Inapproximability Results for Sparsest Cut, Optimal Linear Arrangement, and Precedence Constrained Scheduling

Christoph Ambühl
University of Liverpool
Liverpool, United Kingdom
christoph@csc.liv.ac.uk

Monaldo Mastrolilli
IDSIA
Lugano, Switzerland
monaldo@idsia.ch

Ola Svensson
IDSIA
Lugano, Switzerland
ola@idsia.ch

Abstract

We consider (Uniform) Sparsest Cut, Optimal Linear Arrangement and the precedence constrained scheduling problem $1|prec|\sum w_jC_j$. *So far, these three notorious NP-hard problems have resisted all attempts to prove inapproximability results. We show that they have no Polynomial Time Approximation Scheme (PTAS), unless NP-complete problems can be solved in randomized subexponential time. Furthermore, we prove that the scheduling problem is as hard to approximate as Vertex Cover when the so-called fixed cost, that is present in all feasible solutions, is subtracted from the objective function.*

1. Introduction

Sparsest Cut and Optimal Linear Arrangement [1] (OLA) are typical cases of classical graph problems for which we have neither a hardness of approximation result, nor a 'good' approximation algorithm. For Sparsest Cut, Arora, Rao & Vazirani [5], by using a Semi-Definite Programming (SDP) relaxation, provided a $O(\sqrt{\log n})$-approximation algorithm, improving an $O(\log n)$-approximation algorithm by Leighton & Rao [19]. For OLA, Feige & Lee [13] observed that combining the techniques in [5] with the rounding algorithm of Rao and Richa [22] yields an $O(\sqrt{\log n}\log\log n)$-approximation algorithm. This improves over the $O(\log n)$-approximation of Rao and Richa. These SDP relaxations were shown to have integrality gap $\Omega(\log\log n)$ by Devanur, Khot, Saket & Vishnoi [10]. As noted in [10], it remains a challenging open problem to prove a hardness of approximation result for Sparsest Cut and OLA. The currently known hardness results apply only to the non-uniform case of Sparsest Cut, and are based on the Unique Games Conjecture [6, 17].

The third problem we address is the classical problem of scheduling precedence constrained jobs on a single machine to minimize the weighted completion time, known as $1|prec|\sum w_jC_j$ in standard scheduling notation [14, 18]. While currently no inapproximability result is known (other than that the problem does not admit a fully polynomial time approximation scheme), there are several 2-approximation algorithms [21, 23, 15, 8, 7, 20, 1]. Narrowing this approximability gap is considered one of the most prominent open problems in scheduling (see e.g. [24]).

In this paper, we show that no PTAS is possible for these problems assuming only the fairly standard assumption $NP \not\subseteq \cap_{\epsilon>0} BPTIME(2^{n^{\epsilon}})$ (i.e. NP-complete problems cannot be solved in randomized sub-exponential time). Our results use the recent Quasi-random PCP construction of Khot [16], who proved important inapproximability results for Graph Min-Bisection, Densest Subgraph, and Bipartite Clique. However, the results in [16] or even the stronger average-case assumptions used by Feige in [11] are not known to generalize to Sparsest Cut and OLA (see e.g. [10, 25]).

We show that the Quasi-random PCP [16], and careful constructions provided in this paper, suffice to rule out the existence of a PTAS for Sparsest Cut (Section 3) and OLA (Section 4). Moreover, we prove that $1|prec|\sum w_jC_j$ has no PTAS by presenting a gap preserving reduction from Maximum Edge Biclique (MEB) (Section 5.1) .

Feige [11] already showed that MEB is hard to approximate by assuming a hypothesis about average-case hardness of Random 3SAT. Improving on the weaker result in [12], we present the inapproximability of MEB (Section 2) based on the more standard assumptions of Khot [16].

To summarize, the reductions presented in this paper are the following.

Understanding the approximability of $1|prec|\sum w_jC_j$ is also interesting because of its relation to Vertex Cover. The objective function of the scheduling problem is split

[1] Also known as Minimum Linear Arrangement.

