that they have no polynomial time approximation scheme, unless NP-complete problems can be solved in randomized subexponential time. Furthermore, we show that the same techniques can be used for the Maximum Edge Biclique problem, for which we obtain a hardness factor similar to previous results but under a more standard assumption.

Key words. hardness of approximation, graph theory

AMS subject classifications. 68Q17, 68Q25

DOI. 10.1137/080729256

1. Introduction.

Maximum Edge Biclique, Sparsest Cut, and Minimum Linear Arrangement are fundamental combinatorial problems. They have a rich number of applications in areas such as computational biology, circuit design, manufacturing optimization, and graph drawing (see, e.g., [11, 12, 15, 30]). Moreover, as they often appear as subproblems in more complex settings, it is important to understand whether we can efficiently find "good" solutions to these problems. For example, suppose we have a "good" algorithm for the Sparsest Cut problem. Then we can partition a graph into large pieces while minimizing the size of the "interface" between them, a property that is very useful when designing graph theoretic algorithms via the divide-and-conquer paradigm (see [30] for a comprehensive discussion).

Since the addressed optimization problems are NP-hard [19, 25, 26], one is forced to settle for approximation algorithms. Unfortunately, there is no known approximation algorithm for the Maximum Edge Biclique problem that achieves an approximation guarantee significantly better than the inverse of the number of edges in the bipartite graph. The situation for the Sparsest Cut problem and the Minimum Linear Arrangement problem is more hopeful. Leighton and Rao [24] showed that the Sparsest Cut problem can be approximated within a factor $O(\log n)$ by using a linear programming relaxation. The approximation guarantee is tight in the sense that it

*Received by the editors July 3, 2008; accepted for publication (in revised form) February 6, 2011; published electronically April 27, 2011. An extended abstract of this paper appeared in *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science* (FOCS), 2007, pp. 329–337.

http://www.siam.org/journals/sicomp/40-2/72925.html

†Department of Computer Science, University of Liverpool, Liverpool, L69 3BX, Great Britain (christoph@csc.liv.ac.uk). This author's work was supported by Nuffield Foundation grant NAL32608.

‡IDSIA, 6928 Lugano, Switzerland (monaldo@idsia.ch). This author's work was supported by Swiss National Science Foundation project 200021-104017/1, "Power Aware Computing," and by Swiss National Science Foundation project 200020-122110/1, "Approximation Algorithms for Machine Scheduling Through Theory and Experiments III."

§KTH - Royal Institute of Technology, SE-100 44 Stockholm, Sweden (osven@kth.se). This author's work was supported by Swiss National Science Foundation project 200021-104017/1, "Power Aware Computing," by Swiss National Science Foundation project 200020-122110/1, "Approximation Algorithms for Machine Scheduling Through Theory and Experiments III," and by ERC Advanced Investigator grant 226203.

matches the lower bound on the integrality gap of the corresponding relaxation up to constant factors [24]. Recently, Arora, Rao, and Vazirani [5] used semidefinite programming to obtain the best known approximation algorithm for Uniform Sparsest Cut with performance guarantee $O(\sqrt{\log n})$.[1] Subsequently, these techniques were also used to obtain the best known approximation algorithm for the non-Uniform Sparsest Cut problem [9], which is a generalization of the Uniform Sparsest Cut problem.

The situation is similar for Minimum Linear Arrangement. Feige and Lee [18] and Charikar et al. [8] independently showed that combining the techniques in [5] with the rounding algorithm of Rao and Richa [29] yields an $O(\sqrt{\log n}\log\log n)$-approximation algorithm for the Minimum Linear Arrangement problem. This improves over the $O(\log n)$-approximation algorithm of Rao and Richa [29]. The semidefinite programming relaxations used for Sparsest Cut and Minimum Linear Arrangement were recently shown to have integrality gap $\Omega(\log\log n)$ by Devanur et al. [13]. This result suggests that we cannot use those relaxations to obtain a constant factor approximation algorithm for the Sparsest Cut problem or the Minimum Linear Arrangement problem.

Despite substantial efforts, it seems difficult to obtain "good" (constant factor) approximation algorithms for the considered problems. Instead, one can hope for negative results, i.e., results that indeed show the problems to be hard to approximate. For Sparsest Cut and Minimum Linear Arrangement, the only previously known hardness results are based on the unique games conjecture [21] and say that the *non-Uniform* Sparsest Cut problem has no constant factor approximation algorithm [10, 23]. Feige and Kogan [17] showed that the Maximum Edge Biclique problem is hard to approximate within a factor of $2^{(\log n)^\delta}$ for some $\delta > 0$ under the plausible assumption that 3-SAT $\notin DTIME(2^{n^{3/4}})$. The hardness factor was later improved by Feige [16], who showed that Maximum Edge Biclique is hard to approximate within $O(n^\epsilon)$, for some $\epsilon > 0$, by assuming a hypothesis about average-case hardness of Random 3-SAT. (The formal definition of the used hypothesis is as follows. For every fixed $\epsilon > 0$, for Δ a sufficiently large constant independent of n, there is no polynomial time algorithm that on most 3CNF formulas with n variables and $m = \Delta n$ clauses outputs "typical" but never outputs "typical" on 3CNF formulas with $(1-\epsilon)m$ satisfiable clauses. The word "typical" comes from the fact that for a large enough Δ, every assignment to the variables of a random 3CNF formula with n variables and $m = \Delta n$ clauses satisfies roughly $7m/8$ clauses.)

In summary, no "good" approximation algorithms are known for Maximum Edge Biclique, Sparsest Cut, and Minimum Linear Arrangement. At the same time, the only known hardness of approximation results use nonstandard assumptions and apply to non-Uniform Sparsest Cut (a more general and thus possibly harder problem than Uniform Sparsest Cut) and Maximum Edge Biclique. Improving our understanding of the approximability of these problems is considered a major open problem in complexity theory (see, e.g., [13, 31, 32]).

Here, we address this problem by giving the first inapproximability results for Sparsest Cut and Minimum Linear Arrangement. We also obtain hardness of approximation results for Maximum Edge Biclique that are comparable to Feige's results [16] but use a more standard assumption. Our results use the recent Quasi-random PCP construction of Khot [22], who proved important inapproximability results for Graph

[1] The same approximation guarantee was later obtained without solving the semidefinite program, and this approach has better running time [2].

Min-Bisection, Densest Subgraph, and Balanced Bipartite Clique. These inapproximability results were obtained under the standard assumption that SAT has no probabilistic algorithm that runs in time 2^{n^ϵ}, where n is the instance size and $\epsilon > 0$ can be made arbitrarily close to 0. Prior to Khot's results, Graph Min-Bisection, Densest Subgraph, and Balanced Bipartite Clique had a status similar to Maximum Edge Biclique, i.e., no "good" approximation guarantees, and the only hardness results were obtained by using nonstandard assumptions [16]. However, the results in [22] and the hypothesis used by Feige in [16] were not known to imply inapproximability results for Sparsest Cut and Minimum Linear Arrangement (see, e.g., [13, 31]). The main contribution of this paper is to show that the Quasi-random PCP [22] and carefully designed reductions indeed suffice to rule out the existence of a polynomial time approximation scheme (PTAS) for Sparsest Cut, Minimum Linear Arrangement, and Maximum Edge Biclique. The hardness factor of Maximum Edge Biclique can then be boosted by using standard techniques (see Theorem 1.4).

1.1. Preliminaries. We start with the definitions of the addressed problems followed by a brief explanation and statement of the Quasi-random PCP.

Maximum Edge Biclique
Input: An n-by-n bipartite graph G.
Output: A k_1-by-k_2 complete bipartite subgraph of G.
Objective function: Maximize $k_1 \cdot k_2$.

(Uniform) Sparsest Cut
Input: A graph $G = (V, E)$.
Output: A cut, i.e., a partition of V into two disjoint sets S and $\bar{S}$.
Objective function: Minimize the *sparsity* $E(S, \bar{S})/\big(|S||\bar{S}|\big)$, where $E(S, \bar{S})$ denotes the number of edges crossing the cut.

Minimum Linear Arrangement
Input: A graph $G = (V, E)$.
Output: A permutation of the vertices, i.e., a one-to-one function $\pi : V \to \{1, 2, \ldots, |V|\}$.
Objective function: Minimize $\sum_{\{u,v\}\in E} |\pi(v) - \pi(u)|$.

The famous PCP theorem, by Arora and Safra [6] and Arora et al. [4], can be stated as follows.

THEOREM 1.1. *Given a SAT formula ϕ of size n we can in time polynomial in n construct a set of M tests satisfying the following:*

1. *Each test queries a constant number d of bits from a proof, and based on the outcome of the queries it either accepts or rejects ϕ.*
2. *(YES Case/Completeness) If ϕ is satisfiable, then there exists a proof so that all tests accept ϕ.*
3. *(NO Case/Soundness) If ϕ is not satisfiable, then no proof will cause more than $M/2$ tests to accept ϕ.*

Note that by picking one test at random, one can look at only a constant number of bits of a given proof and then with good probability know whether the given proof is correct or not. Therefore, such proofs are called probabilistically checkable proofs (PCPs). The algorithm that constructs a set of such tests with the goal of distinguishing between correct and incorrect proofs will be referred to as a PCP verifier.

Khot [22] introduced the notion of Quasi-random PCPs. The idea is to focus on the distribution (as opposed to the outcome) of queries made by the verifier. The distribution is required to depend on whether the input to the PCP verifier is a YES or a NO instance. In the NO case, the queries are required to be distributed randomly over the proof; i.e., given any set B of half the bits, if each test queries d bits from the proof, then only a fraction $(1/2)^d$ of the tests is expected to query bits only from B. In the YES case, the distribution is required to be far from random. Since the verifier does not know whether the input is a YES or NO instance, it seems quite counterintuitive at first sight that he can make his query pattern depend on the YES/NO case. However, consider the PCP verifier by Holmerin and Khot [20]: each test of their verifier queries three bits from a balanced proof, i.e., a proof with an equal number of 1-bits and 0-bits, and accepts if and only if the exclusive-or of the three queried bits is zero. Suppose the tests of this verifier query the same bits of the proof regardless of whether it is a YES or NO instance; then the tests that accept in the YES case will also accept in the NO case (given the same proof). It is thus necessary that the query pattern depend on the YES/NO case, without the verifier knowing which case it is.

The following Quasi-random PCP construction by Khot [22] will be the starting point for our reductions and can be stated as follows.

THEOREM 1.2 (see [22]). *For every $\epsilon > 0$, given a SAT formula ϕ of size n, we can in time $2^{O(n^\epsilon)}$ probabilistically construct a set of $M = 2^{O(n^\epsilon)}$ tests satisfying the following with high probability:*

1. *Each test queries $d = O\left(\frac{1}{\epsilon}\log\frac{1}{\epsilon}\right)$ bits from a proof of length $N = 2^{O(n^\epsilon)}$.*
2. *Each bit of the proof is queried by dM/N tests (queries are uniformly distributed over the proof).*
3. *(YES Case/Completeness) If ϕ is satisfiable, then there exists a set of half the bits (corresponding to the 0-bits in a correct proof) so that βM tests query bits only from this set, where $\beta = (1 - O(1/d))\,\frac{1}{2^{d-1}}$.*
4. *(NO Case/Soundness) For $p > 0$, let B be any subset of bits of size $(1/2+p)N$. If ϕ is not satisfiable, then at most $(\alpha + pd)M$ tests query bits only from B, where $\alpha = \frac{1}{2^d} + \frac{1}{2^{20d}}$.*[2]

1.2. Results and proof ideas. The main results of this paper are summarized in the following theorem.

THEOREM 1.3. *Let $\epsilon > 0$ be an arbitrarily small constant. If there is a PTAS for Sparsest Cut, Minimum Linear Arrangement, or Maximum Edge Biclique, then there is a (probabilistic) algorithm that decides whether a given SAT instance of size n is satisfiable in time 2^{n^ϵ}.*

Proof overview. The hardness of approximation follows by presenting reductions from the Quasi-random PCP [22]. The reductions to Maximum Edge Biclique, Sparsest Cut, and Minimum Linear Arrangement are presented in sections 2, 3, and 4, respectively. They all follow a general pattern that is sketched below. We start by building a graph instance of the addressed problem with vertices corresponding to proof bits and tests of the Quasi-random PCP. The graph is created in such a way that the vertices corresponding to tests ("test-vertices") have a relatively low impact on the total solution cost. This is achieved by having a relatively small number of test-vertices. Moreover, when test-vertices are disregarded, then any optimal solution

[2]We note that in [22] the soundness says that for any set of half the bits, at most αM tests query bits only from this set. The soundness here follows easily by using that each bit is queried by dM/N tests.

is balanced; that is, bit-vertices are evenly partitioned into two parts in the solution. Since test-vertices have low impact on the total cost, one can prove that any "good" solution must be quasi-balanced; i.e., bit-vertices are roughly evenly partitioned into two parts in the solution. By the construction of the graph, test-vertices that correspond to tests that query bits only on one side of the partition have a lower cost (referred to as good test-vertices). The gap then follows by noting that, by Theorem 1.2, it is hard to decide whether there are "many" or "few" good test-vertices.

We remark that since the gaps obtained by using Theorem 1.2 are very small, we have not optimized our reductions in favor of simplicity. ∎

The hardness factor for Maximum Edge Biclique can be boosted, as was done for Balanced Bipartite Clique in [22].

THEOREM 1.4. *Let $\epsilon > 0$ be an arbitrarily small constant. Assume that SAT does not have a probabilistic algorithm that decides whether a given instance of size n is satisfiable in time 2^{n^ϵ}. Then there is no polynomial (possibly randomized) algorithm for Maximum Edge Biclique that achieves an approximation ratio of $1/N^{\epsilon'}$ on graphs of size N, where ϵ' depends only on ϵ.*

The proof of this theorem is omitted, as it is identical to the one given for boosting the hardness for Balanced Bipartite Clique [22], which in turn is based on the techniques used by Berman and Schnitger [7] for the Clique problem.

2. Maximum Edge Biclique.

In this section we present a reduction from the Quasi-random PCP construction given by Theorem 1.2 to the Maximum Edge Biclique problem so that in the completeness case the graph has an edge biclique with "large" value, whereas in the soundness case all edge bicliques have "small" value (see section 2.5 for details on the achieved gap). We first present the construction (section 2.1) followed by an important property of the constructed graph (section 2.2). We then present the completeness and soundness analyses (sections 2.3 and 2.4).

Since the reduction and analysis are relatively easy, this section serves as a good starting point before continuing to the more complex reductions (which follow the same general pattern) in sections 3 and 4.

2.1. Construction.

Let N be the proof size, and let M be the total number of tests of the PCP verifier in Theorem 1.2. Both N and M are bounded by $2^{O(n^\epsilon)}$, where n is the size of the original SAT formula. Let d be the integer as in Theorem 1.2. Select w to be $(\frac{\beta-\alpha}{12\cdot d})^2$ (very small), where β and α are the bounds given by the completeness and soundness of Theorem 1.2. Hence, $\beta = (1 - O(1/d))\frac{1}{2^{d-1}}$ and $\alpha = \frac{1}{2^d} + \frac{1}{2^{20d}}$.

Construct a bipartite graph $G(V, W, E)$ with $|V| = |W|$ as follows (for an overview of the construction see Figure 2.1). The right-hand side (RHS) consists of N bit-vertices corresponding to the bits in the PCP proof and M test-vertices corresponding to the tests of the PCP verifier. The left-hand side (LHS) consists of N bit-vertices corresponding to the bits in the PCP proof and M slack-vertices to keep the bipartite graph balanced. (The slack-vertices are not adjacent to any vertices and are thus not included in any bipartite clique.) Connect an LHS bit-vertex to all RHS bit-vertices *except* the one corresponding to the same bit of the proof. Furthermore, connect it to an RHS test-vertex if and only if the bit is *not queried* by the test. Finally, assume that $w\frac{N}{2} = M$. (This can be achieved by simply copying vertices: every bit-vertex is replaced by c_N copies of itself, and every test-vertex is replaced by c_M copies of itself, such that now $wN/2 = M$ holds. Copies are connected if and only if the original vertices were. Any maximal biclique must take none or all the copies of a vertex on either side of G.)

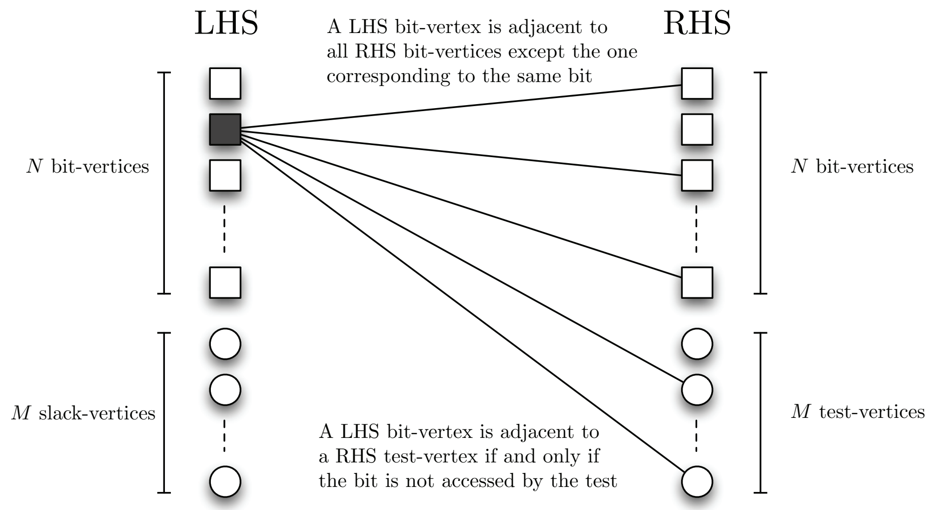

FIG. 2.1. *An example of the construction. Only the edges incident to the dark gray bit-vertex are depicted.*

The intuition behind the construction is the following. As there are many more bit-vertices than test-vertices, any Maximum Edge Biclique must include approximately half of the bit-vertices of the LHS and the remaining bit-vertices of the RHS (see section 2.2). We then use Theorem 1.2 together with the fact that bit-vertices are partitioned into two sets of approximately equal size to analyze the completeness and soundness (see sections 2.3 and 2.4, respectively).

2.2. An optimal edge biclique is quasi-balanced. Given a biclique, let L and R denote the number of bit-vertices of the LHS and bit-vertices of the RHS that are included in the biclique, respectively. Note that in any maximal edge biclique $L+R=N$. We say that a biclique is *quasi-balanced* if $|L-R| \le \frac{\beta-\alpha}{6d}N$.

The following lemma follows in a straightforward manner from the fact that we have many more bit-vertices than test-vertices in our constructed biclique instance.

LEMMA 2.1. *Any optimal edge biclique is quasi-balanced.*

Proof. Any balanced biclique of G, i.e., a biclique with $L=R=N/2$, has value at least $\left(\frac{N}{2}\right)^2$, which serves as a lower bound on the optimal solution. Now consider a biclique with $L=\frac{1+b}{2}N$ and $R=\frac{1-b}{2}N$, where $|b| > \frac{\beta-\alpha}{6d}$. Taking all test-vertices in the biclique gives us the upper bound:

$$L(R+M) = \frac{1+b}{2}N\left(\frac{1-b}{2}N+M\right) = (1-b^2+bw+w)\frac{1}{4}N^2. \tag{2.1}$$

The statement follows by recalling $w=(\frac{\beta-\alpha}{12\cdot d})^2$ and observing the following:

1. The maximum of $f(x)=-x^2+xw+w$ is achieved when $x=\frac{w}{2}<\frac{\beta-\alpha}{6d}$.
2. $f(b)=-b^2+bw+w \le -(\frac{\beta-\alpha}{6d})^2+\frac{\beta-\alpha}{6d}(\frac{\beta-\alpha}{12d})^2+(\frac{\beta-\alpha}{12d})^2<0$.

We thus have that the value of $f(b)$ is always less than 0 when $|b|>\frac{\beta-\alpha}{6d}$. ∎

2.3. Completeness. Here, we assume a YES instance for the Quasi-random PCP; i.e., the given SAT formula ϕ in Theorem 1.2 is satisfiable. We will see that there is an edge biclique of size at least

$$(1+\beta w)\left(\frac{N}{2}\right)^2. \tag{2.2}$$

This will be achieved by constructing a "balanced" solution, that is, a biclique where the bit-vertices are partitioned into two equal sized sets. By Theorem 1.2, half the bits in the proof, namely the 1-bits in a correct proof, are such that a fraction β of tests does *not* query any of them. Let Γ denote the set of all such tests with $|\Gamma| = \beta M = \beta w \frac{N}{2}$. Now consider the biclique, where the LHS consists of the bit-vertices corresponding to the 1-bits in the proof and the RHS consists of the remaining bit-vertices (corresponding to the 0-bits in the proof) and the test-vertices corresponding to the tests in Γ. This gives an edge biclique of size $\frac{N}{2}\cdot\left(\frac{N}{2}+\beta M\right) = \frac{N}{2}\cdot\left(\frac{N}{2}+\beta w\frac{N}{2}\right) = (1+\beta w)\left(\frac{N}{2}\right)^2$.

2.4. Soundness. Here, we assume a NO instance for the Quasi-random PCP; i.e., the given SAT formula ϕ in Theorem 1.2 is *not* satisfiable. We will see that there is no edge biclique of size

$$\left(1+\frac{\alpha+\beta}{2}w\right)\left(\frac{N}{2}\right)^2. \tag{2.3}$$

By Lemma 2.1, it is enough to bound the value of quasi-balanced edge bicliques. Consider such a quasi-balanced biclique, and let L, R, and T denote, respectively, the number of bit-vertices of LHS, bit-vertices of RHS, and test-vertices of RHS that are included in the biclique.

Note that a test-vertex can be included in a biclique only if it is adjacent to all bit-vertices in the LHS of the biclique. In other words, a test-vertex can be included in a biclique only if the corresponding test queries only bits that correspond to bit-vertices included in the RHS of the biclique. The soundness of Theorem 1.2 says that, for any given set of a fraction $1/2+p$ of the bits, at most a fraction $\alpha + p\cdot d$ of the tests queries only those bits. Hence, any edge biclique with $L=\frac{1-b}{2}N$ and $R=\frac{1+b}{2}N$ has $T \le \left(\alpha+\frac{|b|}{2}d\right)M \le (\alpha+|b|d)w\frac{N}{2}$.

Assuming $|b| \le \frac{\beta-\alpha}{6d}$ (Lemma 2.1), we have the following (rough) bound on the value of any edge biclique of G:

$$\begin{aligned}
L(R+T) &\le \frac{1-b}{2}N\left(\frac{1+b}{2}N+(\alpha+|b|d)w\frac{N}{2}\right)\\
&\le (1+(1+|b|)(\alpha+|b|d)w)\left(\frac{N}{2}\right)^2\\
&\le (1+(\alpha+|b|(2d+\alpha))\,w)\left(\frac{N}{2}\right)^2\\
&< \left(1+\frac{\alpha+\beta}{2}w\right)\left(\frac{N}{2}\right)^2.
\end{aligned}$$

The last inequality holds because

$$\alpha+|b|(2d+\alpha) < \frac{\alpha+\beta}{2} \Leftrightarrow 2|b|(2d+\alpha) < \beta-\alpha,$$

which is easily seen to be true by recalling that $|b| \le \frac{\beta-\alpha}{6d}$ and $\alpha < d$.

2.5. Inapproximability gap. Here, we put everything together to obtain the claimed hardness of approximation result, i.e., that a PTAS for Maximum Edge Biclique implies a (probabilistic) algorithm for SAT that runs in time $2^{O(n^\epsilon)}$, where n is the instance size. By using Theorem 1.2, we have provided a probabilistic reduction Γ from SAT to Maximum Edge Biclique. For any fixed $\epsilon > 0$, given an instance ϕ of SAT of size n, Γ produces an edge biclique instance G in time $2^{O(n^\epsilon)}$ satisfying the following with high probability:

- *(Completeness)* If ϕ is satisfiable, then G has an edge biclique of value

$$(1+\beta w)\left(\frac{N}{2}\right)^2. \tag{2.4}$$

- *(Soundness)* If ϕ is not satisfiable, then all edge bicliques of G have value at most

$$\left(1+\frac{\alpha+\beta}{2}w\right)\left(\frac{N}{2}\right)^2. \tag{2.5}$$

As α, β, and w are all functions of the parameter d of Theorem 1.2, which in turn is a function of ϵ, and since $\alpha < \beta$, the quotient $\frac{(2.5)}{(2.4)}$ is less than $1-\zeta(\epsilon)$ for some $\zeta(\epsilon) > 0$.

Now assume that the Maximum Edge Biclique problem admits a PTAS. Then, by definition, it has a polynomial time $(1-\zeta(\epsilon))$-approximation algorithm $\mathcal{A}_{\zeta(\epsilon)}$ for any fixed $\epsilon > 0$. Moreover, the following (probabilistic) algorithm solves SAT in time $2^{O(n^\epsilon)}$ for any fixed $\epsilon > 0$.

Decide SAT instance $\boldsymbol{\phi}$.

1. Run Γ to obtain a Maximum Edge Biclique instance G from ϕ (in time $2^{O(n^\epsilon)}$).
2. Run $\mathcal{A}_{\zeta(\epsilon)}$ on G to obtain a solution with value val (in time polynomial in the size of G, which is $2^{O(n^\epsilon)}$).
3. If $val \geq (1-\zeta(\epsilon)) \cdot (2.4)$, then ϕ is satisfiable; else ϕ is not satisfiable.

3. Sparsest Cut. We present a reduction from the Quasi-random PCP construction given by Theorem 1.2 to Uniform Sparsest Cut so that in the completeness case the constructed graph has a cut with "small" sparsity, whereas in the soundness case all cuts have "large" sparsity (see section 3.5 for details on the achieved gap). We first present the construction (section 3.1) followed by an important property of the constructed graph (section 3.2). We then present the completeness and soundness analyses (sections 3.3 and 3.4).

3.1. Construction. Let N be the proof size, and let M be the total number of tests of the PCP verifier in Theorem 1.2. Both N and M are bounded by $2^{O(n^\epsilon)}$, where n is the size of the original SAT formula. Let d be the number of bits each test queries as in Theorem 1.2. Select $k = (\frac{10d}{\beta-\alpha})^8$ and $h = k\,(k^2+k+\frac{1}{4})$, where β and α are the bounds given by the completeness and soundness of Theorem 1.2. Hence, $\beta = (1-O(1/d))\frac{1}{2^{d-1}}$ and $\alpha = \frac{1}{2^d} + \frac{1}{2^{20d}}$. Note that $h \gg k \gg 1$. We now describe the construction (for an overview see Figure 3.1). The graph $G=(V,E)$ consists of a bipartite graph G_b and two "huge" cliques of size kMN called C_ℓ and C_r. The graph G_b is a bipartite graph where the LHS consists of M test-vertices corresponding to the tests of the PCP verifier. The RHS consists of N clusters, one for each bit in the PCP proof, where each cluster consists of M bit-vertices. Place edges between a test-vertex to *all* vertices of a cluster if and only if the bit corresponding to that cluster is queried by the test.