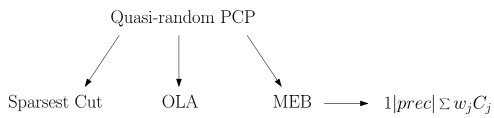

into so-called fixed cost and variable cost (see Section 1.1 for details). Only the variable cost depends on the schedule, whereas the fixed cost is the same for all feasible schedules. In a series of three papers [8, 9, 1], it was established that optimizing the variable cost is a special case of the Vertex Cover problem. In this paper, we show that the variable part is as hard to approximate as the the Vertex Cover problem (Section 5.2). This gives further evidence that the various 2-approximation algorithms for $1|prec|\sum w_j C_j$ might be tight. For a fully satisfactory answer to this problem, a deeper understanding on the interplay between the fixed and the variable costs is needed.

Some of the proofs of these results are quite lengthy and technical. The interested reader can find the omitted details in the full version of the paper [4].

1.1 Preliminaries

We start with the definitions of the addressed problems.

Maximum Edge Biclique (MEB)
Input: A 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* $\frac{E(S,\bar{S})}{|S||\bar{S}|}$, where $E(S, \bar{S})$ denotes the number of edges crossing the cut.

Optimal Linear Arrangement (OLA)
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 $\sum_{\{u,v\}\in E} |\pi(v) - \pi(u)|$.

Single Machine Scheduling $(1|prec|\sum w_j C_j)$
Input: A poset $\mathbf{P} = (N, P)$ consisting of a set of jobs $N = \{1, 2, \ldots, n\}$ and precedence constraints in the form of a partial order P on N. Moreover, each job j has a processing time p_j and a weight w_j, where p_j and w_j are nonnegative integers.
Output: A feasible schedule represented by a linear extension L of P.
Objective function: Minimize

$$\sum_{(i,j)\in L} p_i w_j = \sum_{(i,j)\in P} p_i w_j + \sum_{(i,j)\in L\setminus P} p_i w_j .$$

Remark 1.1 *The first term of the objective function is called* fixed cost *because it is independent of the schedule L. The second term is the* variable part*. It depends on L and is therefore the "interesting" part of the problem.*

Khot [16] introduced the notion of Quasi-random PCPs. The idea is to focus on the distribution 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 input. In the NO case, the queries are required to be distributed randomly over the proof and in the YES case, the distribution is required to be far from random. The following PCP construction will be the starting point for our reductions.

Theorem 1.2 (Khot [16]) *For every $\epsilon > 0$, there exists an integer $d = O(\frac{1}{\epsilon}\log\frac{1}{\epsilon})$ such that there is a PCP verifier for SAT instance of size n satisfying:*

1. *The proof Π for the verifier is of size 2^{n^ϵ}.*

2. *The verifier queries d bits from the proof. Let Q denote the set of query bits.*

3. *Every query is uniformly distributed over Π (two different queries are of course correlated)* [2].

4. *(Yes Case/Completeness:) Suppose the SAT instance is a YES instance and Π is a correct proof. Let Π_0 be the set of 0-bits in the proof (it contains half the bits from the proof). Then*

$$Pr_Q[Q \subseteq \Pi_0] \geq (1 - O(1/d))\frac{1}{2^{d-1}}$$

   The probability is taken over the random test of the verifier.

5. *(No Case/Soundness:) Suppose the SAT instance is a NO instance, and let Π_* be any set of half the bits from Π. Then*

$$\frac{1}{2^d} - \frac{1}{2^{20d}} \leq Pr_Q[Q \subset \Pi_*] \leq \frac{1}{2^d} + \frac{1}{2^{20d}}$$

1.2 Results and Proof Ideas

The main results of this paper are the following Theorem 1.3 and Theorem 1.4.

[2] Since every query is uniformly distributed, all bits in the proof are queried by the same number of tests.

Theorem 1.3 *Let $\epsilon > 0$ be an arbitrarily small constant. If there is a PTAS for Sparsest Cut, Optimal Linear Arrangement, Maximum Edge Biclique or $1|prec|\sum w_j C_j$ then SAT has a (probabilistic) algorithm that runs in time 2^{n^ϵ} where n is the instance size.*

We prove that $1|prec|\sum w_j C_j$ has no PTAS by presenting a gap preserving reduction from Maximum Edge Biclique. The hardness of approximation of the remaining problems follows by presenting reductions from the Quasi-random PCP [16]. They 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 have a relatively low impact on the total solution cost. This is achieved, for the considered problems, by having a relatively small number of test vertices. Moreover, when test vertices are disregarded, then any optimal solution is balanced, that is, bit-vertices are evenly partitioned into two parts in the solution. Furthermore, 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 are adjacent to just one side of the partition have a lower cost (referred to as good test-vertices). The gap follows by noting that, by Theorem 1.2, there are more good test-vertices in the YES-instance; the latter, together with the above construction, is sufficient to introduce a gap.