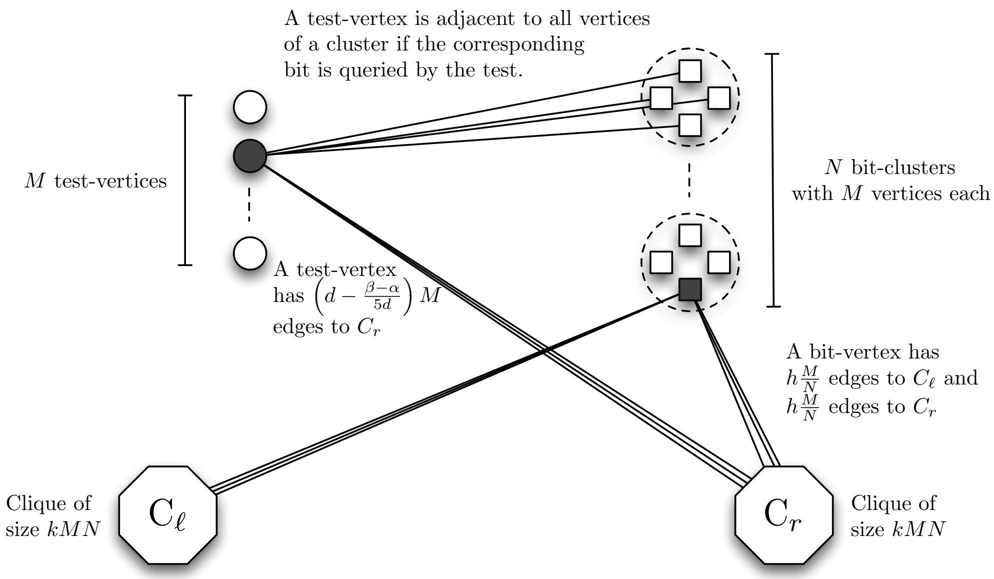

FIG. 3.1. *The graph G for Sparsest Cut. Cliques, bit-vertices, and test-vertices are depicted by polygons, squares, and circles, respectively. For simplicity, only edges incident to dark gray vertices are depicted.*

Finally, we complete the construction of the graph G by connecting the bipartite graph G_b to C_ℓ and C_r as follows. Each bit-vertex has $h\frac{M}{N}$ edges to C_ℓ and $h\frac{M}{N}$ edges to C_r, and each test-vertex has $(d - \frac{\beta-\alpha}{5d})M$ edges to C_r.[3] Furthermore, we distribute the edges incident to the cliques so that the difference of the degree between any two vertices in a clique is at most one.

The intuition behind the construction is the following. For a cut to have low sparsity it is good to divide the vertices into two sets of approximately the same size. As our construction has relatively few test-vertices compared to the number of bit-vertices and the size of the cliques, a cut of small sparsity must place the cliques on different sides and partition the bit-vertices into two sets of approximately the same size (see section 3.2). We then use Theorem 1.2, together with the fact that in any good cut the bit-vertices are partitioned into two sets of approximately equal size, to analyze the completeness and soundness (see sections 3.3 and 3.4, respectively).

3.2. An optimal cut is quasi-balanced. We say that a cut $(S, \bar{S})$ is *quasi-balanced* if it satisfies the following properties:

1. The cliques C_ℓ and C_r are placed on different sides of the cut. Assume, without loss of generality, that the vertices of C_ℓ are included in S and the vertices of C_r are included in $\bar{S}$.
2. Let L and R be the number of bit-vertices in S and $\bar{S}$, respectively. Then $|L - R| < (\frac{\beta-\alpha}{10d})^2 NM$.

[3]Note that no parallel edges are needed since $h\frac{M}{N}$ and $(d - \frac{\beta-\alpha}{5d})M$ are both less than kMN for a sufficiently large N.

The goal of this section is to prove that any optimal sparsest cut must be quasi-balanced. Indeed, if we consider the subgraph induced by all but the test-vertices, then it is easy to see that any sparsest cut is balanced, that is, quasi-balanced with $|L-R|=0$. The intuition is now that the test-vertices have a relatively small impact on the cost and, hence, any optimal sparsest cut must be close to being balanced, i.e., quasi-balanced. For the formal proof, we will need some useful properties of the constructed graph G.

OBSERVATION 3.1.

1. *The number of edges from bit-vertices and test-vertices to a clique is less than* $h\frac{M}{N}\cdot MN + dM^2 = (h+d)M^2$. *By the distribution of edges, a vertex* v *of* C_ℓ *or* C_r *is thus adjacent to at most* $\lceil\frac{(h+d)M^2}{kMN}\rceil = \lceil\frac{(h+d)M}{kN}\rceil < 3h\frac{M}{N}$ *vertices outside the clique.*
2. *A bit-vertex is adjacent to* $2h\frac{M}{N}$ *vertices of the cliques. As queries are uniformly distributed (see Theorem* 1.2*), a bit-vertex is adjacent to at most* $d\frac{M}{N}$ *test-vertices. It follows that a bit-vertex is adjacent to at most* $2h\frac{M}{N}+d\frac{M}{N} \leq 3h\frac{M}{N}$ *vertices.*

LEMMA 3.2. *The graph* G *has a cut* $(S,\bar{S})$ *with sparsity*

$$\frac{E(S,\bar{S})}{|S||\bar{S}|} \leq \frac{1}{N^2}\left(k + \frac{\frac{d}{2}}{k^2+k+\frac{1}{4}}\right). \tag{3.1}$$

Moreover, $E(S,\bar{S}) = O(M^2)$ *in any optimal sparsest cut of* G.

Proof. Consider the cut $(S,\bar{S})$, where S contains all vertices of C_ℓ and the bit-vertices corresponding to half the bits ($\bar{S}$ contains the remaining vertices).

Since the cliques are on different sides of the cut and the solution is "balanced," i.e., the bit-vertices are partitioned into two sets of equal size, we have that $|S||\bar{S}| \geq \left(kMN+\frac{MN}{2}\right)\left(kMN+\frac{MN}{2}\right) = M^2N^2(k^2+k+\frac{1}{4})$.

We continue by calculating $E(S,\bar{S})$. Since all vertices of C_ℓ are in S and all vertices of C_r are in $\bar{S}$, we have that the number of edges between bit-vertices and the cliques that cross the cut is $MN\cdot h\frac{M}{N} = hM^2$. Consider the edges incident to test-vertices. Note that, as each test queries d bits and in G there is a cluster of M bits for each bit, the total number of edges incident to test- and bit-vertices is dM^2. By Theorem 1.2, the queries are uniformly distributed, and thus the total number of edges between the test-vertices and the bit-vertices in S that corresponds to half the bits is $\frac{dM^2}{2}$. Summing up the above observations, we get $E(S,\bar{S}) = M^2\left(h+\frac{d}{2}\right)$. It follows that the sparsity of the cut is

$$\frac{E(S,\bar{S})}{|S||\bar{S}|} \leq \frac{M^2\left(h+\frac{d}{2}\right)}{M^2N^2(k^2+k+\frac{1}{4})},$$

which, by recalling that $h = k\left(k^2+k+\frac{1}{4}\right)$, can be written as

$$\frac{1}{N^2}\left(k+\frac{\frac{d}{2}}{k^2+k+\frac{1}{4}}\right),$$

which is the RHS of (3.1).

Finally, to see that any optimal sparsest cut $(S,\bar{S})$ has $E(S,\bar{S}) = O(M^2)$, note that the total number of vertices of G is $2kNM+NM+M$. Hence $|S||\bar{S}| \leq (|V|/2)^2 = \left(kNM+\frac{NM}{2}+\frac{M}{2}\right)^2 \leq ((k+1)NM)^2 = O\left((NM)^2\right)$ for any cut. Now suppose toward contradiction that there exists an optimal sparsest cut with $E(S,\bar{S}) = \omega(M^2)$. Then $\frac{E(S,\bar{S})}{|S||\bar{S}|} = \omega(1/N^2)$, which contradicts its optimality since we proved that there exists a cut with sparsity $O(1/N^2)$. ∎

We are now ready to prove the main result of this section.

LEMMA 3.3. *Any optimal cut is quasi-balanced.*

Proof. We show that an optimal cut is quasi-balanced by first proving that the cliques are placed on different sides of the cut (Claims 3.4 and 3.5) and then that bit-vertices are partitioned into two sets of almost equal size (Claim 3.6).

We say that a clique is divided in a cut $(S,\bar{S})$ if both sets S and $\bar{S}$ contain vertices of the clique. The intuition behind the following claim is that the cliques are so huge that any cut dividing a clique will have a large number of edges crossing the cut.

CLAIM 3.4. *The cliques C_ℓ and C_r are not divided in any optimal sparsest cut.*

Proof of claim. Given an optimal sparsest cut $(S,\bar{S})$, we prove that all vertices of C_r are placed in either S or $\bar{S}$. (The proof that C_ℓ is not divided is similar and left to the reader.) Let l and r be the number of vertices of C_r in S and $\bar{S}$, respectively. Suppose toward contradiction that $l > 0$ and $r > 0$.

If both l and r are big, say at least $\frac{kNM}{4}$, then we have $E(S,\bar{S}) \geq \left(\frac{kNM}{4}\right)^2$, which contradicts the optimality of the cut since an optimal cut has $E(S,\bar{S}) = O(M^2)$ (see Lemma 3.2).

Now consider case 1: $0 < l < \frac{kNM}{4}$ (case 2: $0 < r < \frac{kNM}{4}$ is symmetric). Let v be a vertex of C_r that is placed in S. We complete the proof by considering the following two subcases.

Case 1.a. Suppose there exists a bit-vertex v_b in $\bar{S}$, and consider what happens with the sparsity if we swap places of v and v_b. As the bit-vertex v_b is adjacent to at most $3h\frac{M}{N}$ vertices in total and v is adjacent to at most $3h\frac{M}{N} + \frac{kNM}{4}$ vertices in S (see Observation 3.1) and to at least $\frac{3kNM}{4}$ vertices in $\bar{S}$ (that belong to C_r), the number of edges that cross the cut will decrease by at least

$$\frac{3kNM}{4} - \frac{kNM}{4} - 3h\frac{M}{N} - 3h\frac{M}{N} > \frac{kNM}{4}$$

(for big enough N). The sizes of the two partitions S and $\bar{S}$ remain unchanged. It follows that the sparsity of the cut will decrease, which contradicts its optimality.

Case 1.b. Suppose there are no bit-vertices in $\bar{S}$. Then all bit-vertices are in S, and we have $|S| \geq NM$, and since $r > 3kNM/4$, we have $|\bar{S}| \geq 3kNM/4$. Consider what happens if we move v to $\bar{S}$. Similar to the case above, the number of edges that cross the cut will decrease by at least $\frac{kNM}{4}$. The new value of the sparsest cut will thus be at most

$$\frac{E(S,\bar{S}) - \frac{kNM}{4}}{(|S|-1)(|\bar{S}|+1)} = \frac{E(S,\bar{S}) - \frac{kNM}{4}}{|S||\bar{S}|(1 - \frac{1}{|S|} + \frac{1}{|\bar{S}|} - \frac{1}{|S||\bar{S}|})}.$$

By using that both $|S|$ and $|\bar{S}|$ are at least NM, we have that the sparsity is at most

$$\frac{E(S,\bar{S}) - \frac{kNM}{4}}{|S||\bar{S}|(1 - \frac{2}{NM})},$$

which is strictly smaller than $\frac{E(S,\bar{S})}{|S||\bar{S}|}$ because (using that we have $E(S,\bar{S}) = O(M^2)$ in an optimal cut)

$$E(S,\bar{S})\left(1 - \frac{2}{NM}\right) \geq E(S,\bar{S}) - O\left(\frac{M}{N}\right) \geq E(S,\bar{S}) - \frac{kNM}{4},$$

again contradicting the optimality of the cut. ◻

Given that the cliques are not divided in an optimal sparsest cut, we now prove that they are placed on different sides. The intuition is that the cliques are so huge that a cut that places them on the same side is very unbalanced, i.e., the product $|S||\bar{S}|$ is small, which in turn will cause the cut to have large sparsity.

CLAIM 3.5. *The cliques C_ℓ and C_r are placed on different sides in any optimal sparsest cut.*

Proof of claim. Suppose toward contradiction that both cliques are placed in, say, S in an optimal sparsest cut $(S, \bar{S})$. Recall that each bit-vertex has $2h\frac{M}{N}$ edges to the cliques and each test-vertex has $(d - \frac{\beta-\alpha}{5d})M$ edges to the clique C_r. It follows that each vertex in $\bar{S}$ has at least $2h\frac{M}{N}$ edges that cross the cut (for big enough N), and the cut has sparsity

$$\frac{E(S,\bar{S})}{|S||\bar{S}|} \geq \frac{2h\frac{M}{N}\cdot|\bar{S}|}{|S||\bar{S}|} \geq \frac{2h\frac{M}{N}}{4kMN} = \frac{k^2+k+\frac{1}{4}}{2N^2}.$$

This contradicts the optimality of the cut by recalling that G has a cut with sparsity (3.1). ◻

By the above claim we can assume that the cliques C_ℓ and C_r are placed on different sides of the cut. We continue by proving that the bit-vertices are partitioned into two sets of almost equal size. The following claim completes the proof of Lemma 3.3.