We remark that since the gaps obtained by using Theorem 1.2 are very small, we have not optimized our reductions in favour of simplicity. Moreover, the inapproximability factor for Maximum Edge Biclique can be boosted as done for Bipartite Clique in [16].

Our second main result shows that the variable part of the scheduling problem is as hard to approximate as Vertex Cover.

Theorem 1.4 *Problem $1|prec|\sum w_j C_j$, where we disregard the fixed cost, is equivalent to Vertex Cover in terms of approximability.*

The proof of Theorem 1.4 heavily relies on the results in [1, 9], which imply that optimizing the variable part of $1|prec|\sum w_j C_j$ is equivalent, under a polynomial reduction, to optimizing a special case of the weighted Vertex Cover problem.

We note that all the currently best known approximation algorithms for several special precedence constraints were obtained by exploiting the Vertex Cover nature of the problem, and deriving better than 2-approximation algorithms for the variable part [3, 2]. It is a natural question to understand if a better than 2-approximate solution for the general version of the problem can be obtained in a similar vein. Theorem 1.4 shows this to be unlikely.

2 Max Edge Biclique

In this section we reduce the PCP construction given by Theorem 1.2 to the Maximum Edge Biclique problem. Since the reduction and analysis are relatively easy, this section serves as a good starting point before continuing to the more complex reductions (that follow the same general pattern) in Sections 3 and 4.

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 $2^{O(n^\epsilon)}$. Let d be the integer as in Theorem 1.2. Select w to be $\left(\frac{\beta-\alpha}{12\cdot d}\right)^2$, where $\alpha := \frac{1}{2^d} + \frac{1}{2^{20d}}$ and $\beta := (1 - O(1/d))\frac{1}{2^{d-1}}$.

Construct a n by n bipartite graph G as follows. 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). 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. Connect a LHS bit-vertex to all RHS bit-vertices *except* the one corresponding to the same bit of the proof and to a RHS test-vertex if and only if the bit is *not accessed* 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 it, and every test-vertex is replaced by c_M copies of it such that now $wN/2 = M$ holds. Any maximal biclique must take none or all the copies of a vertex on either partition of G).

Completeness. We will see that there is an edge biclique of size at least

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

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 the proof, are such that a fraction β of tests do 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 $\frac{N}{2}\cdot\left(\frac{N}{2} + \beta M\right) = \frac{N}{2}\cdot\left(\frac{N}{2} + \beta w \frac{N}{2}\right) = (1 + \beta w)\left(\frac{N}{2}\right)^2$.

Soundness. We will see that there is no edge biclique of size

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

Given a biclique let L, R, and B 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 in any optimal solution $L+R=N$. We say that a biclique is *quasi-balanced* if $||L|-|R|| \leq \frac{\beta-\alpha}{6d}N$.

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. Its proof can be found in the full version [4].

Lemma 2.1 *An optimal edge biclique is quasi-balanced.*

We now proceed by bounding the value of quasi-balanced edge bicliques. Since queries are uniformly distributed, a fraction p of bit-vertices is queried by at most a fraction pd of tests. 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 the test only queries bits included in the RHS of the biclique. By applying Theorem 1.2 we get that any edge biclique with $L=\frac{1-p}{2}N$ and $R=\frac{1+p}{2}N$ has $B \leq \left(\alpha+\frac{|p|}{2}d\right)M \leq (\alpha+|p|d)w\frac{N}{2}$.

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

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

3 Sparsest Cut

We note 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]).

Construction. Let N be the proof size and M be the total number of tests of the PCP verifier in Theorem 1.2. Furthermore, let d be the number of bits each test queries as in that theorem. Note that both M and N are bounded by $2^{O(n^\epsilon)}$. Select $k=\left(\frac{10d}{\beta-\alpha}\right)^8$ and $h=k\left(k^2+k+\frac{1}{4}\right)$, where $\alpha=\frac{1}{2^d}+\frac{1}{2^{20d}}$ and $\beta=(1-O(1/d))\frac{1}{2^{d-1}}$ as in Theorem 1.2. 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. 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 accessed by the test.

Finally, we construct 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. Furthermore, we distribute the edges incident to the cliques so that the difference of degree of two vertices of a clique is at most one. We will prove that there is a gap between the yes/completeness and no/soundness case (see (1) and (2)).