CLAIM 3.6. *Given an optimal cut $(S,\bar{S})$, let L and R be the number of bit-vertices in S and $\bar{S}$, respectively. Then*

$$|L-R| \leq \left(\frac{\beta-\alpha}{10d}\right)^2 NM.$$

Proof of claim. Since the cliques are placed on different sides of the cut, each bit-vertex has at least $h\frac{M}{N}$ incident edges that cross the cut. It follows that $E(S,\bar{S}) \geq h\frac{M}{N}\cdot MN = hM^2$. Suppose toward contradiction that $L = \frac{1+p}{2}NM$ and $R = \frac{1-p}{2}NM$ with $|p| > (\frac{\beta-\alpha}{10d})^2$. Then the calculations below show that the sparsity of such a cut is greater than (3.1), which contradicts its optimality:

$$\begin{aligned}
\frac{E(S,\bar{S})}{|S||\bar{S}|} &\geq \frac{hM^2}{\left(kMN + \frac{1+p}{2}MN + M\right)\left(kMN + \frac{1-p}{2}MN + M\right)} \\
&= \frac{h}{N^2\left(k^2+k+\frac{1-p^2}{4} + O(\frac{1}{N})\right)} \\
&\geq \frac{1}{N^2}\left(\frac{h+d}{k^2+k+\frac{1}{4}}\right) \text{ (for a big enough } N) \\
&= \frac{1}{N^2}\left(k + \frac{d}{k^2+k+\frac{1}{4}}\right) \text{ (recall } h = k(k^2+k+\tfrac{1}{4})).
\end{aligned}$$

The last inequality holds because we assumed $|p| > (\frac{\beta-\alpha}{10d})^2$ and we have

$$\begin{aligned}
&h\cdot\left(k^2+k+\frac{1}{4}\right) \geq (h+d)\cdot\left(k^2+k+\frac{1-p^2}{4}+O\left(\frac{1}{N}\right)\right) \\
\Leftrightarrow\ &h\cdot\left(\frac{p^2}{4} - O\left(\frac{1}{N}\right)\right) \geq d\cdot\left(k^2+k+\frac{1-p^2}{4}+O\left(\frac{1}{N}\right)\right),
\end{aligned}$$

which can easily be seen to be true by recalling that $h = k(k^2 + k + 1/4)$ and $k = (\frac{10d}{\beta-\alpha})^8$. ∎

The proof of the above claim concludes the proof of Lemma 3.3. ∎

3.3. Completeness. Here, we assume a YES instance for the Quasi-random PCP; i.e., the given SAT formula ϕ in Theorem 1.2 is satisfiable. We will see that there is a cut with sparsity at most

$$\frac{1}{N^2}\left(k + \frac{\frac{d}{2} - \beta\frac{\beta-\alpha}{5d}}{k^2 + k + \frac{1}{4}}\right). \tag{3.2}$$

By Theorem 1.2, half the bits in the proof, namely the 0-bits in a correct proof, are such that a fraction β of the tests accesses only these bits in its queries. Let Γ denote the set of all such tests with $|\Gamma| = \beta M$. We now partition the vertices of G as follows (for an overview see Figure 3.2): S contain the vertices of C_ℓ, the vertices of the clusters corresponding to the 0-bits, and the test-vertices of Γ ($\bar{S}$ contains the remaining vertices).

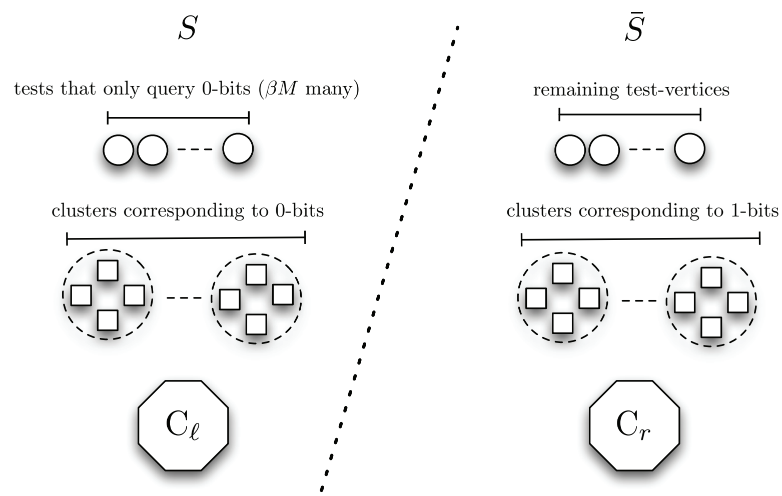

FIG. 3.2. *The cut $(S, \bar{S})$ in the completeness case. For simplicity, no edges are depicted.*

Since the cliques are on different sides of the cut and the solution is "balanced," i.e., the bit-vertices are partitioned into two sets of equal size, we have that $|S||\bar{S}| \geq \left(kMN + \frac{MN}{2}\right)\left(kMN + \frac{MN}{2}\right) = M^2N^2(k^2 + k + \frac{1}{4})$.

We continue by calculating $E(S, \bar{S})$. Since all vertices of C_ℓ are in S and all vertices of C_r are in $\bar{S}$, we have that the number of edges between bit-vertices and the cliques that cross the cut is $MN \cdot h\frac{M}{N} = hM^2$. Consider the edges incident to test-vertices. Note that, as each test queries d bits and G has a cluster of M bits for each bit of the proof, the total number of edges incident to test- and bit-vertices is dM^2. By Theorem 1.2, the queries are uniformly distributed, and thus the total number of edges between the test-vertices and the bit-vertices corresponding to the

0-bits is $\frac{dM^2}{2}$. By observing that the test-vertices of Γ have βdM^2 edges to those bit-vertices and $\beta\,(d - \frac{\beta-\alpha}{5d})\,M^2$ edges to C_r, the total number of edges incident to test-vertices that cross the cut is

$$\frac{dM^2}{2} - \beta dM^2 + \beta\left(d - \frac{\beta-\alpha}{5d}\right)M^2 = M^2\left(\frac{d}{2} - \beta\frac{\beta-\alpha}{5d}\right).$$

Summing up the above observations, we get $E(S,\bar{S}) = M^2\,(h + \frac{d}{2} - \beta\frac{\beta-\alpha}{5d})$, and it follows that the sparsity of the cut is at most

$$\frac{M^2\left(h + \frac{d}{2} - \beta\frac{\beta-\alpha}{5d}\right)}{M^2N^2(k^2+k+\frac{1}{4})},$$

which, by recalling that $h = k\left(k^2+k+\frac{1}{4}\right)$, can be written as

$$\frac{1}{N^2}\left(k + \frac{\frac{d}{2} - \beta\frac{\beta-\alpha}{5d}}{k^2+k+\frac{1}{4}}\right) = \text{ (3.2).}$$

3.4. Soundness. Here, we assume a NO instance for the Quasi-random PCP; i.e., the given SAT formula ϕ in Theorem 1.2 is *not* satisfiable. We will see that all cuts have sparsity at least

$$\frac{1}{N^2}\left(k + \frac{\frac{d}{2} - \frac{\alpha+\beta}{2}\frac{\beta-\alpha}{5d}}{k^2+k+\frac{1}{4}}\right). \tag{3.3}$$

We start by proving a useful property, which is later used to bound the number of "good" test-vertices. Since the construction of G does not necessarily enforce that all bit-vertices of a bit-cluster are placed on the same side of the cut, we cannot apply Theorem 1.2 in a straightforward way. The following lemma is a property of graph G_b (the same bipartite construction and property will be used for Minimum Linear Arrangement in section 4).

LEMMA 3.7. *Consider the bipartite graph G_b, let B be a set of bit-vertices with $|B| \le \frac{1+q}{2}NM$, where $q = (\frac{\beta-\alpha}{10d})^2$, and let T be the set of test-vertices each having at least $(d - \frac{\beta-\alpha}{10d})\,M$ edges to the bit-vertices of B. Then for a NO instance we have that $|T| < \frac{2\alpha+\beta}{3}M$.*

Proof. Note that each bit that is accessed by the test-vertices of T must have at least $(1 - \frac{\beta-\alpha}{10d})M$ bit-vertices in B. Since $|B| \le (1+q)NM/2$, we have that the number of bits in the proof accessed by the tests in T is at most

$$\frac{1+q}{1-\frac{\beta-\alpha}{10d}} \cdot \frac{N}{2} \le \left(1 + q + \frac{\beta-\alpha}{5d}\right)\cdot\frac{N}{2}.$$

The inequality holds because

$$
\begin{aligned}
1+q &\le \left(1-\frac{\beta-\alpha}{10d}\right)\left(1+q+\frac{\beta-\alpha}{5d}\right)\\
&\Leftrightarrow \frac{\beta-\alpha}{10d}\left(1+q+\frac{\beta-\alpha}{5d}\right) \le \frac{\beta-\alpha}{5d},
\end{aligned}
$$

which is true since $q+\frac{\beta-\alpha}{5d}$ is less than one.

The soundness of Theorem 1.2 says that, for any given set of a fraction $(1+q+\frac{\beta-\alpha}{5d})/2$ of the bits, at most a fraction $\alpha+(q+\frac{\beta-\alpha}{5d})\cdot d/2$ of the tests queries only those bits. It follows that

$$
|T| \le \left(\alpha+\left(q+\frac{\beta-\alpha}{5d}\right)\cdot\frac{d}{2}\right)M,
$$

which is less than $\frac{2\alpha+\beta}{3}M$, because

$$
\begin{aligned}
\alpha+\left(q+\frac{\beta-\alpha}{5d}\right)\cdot\frac{d}{2} &< \frac{2\alpha+\beta}{3}\\
\Leftrightarrow \left(q+\frac{\beta-\alpha}{5d}\right)\cdot\frac{3d}{2} &< \beta-\alpha,
\end{aligned}
$$

which can be seen to be true by recalling that $q=(\frac{\beta-\alpha}{10d})^2$. ∎

By Lemma 3.3 we need only consider quasi-balanced cuts. (For an overview of the structure of an optimal cut in the soundness case see Figure 3.3.) We continue by proving that for quasi-balanced cuts the value of $E(S,\bar{S})/(|V|/2)^2$, which is a lower bound on the sparsity of a cut $(S,\bar{S})$, is bounded from below by (3.3).

This is achieved by bounding $E(S,\bar{S})$ as follows. Consider a quasi-balanced cut $(S,\bar{S})$. Let L and R be the bit-vertices in S and $\bar{S}$, respectively. Let Γ be the set of test-vertices each having at least $(d-\frac{\beta-\alpha}{10d})\,M$ edges to the bit-vertices of L. By the fact that the cut is quasi-balanced we have that $\frac{1-q}{2}NM \le |L| \le \frac{1+q}{2}NM$, where $q=(\frac{\beta-\alpha}{10d})^2$, which is sufficient for applying Lemma 3.7, and we get that $|\Gamma| \le \frac{2\alpha+\beta}{3}M$. Since, by Theorem 1.2, the queries are uniformly distributed, the total number of edges between the test-vertices and the bit-vertices of L is at least $\frac{(1-q)dM^2}{2}$. If all test-vertices are placed in $\bar{S}$, all of these edges would cross the cut. The only way to decrease their number is to move test-vertices to S. But since every test-vertex has $(d-\frac{\beta-\alpha}{5d})\,M$ edges to C_r, this is profitable only for test-vertices which have fewer than $\frac{\beta-\alpha}{10d}M$ edges to the bit-vertices of R, i.e., test-vertices that are in Γ. By the above argument we can assume, when calculating a lower bound of $E(S,\bar{S})$, that the only test-vertices placed in S are those in Γ, and it is easy to see that assuming they are not adjacent to any bit-vertices of R might only decrease $E(S,\bar{S})$.

As in the completeness case, we have that the number of edges between bit-vertices and the cliques that cross the cut is $MN\cdot h\frac{M}{N}=hM^2$.

To summarize we have the following:

- The number of edges incident to test-vertices that cross the cut is at least $\frac{(1-q)dM^2}{2}-|\Gamma|dM+|\Gamma|\,(d-\frac{\beta-\alpha}{5d})\,M=\frac{(1-q)dM^2}{2}-|\Gamma|\frac{\beta-\alpha}{5d}M$, which, by using that $|\Gamma|\le\frac{2\alpha+\beta}{3}M$, can be bounded from below by $M^2\,(\frac{(1-q)d}{2}-\frac{2\alpha+\beta}{3}\frac{\beta-\alpha}{5d})$.
- The number of edges between bit-vertices and the cliques that cross the cut is hM^2.

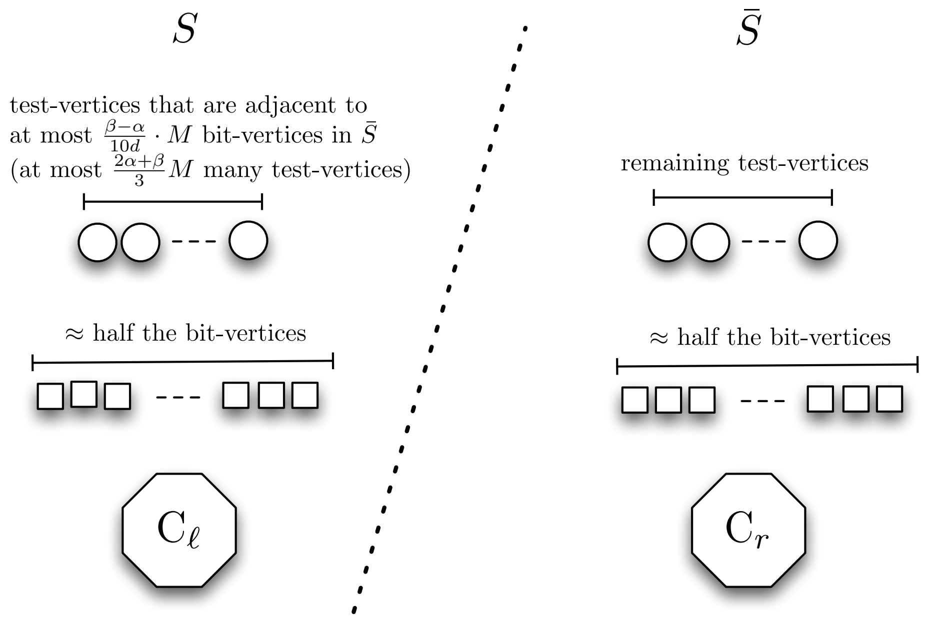

FIG. 3.3. *Structure of an optimal $(S, \bar{S})$ cut in the soundness case. (The edges are not depicted.)*

Since $|S||\bar{S}| \leq (|V|/2)^2$ we have that the sparsity of any cut of G is

$$\begin{aligned}
\frac{E(S,\bar{S})}{|S||\bar{S}|} &\geq \frac{M^2\left(h + \frac{(1-q)d}{2} - \frac{2\alpha+\beta}{3}\frac{\beta-\alpha}{5d}\right)}{\left(kMN + \frac{MN}{2} + M\right)^2} \\
&= \frac{1}{N^2}\left(\frac{h + \frac{(1-q)d}{2} - \frac{2\alpha+\beta}{3}\frac{\beta-\alpha}{5d}}{k^2 + k + \frac{1}{4} + O(\frac{1}{N})}\right) \\
&\geq \frac{1}{N^2}\left(k + \frac{\frac{d}{2} - \frac{\alpha+\beta}{2}\frac{\beta-\alpha}{5d}}{k^2+k+\frac{1}{4}}\right) = (3.3).
\end{aligned}$$

The last inequality holds because $h = k(k^2 + k + 1/4)$ and

$$\begin{aligned}
&\frac{(1-q)d}{2} - \frac{2\alpha+\beta}{3}\frac{\beta-\alpha}{5d} > \frac{d}{2} - \frac{\alpha+\beta}{2}\frac{\beta-\alpha}{5d} \\
&\Leftrightarrow \left(\frac{\alpha+\beta}{2} - \frac{2\alpha+\beta}{3}\right)\frac{\beta-\alpha}{5d} > \frac{qd}{2} \\
&\Leftrightarrow \frac{\beta-\alpha}{6}\frac{\beta-\alpha}{5d} > \frac{qd}{2},
\end{aligned}$$

which is true since $q = (\frac{\beta-\alpha}{10d})^2$.

3.5. Inapproximability gap. Here, we put everything together to obtain the claimed hardness of approximation result, i.e., that a PTAS for (Uniform) Sparsest Cut implies a (probabilistic) algorithm for SAT that runs in time $2^{O(n^\epsilon)}$, where n is the instance size. By using Theorem 1.2, we have provided a probabilistic reduction

Γ from SAT to (Uniform) Sparsest Cut. For any fixed $\epsilon > 0$, given an instance ϕ of SAT of size n, Γ produces a sparsest cut instance G in time $2^{O(n^\epsilon)}$ satisfying the following with high probability:

- *(Completeness)* If ϕ is satisfiable, then G has a cut of sparsity at most

$$\frac{1}{N^2}\left(k + \frac{\frac{d}{2} - \beta\frac{\beta-\alpha}{5d}}{k^2+k+\frac{1}{4}}\right). \tag{3.4}$$

- *(Soundness)* If ϕ is not satisfiable, then all cuts have sparsity at least

$$\frac{1}{N^2}\left(k + \frac{\frac{d}{2} - \frac{\alpha+\beta}{2}\frac{\beta-\alpha}{5d}}{k^2+k+\frac{1}{4}}\right). \tag{3.5}$$

As α, β, and k are all functions of parameter d of Theorem 1.2, which in turn is a function of ϵ, and since $\alpha < \beta$, the quotient $\frac{(3.5)}{(3.4)}$ is greater than $1 + \zeta(\epsilon)$ for some $\zeta(\epsilon) > 0$. The claimed hardness of approximation result now follows from the same arguments as given in section 2.5.

Finally, we mention that the reduction presented in this section is also valid, with almost the same analysis, for the related problem of finding a cut that minimizes the flux $\frac{E(S,\bar{S})}{\min(|S|,|\bar{S}|)}$ (see, e.g., [5]).

4. Minimum Linear Arrangement. For simplicity, we first consider the weighted version of the Minimum Linear Arrangement problem. That is, an edge $\{u,v\} \in E$ has weight w_{uv}, and the objective is to find a permutation π of the vertices that minimizes $\sum_{\{u,v\}\in E} w_{uv}|\pi(u) - \pi(v)|$. We present a reduction from the Quasi-random PCP construction given by Theorem 1.2 to weighted Minimum Linear Arrangement so that in the completeness case the constructed graph has a linear arrangement with "small" cost, whereas in the soundness case all linear arrangements have "large" cost (see section 4.5 for details on the achieved gap). We first present the construction (section 4.1) followed by an important property of the constructed graph (section 4.2). We then present the completeness and soundness analyses (sections 4.3 and 4.4). Finally, we note in section 4.6 that the arguments generalize in a straightforward manner to the unweighted case.

4.1. Construction. Let N be the proof size and M be the total number of tests of the PCP verifier in Theorem 1.2. Both N and M are bounded by $2^{O(n^\epsilon)}$, where n is the size of the original SAT formula. Let d be the number of bits each test queries in the Quasi-random PCP construction. Select k to be $(\frac{10d}{\beta-\alpha})^8$, where β and α are the bounds given by the completeness and soundness of Theorem 1.2. Hence, $\beta = (1 - O(1/d))\frac{1}{2^{d-1}}$ and $\alpha = \frac{1}{2^d} + \frac{1}{2^{20d}}$. Note that $k \gg 1$. We now describe the construction (for an overview see Figure 4.1). The final graph G consists of the graphs G_b, G_ℓ, and G_r and is constructed as follows:

- The graph G_b is a bipartite graph where the LHS consists of M test-vertices corresponding to the tests of the PCP verifier. The RHS consists of N clusters, one for each bit in the PCP proof, where each cluster consists of M bit-vertices. Place edges, weighted by 1, between a test-vertex to *all* vertices of a cluster if and only if the bit corresponding to that cluster is queried by the test. (Note that G_b is the same bipartite graph as in section 3.)
- The graph G_ℓ consists of a vertex C_ℓ and $2kMN$ additional slack-vertices. We place an edge from each slack-vertex to C_ℓ and weight these edges by $k^4\frac{M}{N}$.

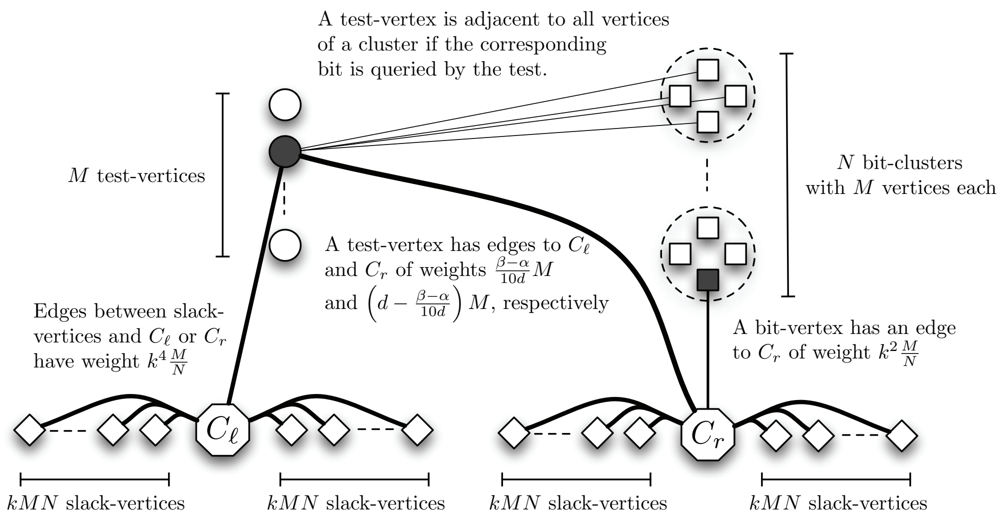

FIG. 4.1. *The graph G for Minimum Linear Arrangement. Slack-vertices, bit-vertices, and test-vertices are depicted by diamonds, squares, and circles, respectively. For simplicity only some edges are depicted and the thickness of an edge is relative to its weight.*

- The graph G_r is constructed as G_ℓ, where instead of C_ℓ we have C_r.

Finally, we construct the graph G by connecting the bipartite graph G_b to G_ℓ and G_r as follows. Each test-vertex has edges to C_ℓ and C_r, weighted by $\frac{\beta-\alpha}{10d}M$ and $(d-\frac{\beta-\alpha}{10d})M$, respectively. Each bit-vertex has an edge to C_r of weight $k^2\frac{M}{N}$.

The intuition behind the construction is the following. As slack-vertices and bit-vertices have edges to C_ℓ and C_r of very large weight, any good linear arrangement will locate these vertices evenly before and after the vertices of C_ℓ and C_r (see Figure 4.2). With this intuition in mind, we prove the important property that any good linear arrangement will partition the bit-vertices into two sets of approximately the same size (see section 4.2). We then use Theorem 1.2 to analyze the completeness and soundness (see sections 4.3 and 4.4, respectively).

Throughout the analyses, we restrict ourselves without loss of generality to linear arrangements where C_ℓ is placed to the left of C_r. The case when C_l is to the right of C_r is symmetric. Moreover, we use the following convention to simplify notation. Let π be a linear arrangement of G. For sets A, B of vertices we write $A <_\pi B$ (subscript omitted when π is clear from the context) whenever $\forall u \in A, \forall v \in B : \pi(u) < \pi(v)$.

4.2. An optimal linear arrangement is quasi-balanced. Select $q = (\frac{\beta-\alpha}{10d})^2$, i.e., a "small" number. We say that a linear arrangement π of G is *quasi-balanced* if (see also Figure 4.2) the following hold:

- The slack-vertices of G_i can be partitioned into two sets S_L^i, S_R^i with $||S_L^i| - |S_R^i|| \le qkNM$ for $i \in \{l, r\}$.
- The bit-vertices can be partitioned into two sets B_L and B_R with $||B_L| - |B_R|| \le qNM$ so that

$$S_L^\ell < \{C_\ell\} < S_R^\ell < B_L < S_L^r < \{C_r\} < S_R^r < B_R.$$

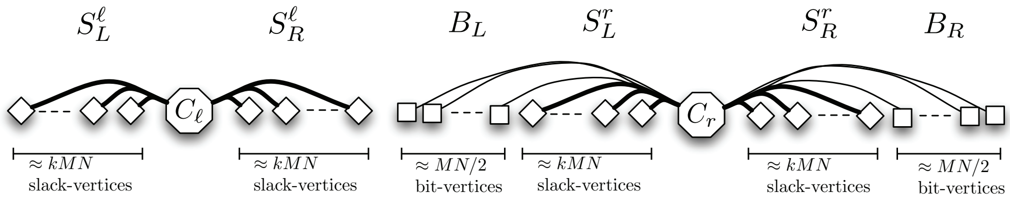

FIG. 4.2. *A quasi-balanced linear arrangement. (The test-vertices are not depicted.)*

The goal of this section is to prove that any optimal linear arrangement is quasi-balanced. Indeed, if we consider the subgraph induced on all but the test-vertices, then it is easy to see that any optimal linear arrangement is balanced, that is, quasi-balanced with $|S_L^i| - |S_R^i| = 0$ for $i \in \{\ell, r\}$ and $|B_L| - |B_R| = 0$. The intuition is that the test-vertices have a relatively small impact on the cost and, hence, any optimal linear arrangement must be close to being balanced, i.e., quasi-balanced. For the formal proof, we will need the following upper bound on the cost of an optimal linear arrangement.

LEMMA 4.1. *The graph G has a linear arrangement with cost at most*

$$M^3 N\left(2k^6 + k^3 + \frac{k^2}{4} + 2dk\right). \tag{4.1}$$

Proof. Partition the slack-vertices of G_i into two sets S_L^i and S_R^i with $|S_L^i| = |S_R^i| = kNM$ for $i \in \{\ell, r\}$. Let B_L be the set of bit-vertices corresponding to a set of half the bits, and let B_R be the remaining bit-vertices. Note that $|B_L| = |B_R| = NM/2$. We also let Γ with $|\Gamma| = M$ be all the test-vertices. Now consider a linear arrangement π of G so that (see Figure 4.2)

$$S_L^\ell < \{C_\ell\} < S_R^\ell < B_L < S_L^r < \Gamma < \{C_r\} < S_R^r < B_R.$$

We proceed by bounding the cost of π by considering the different edges:

- The edges incident to slack-vertices have cost at most

$$4 \cdot k^4 \frac{M}{N} \sum_{i=1}^{kNM} (i + M) = 4 \cdot k^4 \frac{M}{N}\left(\frac{kNM(kNM+1)}{2} + kNM^2\right),$$

  which is bounded from above by $2k^6 M^3 N + o(M^3 N)$.
- The edges between the bit-vertices and C_r have cost at most

$$2 \cdot k^2 \frac{M}{N} \sum_{i=1}^{NM/2} (i + kNM + M).$$

  Since
  (i) $\sum_{i=1}^{NM/2} i = (NM/2)(NM/2+1)/2 = \frac{(NM)^2}{8} + o\left((NM)^2\right)$,
  (ii) $\sum_{i=1}^{NM/2} kNM = k(NM)^2/2$,
  (iii) $\sum_{i=1}^{NM/2} M = o\left((MN)^2\right)$,
  the cost of the edges between bit-vertices and C_r is bounded from above by $M^3 N(\frac{k^2}{4} + k^3) + o(M^3 N)$.