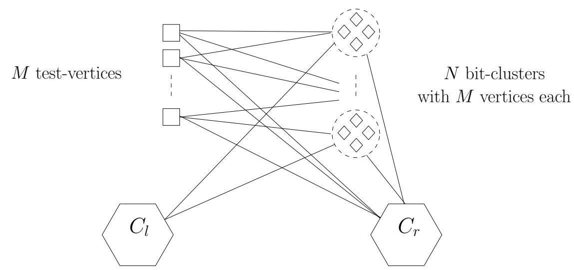

Figure 1. The graph G for Sparsest Cut. Cliques, Bit-vertices, and test-vertices are depicted by polygons, squares and diamonds, respectively.

Completeness. We will see that there is a cut with sparsity at most

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

By Theorem 1.2, half the bits in the proof, namely the 0-bits in the proof, are such that a fraction β of tests do access all queries from them. Let Γ denote the set of all such tests with $|\Gamma|=\beta M$. We now partition the vertices of G as follows. Let S contain the vertices of C_l, the vertices of the clusters corresponding to the 0-bits, and the test-vertices of Γ. 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 constinue 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 crosses the cut is $MN \cdot h\frac{M}{N} = hM^2$. Consider the edges incident to test-vertices. 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 $\beta\left(d - \frac{\beta-\alpha}{5d}\right)M^2$ 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 $E(S,\bar{S}) = M^2\left(h + \frac{d}{2} - \beta\frac{\beta-\alpha}{5d}\right)$ and it follows that the sparsity of the cut is at most (1).

Soundness. We will see that all cuts have sparsity at least

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

We start by proving a useful property, which is later used in the soundness analysis 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 OLA in Section 4).

Lemma 3.1 *Consider the bipartite graph G_b, let B be a set of bit-vertices with $|B| \leq \frac{1+q}{2}NM$, where $q = \left(\frac{\beta-\alpha}{10d}\right)^2$, and let T be the set of test-vertices that have at least $\left(d - \frac{\beta-\alpha}{10d}\right)M$ edges to the bit-vertices of B. Then for a NO-instance we have that $|T| < \frac{2\alpha+\beta}{3}M$.*

In the first step of the soundness analysis, we will prove that for a cut to be small, it needs to be what we call *quasi-balanced*. We then prove 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 (2).

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

1. The cliques C_l and C_r are placed on different sides of the cut. Assume for simplicity that the vertices of C_l are included in S and thus the vertices of C_r are included in $\bar{S}$.
2. Let L and R be the bit-vertices in S and $\bar{S}$, respectively, then $||L| - |R|| < \left(\frac{\beta-\alpha}{10d}\right)^2 NM$.

The proof of the following lemma can be found in the full version [4].

Lemma 3.2 *An optimal cut is quasi-balanced.*

By the above lemma we only need to consider quasi-balanced cuts. We continue by proving that the sparsity of such a cut is at least (2). This is achieved by bounding $E(S,\bar{S})$ as follows. Let $S,\bar{S}$ be a quasi-balanced cut and let Γ be the set of test-vertices that have at least $\left(d - \frac{\beta-\alpha}{10d}\right)M$ edges to the bit-vertices of L. By the fact that the cut is quasi-balanced we have that $\frac{1-q}{2}NM \leq |L| \leq \frac{1+q}{2}NM$, where $q = \left(\frac{\beta-\alpha}{10d}\right)^2$, which is sufficient for applying Lemma 3.1 and we get that $|\Gamma| \leq \frac{2\alpha+\beta}{3}M$. 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 $\left(d - \frac{\beta-\alpha}{5d}\right)M$ edges to C_r, this is only profitable for test-vertices which have less 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 crosses 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 $M^2\left(\frac{(1-q)d}{2} - \frac{2\alpha+\beta}{3}\frac{\beta-\alpha}{5d}\right)$.
- The number of edges, between bit-vertices and the cliques that cross the cut, is hM^2.

Since q is very small and $|S||\bar{S}| \leq (|V|/2)^2$ we have that the sparsity of any cut of G can be bounded from below by (2) (calculations omitted).

4 Optimal Linear Arrangement

For simplicity, we consider the weighted version of OLA. That is, edges have weights and the objective is to find a permutation π of the vertices to minimize $\sum_{\{u,v\}\in E} w_{uv}|\pi(v) - \pi(u)|$. We first present the construction of the OLA instance and then we sketch the completeness and soundness analyzes. Details together with the generalization to the unweighted version can be found in the full version of the paper.