- Now consider the edges incident to test-vertices. As an edge from a test-vertex to C_r has weight $(d - \frac{\beta-\alpha}{10d})M$ and the length of such an edge is at most M in π, the cost of such an edge is at most $(d - \frac{\beta-\alpha}{10d})M^2$. As there are M test-vertices, the cost of all edges from test-vertices to C_r is at most $(d - \frac{\beta-\alpha}{10d})M^3 = o(M^3N)$. Each test-vertex also has an edge of weight $\frac{\beta-\alpha}{10d}M$ to C_ℓ, and such an edge has length at most $(2kMN + MN/2 + M)$ in π. Hence, the cost of all edges from the M test-vertices to C_ℓ is at most $M^3N\frac{\beta-\alpha}{10d}\left(2k + \frac{1}{2}\right) + o(M^3N) \leq M^3Nk + o(M^3N)$. Finally, a test-vertex has at most dM edges to the bit-vertices, each of length at most $(kMN + MN/2 + M)$ in π. Thus the cost of all edges from the M test-vertices to bit-vertices is at most $M^3N\left(dk + \frac{d}{2}\right) + o(M^3N) \leq M^3N(d+1)k + o(M^3N)$. In summary, the total cost of the edges incident to test-vertices is at most $M^3Nk + M^3N(d+1)k + o(M^3N) = M^3N(d+2)k + o(M^3N)$.

Summing up the above observations gives us that the cost of π is at most

$$M^3N\left(2k^6 + k^3 + \frac{k^2}{4} + (d+2)k\right) + o(M^3N),$$

which is (for large enough N and M) less than

$$M^3N\left(2k^6 + k^3 + \frac{k^2}{4} + 2dk\right) = (4.1). \qquad \square$$

We are now ready to prove the main result of this section.

LEMMA 4.2. *Any optimal linear arrangement of G is quasi-balanced.*

Proof. We first prove (Claim 4.3) that in any optimal linear arrangement of G, G_i's slack-vertices can be partitioned into two sets S_L^i, S_R^i for $i \in \{l, r\}$; and bit-vertices can be partitioned into two sets B_L and B_R so that

$$S_L^\ell < \{C_\ell\} < S_R^\ell < B_L < S_L^r < \{C_r\} < S_R^r < B_R. \tag{4.2}$$

Second (Claim 4.4), we will see that the sets must be almost "balanced" in an optimal linear arrangement, that is, $||S_L^i| - |S_R^i|| \leq qkNM$ for $i \in \{l, r\}$ and $||B_L| - |B_R|| \leq qNM$.

CLAIM 4.3. *In any optimal linear arrangement π of G, vertices must be ordered as in* (4.2).

Proof of claim. Since we consider only linear arrangements with C_ℓ to the left of C_r, it is easy to see that

$$S_L^\ell <_\pi \{C_\ell\} <_\pi S_R^\ell <_\pi S_L^r <_\pi \{C_r\} <_\pi S_R^r.$$

Let v_b be a bit-vertex and v_s be a slack-vertex of G_r. Suppose, toward contradiction, that v_b are placed between v_s and C_r, for example, $\pi(v_s) < \pi(v_b) < \pi(C_r)$ (the remaining cases are symmetric and omitted). Consider what happens with the cost if we swap the places of v_b and v_s:

- Vertex v_s is adjacent only to C_r, and this edge has weight $k^4\frac{M}{N}$.
- Vertex v_b has one edge to C_r of weight $k^2\frac{M}{N}$. Since queries are uniformly distributed (see Theorem 1.2), v_b has $d\frac{M}{N}$ edges to test-vertices, each of weight 1. Thus, the total weight of the edges incident to v_b is $d\frac{M}{N} + k^2\frac{M}{N}$.

It follows that by swapping v_b and v_s we *decrease* the cost by at least $(\pi(v_b) - \pi(v_s))\left(k^4\frac{M}{N} - \left(d\frac{M}{N} + k^2\frac{M}{N}\right)\right) > 0$, contradicting the optimality of π.

By the above arguments there are no bit-vertices placed in between slack-vertices and the corresponding vertex C_r (or C_ℓ). We can thus partition the bit-vertices into three sets B_1, B_L, B_R so that

$$B_1 < S_L^\ell < \{C_\ell\} < S_R^\ell < B_L < S_L^r < \{C_r\} < S_R^r < B_R.$$

We complete the proof of this claim by proving that $B_1 = \emptyset$ in an optimal linear arrangement π of G. Suppose, toward contradiction, that $B_1 \neq \emptyset$ in π. Recall that a bit-vertex has an edge of weight $k^2\frac{M}{N}$ to C_r and the total weight of its remaining edges (to test-vertices) is $d\frac{M}{N}$. Furthermore, the total weight of the edges incident to test-vertices is $2dM^2$ (the cost of the edges that are incident to test-vertices and not to the bit-vertices in B_1 might also increase since their length can increase when the bit-vertices in B_1 are moved). Let π' be the linear arrangement

$$S_L^\ell < \{C_\ell\} < S_R^\ell < B_1 < B_L < S_L^r < \{C_r\} < S_R^r < B_R.$$

As $|S_L^\ell| + |S_R^\ell| = 2kNM$, the cost of π' is smaller than the cost of π by at least

$$|B_1|\left(2kMN\left(k^2\frac{M}{N} - d\frac{M}{N}\right) - 2dM^2\right),$$

which is positive whenever $B_1 \neq \emptyset$. $\square$

The following claim completes the proof of Lemma 4.2.

CLAIM 4.4. *In any optimal linear arrangement π of G we have the following:*

- $||S_L^i| - |S_R^i|| \leq qkNM$ *for* $i \in \{l, r\}$.
- $||B_L| - |B_R|| \leq qNM$.

Proof of claim. Let $|S_L^\ell| = (1+s_\ell)kMN$, $|S_R^\ell| = (1-s_\ell)kMN$, $|S_L^r| = (1+s_r)kMN$, $|S_R^r| = (1-s_r)kMN$, $|B_L| = (1+b)MN/2$, and $|B_R| = (1-b)MN/2$, where s_ℓ, s_r, and b may assume negative values.

We proceed by calculating a lower bound on the cost of π by considering the different types of edges:

- The cost of the edges incident to slack-vertices is at least

$$\begin{aligned}&k^4\frac{M}{N}\left(\sum_{i=1}^{(1+s_\ell)kMN} i + \sum_{i=1}^{(1-s_\ell)kMN} i + \sum_{i=1}^{(1+s_r)kMN} i + \sum_{i=1}^{(1-s_r)kMN} i\right)\\ &\geq k^6M^3N\left((1+s_\ell)^2/2 + (1-s_\ell)^2/2 + (1+s_r)^2/2 + (1-s_r)^2/2\right),\end{aligned}$$

  which is equal to $k^6M^3N(2 + s_\ell^2 + s_r^2)$.
- The cost of the edges incident to bit-vertices is at least

$$k^2\frac{M}{N}\left(\sum_{i=1}^{(1+b)MN/2}(i + (1+s_r)kNM) + \sum_{i=1}^{(1-b)MN/2}(i + (1-s_r)kNM)\right).$$

  Since
  (i) $\sum_{i=1}^{(1+b)NM/2} i + \sum_{i=1}^{(1-b)NM/2} i \geq \frac{(NM)^2}{4}\left(\frac{(1+b)^2}{2} + \frac{(1-b)^2}{2}\right) = (NM)^2\frac{1+b^2}{4}$,
  (ii) $\sum_{i=1}^{(1+b)NM/2}(1+s_r)kNM + \sum_{i=1}^{(1-b)NM/2}(1-s_r)kNM \geq (NM)^2(1 - |s_r|)k$,

  the cost of the edges incident to bit-vertices is then bounded from below by $M^3N((1-|s_r|)k^3 + \frac{1+b^2}{4}k^2)$.

Summing up the above observations we have that the cost of π is at least

$$M^3N\left((2+s_\ell^2+s_r^2)k^6+(1-|s_r|)k^3+\frac{1+b^2}{4}k^2\right).$$

As $k=(\frac{10d}{\beta-\alpha})^8$ (a huge number) and $q=(\frac{\beta-\alpha}{10d})^2=(\frac{1}{k})^{1/4}$, the cost of π is greater than the upper bound on an optimal linear arrangement (4.1) whenever $|s_\ell|>q/2$, $|s_r|>q/2$, or $|b|>q$, and the statement follows. □

The proof of Claim 4.4 concludes the proof of Lemma 4.2. □

4.3. Completeness. We will see that there is a linear arrangement with value at most

$$M^3N\left[2k^6+k^3+\frac{k^2}{4}+\left(d+\left(1-\frac{2\beta+\alpha}{3}\right)\frac{\beta-\alpha}{5d}\right)k\right]. \tag{4.3}$$

This will be achieved by constructing a "balanced" linear arrangement. Partition the slack-vertices of G_i into two sets S_L^i and S_R^i with $|S_L^i|=|S_R^i|=kNM$ for $i\in\{\ell,r\}$. Let B_L be the bit-vertices corresponding to the 0-bits in a correct proof, and let B_R be the remaining bit-vertices. Note that $|B_L|=|B_R|=NM/2$. By the completeness of Theorem 1.2, half the bits in the proof, namely the 0-bits in a correct proof, are such that a fraction β of tests accesses only them in their queries. Let Γ denote the set of all such test-vertices with $|\Gamma|=\beta M$, and let $\bar\Gamma$ be the set of the remaining test-vertices.

Now consider the balanced linear arrangement π of G (see Figure 4.3):

$$S_L^\ell<\{C_\ell\}<S_R^\ell<B_L<\Gamma<S_L^r<\bar\Gamma<\{C_r\}<S_R^r<B_R.$$

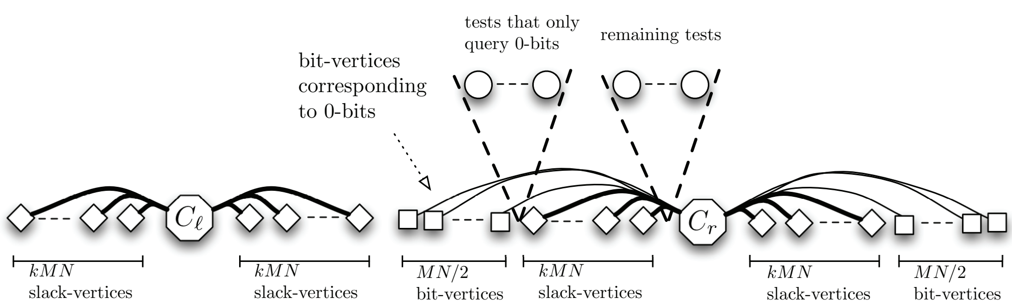

FIG. 4.3. *The linear arrangement π in the completeness case.*

The following lemma concludes the completeness analysis.

LEMMA 4.5. *The cost of π is at most* (4.3) *(for big enough M and N).*

Proof. We need to bound the cost of each edge in the linear arrangement π:

1. As in the proof of Lemma 4.1, both the cost of edges incident to slack-vertices and the cost of edges between the bit-vertices and C_r can be seen to be at most $M^3N(2k^6+k^3+\frac{k^2}{4})+o(M^3N)$.
2. Consider a test-vertex $t\in\Gamma$. As the weight of the edge $\{t,C_r\}$ is $(d-\frac{\beta-\alpha}{10d})M$ and its length in π is at most $kMN+M$, the cost of edge $\{t,C_r\}$ is at most $(d-\frac{\beta-\alpha}{10d})kM^2N+o(M^2N)$. Similarly, as edge $\{t,C_\ell\}$ has weight $\frac{\beta-\alpha}{10d}M$

and its length in π is at most $(kMN + MN/2 + M)$, the cost of $\{t, C_\ell\}$ is at most $M^2N\frac{\beta-\alpha}{10d}(k+1/2) + o(M^2N)$. Finally t has dM edges of weight 1 to bit-vertices in B_L. Since these edges have length at most $(MN/2 + M)$ in π, their cost is at most $M^2Nd/2 + o(M^2N)$.

By the above arguments, the cost of the edges incident to the test-vertices in Γ is at most

$$|\Gamma|M^2N\left(\left(d-\frac{\beta-\alpha}{10d}\right)k+\frac{\beta-\alpha}{10d}\left(k+\frac{1}{2}\right)+\frac{d}{2}\right)+|\Gamma|o(M^2N),$$

which is less than $|\Gamma|M^2N\,(dk+d) + |\Gamma|o(M^2N)$. Using $|\Gamma| = \beta M$, we have that the edges incident to the test-vertices in Γ have cost at most $\beta M^3N\,(dk+d) + o(M^3N)$.

3. Similarly to the above calculations for test-vertices in Γ, the cost of edges incident to test-vertices in $\bar{\Gamma}$ can be seen to be at most

$$\begin{aligned}&(1-\beta)M\left[\underbrace{\frac{\beta-\alpha}{10d}M\left(2kMN+\frac{MN}{2}\right)}_{\text{edges to } C_\ell}+\underbrace{dM\left(kMN+\frac{MN}{2}\right)}_{\text{edges to bit-vertices}}\right]+o(M^3N)\\&=(1-\beta)M\left[\left(d+\frac{\beta-\alpha}{5d}\right)kM^2N+\left(\frac{\beta-\alpha}{20d}+\frac{d}{2}\right)M^2N\right]+o(M^3N),\end{aligned}$$

which is less than $(1-\beta)M^3N\,((d+\frac{\beta-\alpha}{5d})\,k+d)+o(M^3N)$.

We have considered all types of edges of G, and by summing up the above costs we get that the total cost of π is at most

$$\begin{aligned}&M^3N\left[2k^6+k^3+\frac{k^2}{4}+\left(d+(1-\beta)\frac{\beta-\alpha}{5d}\right)k+d\right]+o(M^3N)\\&<M^3N\left[2k^6+k^3+\frac{k^2}{4}+\left(d+\left(1-\frac{2\beta+\alpha}{3}\right)\frac{\beta-\alpha}{5d}\right)k\right]=(4.3).\end{aligned}$$

The last inequality holds because

$$\begin{aligned}\left(d+(1-\beta)\frac{\beta-\alpha}{5d}\right)k+d&<\left(d+\left(1-\frac{2\beta+\alpha}{3}\right)\frac{\beta-\alpha}{5d}\right)k\\\Leftrightarrow d&<\left(\beta-\frac{2\beta+\alpha}{3}\right)\frac{\beta-\alpha}{5d}k\\\Leftrightarrow d&<\frac{\beta-\alpha}{3}\cdot\frac{\beta-\alpha}{5d}k,\end{aligned}$$

which is easily seen to be true by recalling that $k = (\frac{10d}{\beta-\alpha})^8$. ∎

4.4. Soundness. We will see that all linear arrangements of G have value at least

$$M^3N\left[2k^6+k^3+\frac{k^2}{4}+\left(d+\left(1-\frac{\alpha+\beta}{2}\right)\frac{\beta-\alpha}{5d}\right)k\right]. \tag{4.4}$$

By Lemma 4.2 we need only consider quasi-balanced linear arrangements. We proceed by bounding the cost of such linear arrangements from below by (4.4). Given

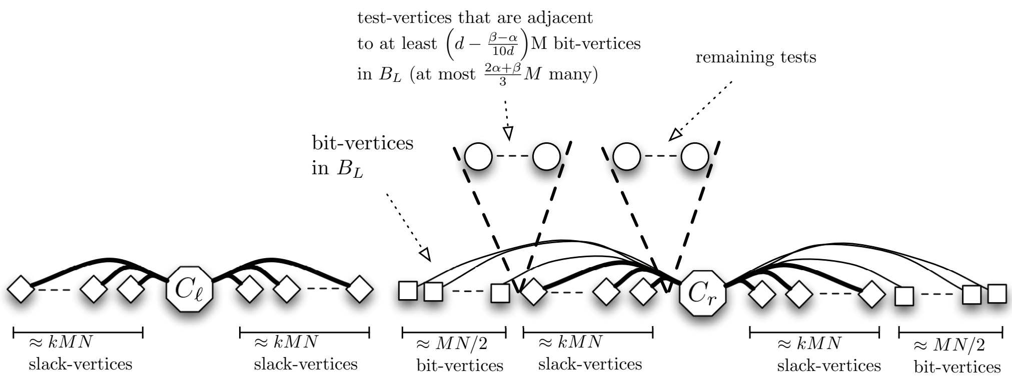

FIG. 4.4. *The structure of an optimal linear arrangement π in the soundness case.*

a quasi-balanced linear arrangement π of G (see Figure 4.4), let Γ be the set of test-vertices that have at least $(d - \frac{\beta-\alpha}{10d})\,M$ edges to B_L in π. Since $|B_L| \le \frac{1+q}{2}NM$, we can apply Lemma 3.7 and get $|\Gamma| < \frac{2\alpha+\beta}{3}M$.

The following lemma follows from an easy case analysis, and its proof is given in the next subsection.

LEMMA 4.6. *In any quasi-balanced linear arrangement π of G, the cost of the edges incident to a test-vertex t is at least*

$$\begin{cases} (1-q)M^2Ndk & \text{if } t \in \Gamma, \\ (1-q)M^2N\left(d + \frac{\beta-\alpha}{5d}\right)k & \text{if } t \notin \Gamma. \end{cases}$$

The above lemma, together with $|\Gamma| < \frac{2\alpha+\beta}{3}M$, implies that the total cost of the edges incident to the M test-vertices is at least

$$(1-q)M^3N\left(d + \left(1 - \frac{2\alpha+\beta}{3}\right)\frac{\beta-\alpha}{5d}\right)k. \tag{4.5}$$

As noted in section 4.2, the cost of the edges not incident to test-vertices is minimized by a balanced linear arrangement (see Figure 4.2) and is thus bounded from below by

$$4k^4\frac{M}{N}\sum_{i=1}^{kMN} i + 2k^2\frac{M}{N}\sum_{i=1}^{MN/2}(i + kMN), \tag{4.6}$$

which is greater than $M^3N(2k^6 + k^3 + \frac{k^2}{4})$.

Summing up (4.5) and (4.6), we have that the total cost of a quasi-balanced linear arrangement is at least

$$\begin{aligned} &M^3N\left[2k^6 + k^3 + \frac{k^2}{4} + (1-q)\left(d + \left(1 - \frac{2\alpha+\beta}{3}\right)\frac{\beta-\alpha}{5d}\right)k\right] \\ &> M^3N\left[2k^6 + k^3 + \frac{k^2}{4} + \left(d + \left(1 - \frac{\alpha+\beta}{2}\right)\frac{\beta-\alpha}{5d}\right)k\right] = \text{(4.4)}. \end{aligned}$$

The last inequality holds because

$$
\begin{aligned}
(1-q)\left(d+\left(1-\frac{2\alpha+\beta}{3}\right)\frac{\beta-\alpha}{5d}\right) &> \left(d+\left(1-\frac{\alpha+\beta}{2}\right)\frac{\beta-\alpha}{5d}\right)\\
\Leftrightarrow \left(\frac{\alpha+\beta}{2}-\frac{2\alpha+\beta}{3}\right)\frac{\beta-\alpha}{5d} &> q\left(d+\left(1-\frac{2\alpha+\beta}{3}\right)\frac{\beta-\alpha}{5d}\right)\\
\Leftrightarrow \frac{\beta-\alpha}{6}\cdot\frac{\beta-\alpha}{5d} &> q\left(d+\left(1-\frac{2\alpha+\beta}{3}\right)\frac{\beta-\alpha}{5d}\right),
\end{aligned}
$$

which is true since $(1-\frac{2\alpha+\beta}{3})\,\frac{\beta-\alpha}{5d} < 1$ and $q = (\frac{\beta-\alpha}{10d})^2$.

4.4.1. Proof of Lemma 4.6. We will repeatedly use the fact that, in any quasi-balanced linear arrangement, S_L^ℓ, S_R^ℓ, S_L^r, and S_R^r all have size at least $(1-q)kNM$, where $q = (\frac{\beta-\alpha}{10d})^2$.

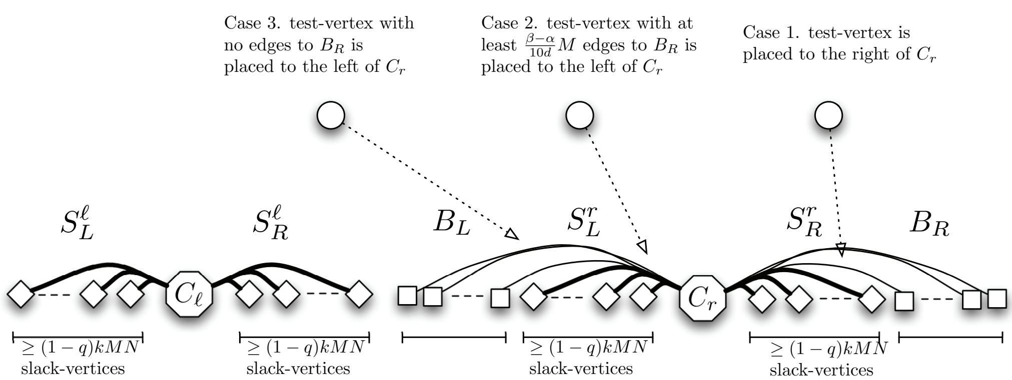

FIG. 4.5. *Overview of the cases considered in the proof of Lemma* 4.6.

We start by proving that any test-vertex that is placed to the right of C_r (Case 1 in Figure 4.5) will have edges of total value at least $(1-q)M^2N\,(d+\frac{\beta-\alpha}{5d})\,k$. Let $p > 0$, and suppose that test-vertex t is placed to the right of $p(1-q)kMN$ slack-vertices of C_r. Since t is placed to the right of C_r, we might only decrease the cost by assuming that all bit-vertices adjacent to t are in B_R. Then the cost of the edges incident to t is at least

$$
\begin{aligned}
&\underbrace{dM(1-p)(1-q)kNM}_{\text{edges to bit-vertices}} + \underbrace{\left(d-\frac{\beta-\alpha}{10d}\right)Mp(1-q)kNM}_{\text{edge to } C_r}\\
&\quad + \underbrace{\frac{\beta-\alpha}{10d}M(2+p)(1-q)kNM}_{\text{edge to } C_\ell}\\
&= M^2N(1-q)k\left(d(1-p)+\left(d-\frac{\beta-\alpha}{10d}\right)p+\frac{\beta-\alpha}{10d}(2+p)\right)\\
&= M^2N(1-q)k\left(d+\frac{\beta-\alpha}{5d}\right).
\end{aligned}
$$

Recall that Γ is the set of test-vertices with at least $(d - \frac{\beta-\alpha}{10d})M$ edges to B_L. Now let $p > 0$, and suppose that t is placed to the left of $p(1-q)kMN$ slack-vertices that are to the left of C_r. On the one hand, if t is not in Γ, then it has at least $\frac{\beta-\alpha}{10d}M$ edges to B_R (Case 2 in Figure 4.5) and the cost of the edges incident to t is at least

$$\underbrace{\left(d - \frac{\beta-\alpha}{10d}\right) M \max[1-p,0](1-q)kNM}_{\text{edges to } B_L} + \underbrace{\left(\frac{\beta-\alpha}{10d}\right) M(1+p)(1-q)kNM}_{\text{edges to } B_R}$$
$$+ \underbrace{\left(d - \frac{\beta-\alpha}{10d}\right) Mp(1-q)kNM}_{\text{edge to } C_r} + \underbrace{\frac{\beta-\alpha}{10d} M(2-p)(1-q)kNM}_{\text{edge to } C_\ell},$$

which can be written as

$$M^2N(1-q)k\left(\left(d - \frac{\beta-\alpha}{10d}\right)\max[1-p,0] + \frac{\beta-\alpha}{10d}(1+p)\right)$$
$$+M^2N(1-q)k\left(\left(d - \frac{\beta-\alpha}{10d}\right)p + \frac{\beta-\alpha}{10d}(2-p)\right),$$

which is easily seen to be at least

$$M^2N(1-q)k\left(d + \frac{\beta-\alpha}{5d}\right).$$

On the other hand, if t is in Γ (Case 3 in Figure 4.5), the cost of the edges incident to t is at least

$$\underbrace{\left(d - \frac{\beta-\alpha}{10d}\right) M \max[1-p,0](1-q)kNM}_{\text{edges to } B_L} + \underbrace{\left(d - \frac{\beta-\alpha}{10d}\right) Mp(1-q)kNM}_{\text{edge to } C_r}$$
$$+ \underbrace{\frac{\beta-\alpha}{10d} M(2-p)(1-q)kNM}_{\text{edge to } C_\ell},$$

which can be written as

$$M^2N(1-q)k\left(\left(d - \frac{\beta-\alpha}{10d}\right)\max[1-p,0] + \left(d - \frac{\beta-\alpha}{10d}\right)p + \frac{\beta-\alpha}{10d}(2-p)\right),$$

and this is easily seen to be at least (bound tight when $p = 1$)

$$M^2N(1-q)kd.$$

The above case distinction concludes the proof of Lemma 4.6.

4.5. Inapproximability gap. Here, we put everything together to obtain the hardness of approximation result that a PTAS for weighted Minimum Linear Arrangement implies a (probabilistic) algorithm for SAT that runs in time $2^{O(n^\epsilon)}$, where n is the instance size. By using Theorem 1.2, we have provided a probabilistic reduction Γ from SAT to weighted Minimum Linear Arrangement. For any fixed $\epsilon > 0$, given an instance ϕ of SAT of size n, Γ produces a weighted Minimum Linear Arrangement instance G in time $2^{O(n^\epsilon)}$ satisfying the following with high probability:

- *(Completeness)* If ϕ is satisfiable, then G has a linear arrangement with cost at most

$$M^3N\left[2k^6+k^3+\frac{k^2}{4}+\left(d+\left(1-\frac{2\beta+\alpha}{3}\right)\frac{\beta-\alpha}{5d}\right)k\right]. \tag{4.7}$$

- *(Soundness)* If ϕ is not satisfiable, then all linear arrangements have cost at least

$$M^3N\left[2k^6+k^3+\frac{k^2}{4}+\left(d+\left(1-\frac{\alpha+\beta}{2}\right)\frac{\beta-\alpha}{5d}\right)k\right]. \tag{4.8}$$

As α, β, and k are all functions of parameter d of Theorem 1.2, which in turn is a function of ϵ, and since $\alpha<\beta$, the quotient $\frac{(4.8)}{(4.7)}$ is greater than $1+\zeta(\epsilon)$ for some $\zeta(\epsilon)>0$. The claimed hardness of approximation result now follows from the same arguments as given in section 2.5.