Construction. Let N be the proof size and M be the total number of tests of the PCP verifier in Theorem 1.2. Furthermore, let d be the number of bits each test queries as in that theorem. Note that both M and N are bounded by $2^{O(n^\epsilon)}$. Select k to be $\left(\frac{10d}{\beta-\alpha}\right)^8$, where $\alpha = \frac{1}{2^d} + \frac{1}{2^{20d}}$ and

$\beta = (1 - O(1/d))\frac{1}{2^{d-1}}$ as in Theorem 1.2. The final graph G consists of the graphs G_b, G_l, and G_r 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. 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 accessed 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 $k^4 \frac{M}{N}$.
- 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_r and C_l, weighted by $(d - \frac{\beta-\alpha}{10d})M$ and $\frac{\beta-\alpha}{10d}M$, respectively. Each bit-vertex has an edge to C_r of weight $k^2 \frac{M}{N}$.

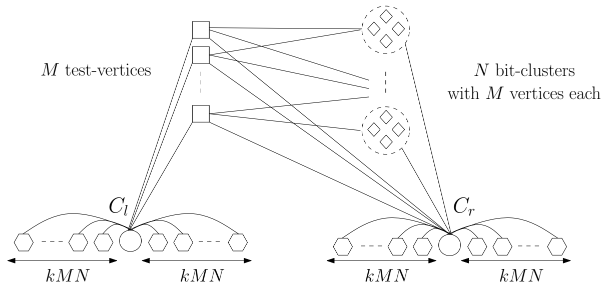

Figure 2. The graph G for OLA. Slack-vertices, bit-vertices, and test-vertices are depicted by polygons, diamonds, and squares, respectively.

We will prove that there is a gap between the yes/completeness and no/soundness case (see (3) and (4)). W.l.o.g, we restrict ourself to only consider 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. 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 $\forall u \in A, \forall v \in B : \pi(u) < \pi(v)$.

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

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

This will be achieved by constructing a so called balanced linear arrangement. We say that a linear arrangement π is *balanced* if the slack-vertices of G_i can be partitioned into two sets S_L^i, S_R^i of equal size, for $i \in \{l, r\}$; and bit-vertices can be partitioned into two equal sized sets B_L and B_R so that

$$S_L^l < \{C_l\} < S_R^l < B_L < S_L^r < \{C_r\} < S_R^r < B_R.$$

(Note that an optimal linear arrangement of the subgraph of G induced by all but the test-vertices is balanced).

By Theorem 1.2, half the bits in the proof, namely the 0-bits in the proof, are such that a fraction β of tests do access all queries from them. 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:

$$S_L^l < \{C_l\} < S_R^l < B_L < \Gamma < S_L^r < \bar{\Gamma} < \{C_r\} < S_R^r < B_R$$

where we let B_L and B_R be the sets of bit-vertices corresponding to 0-bits and 1-bits in the proof, respectively.

Lemma 4.1 *The cost of π is at most* (3) *(for big enough M and N).*

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

1. The cost of edges incident to slack-vertices is at most $4k^4 \frac{M}{N} \sum_{i=1}^{kMN} (i + M) = M^3 N \cdot 2k^6 + o(M^3 N)$.
2. The cost of edge between bit-vertices and C_r is at most $2k^2 \frac{M}{N} \sum_{i=1}^{MN/2} (i + kMN + M) = M^3 N(k^3 + \frac{k^2}{4}) + o(M^3 N)$.
3. The cost of edges incident to test-vertices of Γ is at most (calculations omitted) $\beta M^3 N(dk + d) + o(M^3 N)$.
4. The cost of edges incident to test-vertices of $\bar{\Gamma}$ is at most (calculations omitted) $(1-\beta)M^3 N\left(\left(d + \frac{\beta-\alpha}{5d}\right)k + d\right) + o(M^3 N)$.

We have considered all types of edges of G and the statement follows by summing up the above costs. $\square$

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

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

In the first step of the soundness analysis, we will prove that for a linear arrangement to have low cost, it needs to be what we call *quasi-balanced*. We then prove that the cost of a *quasi-balanced* linear arrangement is bounded from below by (4).

Select $q = \left(\frac{\beta-\alpha}{10d}\right)^2$, i.e., a small number. We say that a linear arrangement π is *quasi-balanced* if the slack-vertices of G_i can be partitioned into two sets S^i_L, S^i_R with $||S^i_L| - |S^i_R|| \le qkNM$, for $i \in \{l, r\}$; and the bit-vertices can be partitioned into two sets B_L and B_R with $||B_L| - |B_R|| \le qNM$ so that

$$S^l_L < \{C_l\} < S^l_R < B_L < S^r_L < \{C_r\} < S^r_R < B_R.$$

Note that a balanced linear arrangement is quasi-balanced with $||S^l_L| - |S^l_R|| = ||S^r_L| - |S^r_R|| = ||B_L| - |B_R|| = 0$. The following lemma implies that we only have to consider quasi-balanced linear arrangements of G (its proof can be found in the full version [4]).

Lemma 4.2 *An optimal linear arrangement of G is quasi-balanced.*

We proceed by bounding the cost of a *quasi-balanced* linear arrangement from below by (4). Given a quasi-balanced linear arrangement π of G, let Γ be the set of test-vertices that have at least $\left(d - \frac{\beta-\alpha}{10d}\right) M$ edges to B_L in π. Since $|B_L| \le \frac{1+q}{2} NM$, we can apply Lemma 3.1 and get $|\Gamma| < \frac{2\alpha+\beta}{3} M$. The following lemma can easily be verified by considering the cost of all different positions of a test-vertex.

Lemma 4.3 *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 $|\Gamma| < \frac{2\alpha+\beta}{3} M$, imply that the total cost of the edges incident to test-vertices is at least

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

As noted in the completeness analysis (an optimal solution to the subgraph of G induced by all but the test-vertices is balanced), the cost of the remaining edges is minimized by a balanced linear arrangement and is thus bounded from below by

$$4k^4 \frac{M}{N} \sum_{i=1}^{kMN} i + 2k^2 \frac{M}{N} \sum_{i=1}^{MN/2} (i + kMN), \tag{6}$$

which is greater than $M^3N(2k^6 + k^3 + \frac{k^2}{4})$. Since q is selected to be very small the total cost (the sum of (5) and (6)) of any linear arrangement π of G can be bounded to be at least (4).

5 Single Machine Scheduling with Precedence Constraints

In the first part of this section we rule out a PTAS for $1|prec|\sum w_jC_j$, by presenting a gap-preserving reduction from MEB. The claim follows by proving (see Section 2) that for MEB it is hard to distinguish between graphs with an edge biclique of size $\ge an^2$ from those having value $< bn^2$, for some $a > b$. In the second part, we show that the variable part of the scheduling problem is equivalent to Vertex Cover in terms of approximability.

5.1 Ruling out a PTAS for the Scheduling Problem

Given a MEB instance $G = (V, W, E)$ with $|V| = |W| = n$, we construct (in polynomial time) an instance S of $1|prec|\sum w_jC_j$ such that the following holds.

Lemma 5.1 *Let opt_m denote the value of the largest edge biclique in G and let opt_s be the value of the optimum schedule for S. Then for any $k > 0$*

$$opt_m < bn^2 \Rightarrow opt_s > \left(k^2 + 2k + 2 - b\right)n^2 \tag{7}$$

$$opt_m \ge an^2 \Rightarrow$$
$$opt_s \le \left(k^2 + 2k + 2 - a + \frac{1}{k^2} + \frac{2}{k}\right)n^2. \tag{8}$$

Hence choosing k large enough such that $\frac{1}{k^2} + \frac{2}{k} < a - b$ will result in a gap-preserving reduction.

Construction. Let $V = \{v_1, v_2, \ldots, v_n\}$ and $W = \{w_1, w_2, \ldots, w_n\}$. We construct the scheduling instance S as follows. The set of jobs are given by $N = D \cup U$, where

$$\begin{aligned} U &= \{w_1, w_2, \ldots, w_n\} \cup \{\mathcal{V}\} \cup \{\bar{v}_1, \bar{v}_2, \ldots, \bar{v}_n\} \text{ and} \\ D &= \{v_1, v_2, \ldots, v_n\} \cup \{\mathcal{W}\} \cup \{\bar{w}_1, \bar{w}_2, \ldots, \bar{w}_n\}. \end{aligned}$$

Note that $|D| = |U| = 2n + 1$ and let k be a large number to be determined later. The processing times and weights of the jobs are given by the following table.