4.6. Unweighted Minimum Linear Arrangement. In this section we will show that the analysis for weighted Minimum Linear Arrangement can also be used in the unweighted case. Let the graph G be defined as in the construction of weighted Minimum Linear Arrangement (see section 4.1). Note that the edges with weight other than 1 are incident to either C_r or C_ℓ. Recall that $k=(\frac{10d}{\beta-\alpha})^8$. Now consider the graph G_U obtained from G, where we do the following:

1. vertices C_r and C_ℓ are replaced by two "huge" cliques of size k^6M, called C'_r and C'_ℓ, respectively;
2. each edge from a vertex v to C_i with weight w is replaced by w edges from v to w different vertices of C'_i for $i\in\{c,l\}$; and
3. edges are distributed to a clique C'_i so that the difference in the degree of two vertices of a clique is no bigger than one.

With this construction, there are at most $2kMN\cdot k^4\frac{M}{N}+MN\cdot k^2\frac{M}{N}+dM^2=M^2(2k^5+k^2+d)$ edges adjacent to C'_i for $i=\{c,l\}$. Since the edges adjacent to a clique are evenly distributed among its vertices, we have that a vertex of C'_r or C'_ℓ has fewer than M edges to vertices not belonging to the cliques.

We will now see that the soundness and completeness analyses for G_U do not differ much from the analyses done for G.

Completeness. Let π' be the linear arrangement of G_U obtained from the linear arrangement π of G as defined in the completeness analysis of Minimum Linear Arrangement (section 4.3), where the vertices of C'_ℓ and C'_r are placed on the location of C_ℓ and C_r, respectively. By noting that the number of vertices of the cliques is relatively small (of order M) and that the total number of edges is $4kMN\cdot k^4\frac{M}{N}+NM\cdot k^2\frac{M}{N}+M2dM=O(M^2)$, it follows that the value of π' of G_U is only $o(M^3N)$ greater than the value of π of G and the same bound (4.3) holds (for big enough M and N).

Soundness. We say that a clique is *divided* in a linear arrangement π if there exist a bit-, slack-, or test-vertex w and two vertices of the clique u and v such that $\pi(u)<\pi(w)<\pi(v)$. Note that if neither C'_ℓ nor C'_r is divided in an optimal solution of G_U, it follows, by treating the cliques as the vertices C_ℓ and C_r, respectively, that the value of an optimal linear arrangement of G_U must be at least as big as the value of an optimal linear arrangement of G. Thus, the following lemma is enough to complete the soundness analysis of G_U.

LEMMA 4.7. *In any optimal linear arrangement π of G_U, the cliques C'_r and C'_ℓ are not divided.*

Proof. We will present our arguments for the clique C'_r. Since the arguments are the same for C'_ℓ, we leave this case to the reader. Given an optimal linear arrangement π of G_U, let l and r denote the leftmost and rightmost vertices of C'_r in π, respectively, and let $S = \{v \text{ is a slack-, test-, or bit-vertex} : \pi(l) < \pi(v) < \pi(r)\}$. Suppose, toward contradiction, that S is nonempty. Select

$$v_L = \arg\min_{v \in S}(\pi(v)) \quad \text{and} \quad v_R = \arg\max_{v \in S}(\pi(v))$$

(the leftmost vertex and rightmost vertex of S, respectively). Let A denote the set of vertices of C'_r that are placed between l and v_L, i.e., $A = \{v \text{ is a vertex of } C'_r : \pi(l) < \pi(v) < \pi(v_L)\}$. Similarly, let B denote the set of vertices of C'_r that are placed between v_R and r.

Note that the selection of v_L and v_R implies that either $|A|$ or $|B|$ is less than $k^6M/2$. Suppose $|A| \leq k^6M/2$, and consider what happens with the cost if we swap the places of l and v_L:

1. *Edges leaving l.* The number of edges from l to vertices outside the clique is at most M. The cost of these edges will thus *increase* by at most $(\pi(v_L) - \pi(l))M$. The cost of the edges from l to vertices in A will *increase* by

$$\sum_{i \in A} \underbrace{(\pi(v_L) - \pi(i))}_{\text{new cost}} - \underbrace{(\pi(i) - \pi(l))}_{\text{old cost}} \leq \sum_{i \in A} (\pi(v_L) - \pi(l)) + \pi(l) - \pi(i),$$

   which is bounded from above by

$$\sum_{i=1}^{|A|} (\pi(v_L) - \pi(l)) - i \leq (\pi(v_L) - \pi(l))|A| - |A|^2/2.$$

   Finally, the cost of the edges from l to the vertices of C'_r that are not in A will *decrease* by $(\pi(v_L) - \pi(l))(k^6M - |A|) \geq (\pi(v_L) - \pi(l))k^6M/2$.
2. *Edges leaving v_L.* Note that slack-, bit-, and test-vertices have degree at most $2dM$ (for large enough N). The cost of the edges incident to v_L will thus *increase* by at most $(\pi(v_L) - \pi(l))2dM$.

Summing up the above observations we have that the increase of cost will be at most

$$(\pi(v_L) - \pi(l))(M + 2dM + |A| - k^6M/2) - |A|^2/2 < 0,$$

i.e., the cost will decrease, which contradicts the optimality of the linear arrangement. The last inequality follows easily by recalling that $\pi(v_L) - \pi(l) \leq 2k^6M$ (since by the definition of v_L there can be only vertices belonging to the cliques that are placed between l and v_L) and $|A| \leq k^6M/2$ (by assumption).

The remaining case when $|B| \leq k^6M/2$ is symmetric and is omitted. ◻

5. Conclusions and discussion. We have proved the first hardness of approximation results for the classical Minimum Linear Arrangement and (Uniform) Sparsest Cut graph problems. We also obtained hardness results for the Maximum Edge Biclique problem by using a more standard assumption.

All our results are obtained by using the Quasi-random PCP construction by Khot [22]. Hence, our results are under the assumption that SAT is not solvable in probabilistic time 2^{n^ϵ}, where n is the instance size and $\epsilon > 0$ can be made arbitrarily close to 0. Moreover, the hardness factors obtained for Minimum Linear Arrangement and Sparsest Cut by using our reductions from the Quasi-random PCP are tiny. This raises two prominent open problems:

1. Show that it is hard to approximate the addressed problems by using a weaker assumption (ideally $P \neq NP$).
2. Provide a constant factor approximation algorithm for Minimum Linear Arrangement and Uniform Sparsest Cut, or rule out this possibility.

A natural approach for proving that Uniform Sparsest Cut and Minimum Linear Arrangement have no constant approximation algorithms would be to assume the Unique Games Conjecture [21]. Results of this kind have been shown with the stronger assumption that the Unique Games Conjecture is true on expanding graphs (see [3] for Uniform Sparsest Cut and [14] for Minimum Linear Arrangement). However, Arora et al. [3] showed that the Unique Games Conjecture on graphs with relatively high expansion is false. A natural continuation of their work is to understand exactly what kind of expansion is required to prove inapproximability of Sparsest Cut and Minimum Linear Arrangement and whether one can expect Unique Games to be hard with the required expansion. The relation between the Small Set Expansion problem and the Unique Games Conjecture [27] has shed light on this issue, and it was proved [28] that the Minimum Linear Arrangement problem is NP-hard to approximate within any constant factor if the Unique Games Conjecture is true when restricted to instances where all "small" subsets of vertices have high expansion.

Acknowledgments. We are grateful to Andreas Schulz for many helpful discussions during his visit to IDSIA and to Uriel Feige for directing us to the paper [22]. We also thank the anonymous referees for their valuable comments and suggestions.

REFERENCES

[1] C. Ambühl, M. Mastrolilli, and O. Svensson, *Inapproximability results for sparsest cut, optimal linear arrangement, and precedence constrained scheduling*, in Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2007, pp. 329–337.

[2] S. Arora, E. Hazan, and S. Kale, *$O(\sqrt{\log n})$-approximation to sparsest cut in $O(n^2)$ time*, in Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2004, pp. 238–247.

[3] S. Arora, S. Khot, A. Kolla, D. Steurer, M. Tulsiani, and N. K. Vishnoi, *Unique games on expanding constraint graphs are easy: Extended abstract*, in Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC), 2008, pp. 21–28.

[4] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy, *Proof verification and the hardness of approximation problems*, J. ACM, 45 (1998), pp. 501–555.

[5] S. Arora, S. Rao, and U. V. Vazirani, *Expander flows, geometric embeddings and graph partitioning*, in Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC), 2004, pp. 222–231.

[6] S. Arora and S. Safra, *Probabilistic checking of proofs: A new characterization of NP*, J. ACM, 45 (1998), pp. 70–122.

[7] P. Berman and G. Schnitger, *On the complexity of approximating the independent set problem*, Inform. and Comput., 96 (1992), pp. 77–94.

[8] M. Charikar, M. T. Hajiaghayi, H. Karloff, and S. Rao, *l_2^2 spreading metrics for vertex ordering problems*, in Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2006, pp. 1018–1027.

[9] S. Chawla, A. Gupta, and H. Räcke, *Embeddings of negative-type metrics and an improved approximation to generalized sparsest cut*, ACM Trans. Algorithms, 4 (2008), pp. 1–18.

[10] S. Chawla, R. Krauthgamer, R. Kumar, Y. Rabani, and D. Sivakumar, *On the hardness of approximating multicut and sparsest-cut*, Comput. Complexity, 15 (2006), pp. 94–114.

[11] Y. Cheng and G. M. Church, *Biclustering of expression data*, in Proceedings of the Eighth International Conference on Intelligent Systems for Molecular Biology, 2000, pp. 93–103.

[12] M. Dawande, P. Keskinocak, J. M. Swaminathan, and S. Tayur, *On bipartite and multipartite clique problems*, J. Algorithms, 41 (2001), pp. 388–403.

[13] N. R. Devanur, S. Khot, R. Saket, and N. K. Vishnoi, *Integrality gaps for sparsest cut and minimum linear arrangement problems*, in Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC), 2006, pp. 537–546.

[14] N. R. Devanur, S. A. Khot, R. Saket, and N. K. Vishnoi, *On the Hardness of Minimum Linear Arrangement*, manuscript, 2005.

[15] J. Díaz, J. Petit, and M. Serna, *A survey of graph layout problems*, ACM Comput. Surveys, 34 (2002), pp. 313–356.

[16] U. Feige, *Relations between average case complexity and approximation complexity*, in Proceedings of the 34th Annual ACM Symposium on Theory of Computing (STOC), 2002, pp. 534–543.

[17] U. Feige and S. Kogan, *Hardness of Approximation of the Balanced Complete Bipartite Subgraph Problem*, Technical report MCS04-04, Department of Computer Science and Applied Mathematics, The Weizmann Institute of Science, Rehovot, Israel, 2004.

[18] U. Feige and J. R. Lee, *An improved approximation ratio for the minimum linear arrangement problem*, Inform. Process. Lett., 101 (2007), pp. 26–29.

[19] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W. H. Freeman, San Francisco, 1979.

[20] J. Holmerin and S. Khot, *A strong inapproximability gap for a generalization of minimum bisection*, in Proceedings of the 18th Annual IEEE Conference on Computational Complexity (CCC), 2003, pp. 371–378.

[21] S. Khot, *On the power of unique* 2*-prover* 1*-round games*, in Proceedings of the 34th Annual ACM Symposium on Theory of Computing (STOC), 2002, pp. 767–775.

[22] S. Khot, *Ruling out PTAS for graph min-bisection, dense k-subgraph, and bipartite clique*, SIAM J. Comput., 36 (2006), pp. 1025–1071.

[23] S. Khot and N. K. Vishnoi, *The unique games conjecture, integrality gap for cut problems and embeddability of negative type metrics into l_1*, in Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2005, pp. 53–62.

[24] F. T. Leighton and S. Rao, *Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms*, J. ACM, 46 (1999), pp. 787–832.

[25] D. W. Matula and F. Shahrokhi, *Sparsest cuts and bottlenecks in graphs*, Discrete Appl. Math., 27 (1990), pp. 113–123.

[26] R. Peeters, *The maximum edge biclique problem is NP-complete*, Discrete Appl. Math., 131 (2003), pp. 651–654.

[27] P. Raghavendra and D. Steurer, *Graph expansion and the unique games conjecture*, in Proceedings of the 42nd Annual ACM Symposium on Theory of Computing (STOC), 2010, pp. 755–764.

[28] P. Raghavendra, D. Steurer, and M. Tulsiani, *Reductions between expansion problems*, Electron. Colloq. Comput. Complexity, 17 (2010), article 172.

[29] S. Rao and A. W. Richa, *New approximation techniques for some linear ordering problems*, SIAM J. Comput., 34 (2004), pp. 388–404.

[30] D. B. Shmoys, *Cut problems and their application to divide-and-conquer*, in Approximation Algorithms for NP-Hard Problems, PWS, Boston, 1997, pp. 192–235.

[31] L. Trevisan, *Inapproximability of Combinatorial Optimization Problems*, ECCC report TR04-065; available online from http://www.eccc.uni-trier.de/report/2004/065/ (2004).

[32] V. V. Vazirani, *Approximation Algorithms*, Springer-Verlag, Berlin, 2001.