Job	Proc. time	Weight
$\mathcal{V}$	0	kn
$w_1, w_2, \ldots, w_n$	0	$1/k$
$\bar{v}_1, \bar{v}_2, \ldots, \bar{v}_n$	0	1
$\mathcal{W}$	kn	0
$v_1, v_2, \ldots, v_n$	$1/k$	0
$\bar{w}_1, \bar{w}_2, \ldots, \bar{w}_n$	1	0

The precedence constraints are given by the poset $\mathbf{P} = (N, P)$, where

$$\begin{aligned} P \quad = \quad & \{(v_i, w_j) : \{v_i, w_j\} \notin E\} \\ \cup \quad & \{(\bar{w}_i, \mathcal{V}) : i = 1, \ldots, n\} \\ \cup \quad & \{(\mathcal{W}, \bar{v}_i) : i = 1, \ldots, n\} \\ \cup \quad & \{(v_i, \bar{v}_i) : i = 1, \ldots, n\} \\ \cup \quad & \{(\bar{w}_i, w_i) : i = 1, \ldots, n)\} \\ \cup \quad & \{(\mathcal{W}, \mathcal{V})\}. \end{aligned}$$

See Figure 3 for an example.

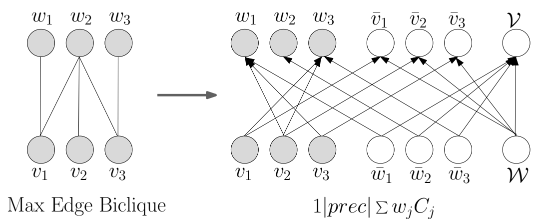

Figure 3. Construction of scheduling instance.

By the processing times and weights we have that $p_i \cdot w_j \leq 1/k$ or $(i,j) \in P$ for all pairs of jobs $(i,j) \in N \times N \setminus H$ where

$$H = \{(v_i, \mathcal{V}), (\mathcal{W}, w_j), (\bar{w}_j, \bar{v}_i) : 1 \leq i \leq n, 1 \leq j \leq m\}.$$

The following lemma shows that large edge bicliques in G imply schedules with low cost in S and vice versa.

Lemma 5.2 *There exists an edge biclique in G with value at least μ if and only if there exists a linear extension L of P so that* $\sum_{(i,j) \in H \cap L} p_i w_j \leq 2n^2 - \mu.$

Since total the cost of all incomparable pairs $(i,j) \notin H$ is $O(n^2/k)$, this lemma allows to prove the bounds (8) and (7) quite easily.

See the full version of the paper [4] for the proofs of Lemmas 5.2 and 5.1.

5.2 Hardness of Variable Part

In this section we prove the following theorem.

Theorem 5.3 *Approximating the variable cost of $1|prec|\sum w_j C_j$ is as hard as approximating Vertex Cover.*

The proof of this theorem is based on quite recent results which are outlined here. Consider any scheduling instance and let $\mathbf{P} = (N, P)$ be the poset that specifies the corresponding precedence constraints. We can associate with $\mathbf{P}$ a graph $G_{\mathbf{P}}$, called the *graph of incomparable pairs*, defined as follows (see also [26, 9, 3]). The vertices of $G_{\mathbf{P}}$ are the incomparable pairs in $\mathbf{P}$. In $G_{\mathbf{P}}$ there is an edge between two incomparable pairs if no linear extension of P reverses both pairs. Given a scheduling instance S with poset P, let $G_{\mathbf{S}}^W$ be the graph of incomparable pairs $G_{\mathbf{P}}$ where each vertex (i,j) has weight $p_i w_j$.

In a series of three papers [8, 9, 1], the following theorem was proven.

Theorem 5.4 *Any vertex cover of $G_{\mathbf{S}}^W$ with weight ω can be turned into a solution to S with variable cost at most ω and vice versa. Both transformations are polynomial.*

Proof of Theorem 5.3. Theorem 5.4 immediately implies that minimizing the variable part of $1|prec|\sum w_j C_j$ is a special case of Vertex Cover and therefore is not harder to approximate.

It remains to prove the other direction. Let $G = (V, E)$ be a Vertex Cover instance and let $n = |V|$. We will construct a scheduling instance S as follows. The construction is inspired by the so-called *adjacency poset* of G. Choose $k > n^2 r/\epsilon$. For each vertex $v_i \in V$, there are two jobs v_i' and v_i''. The processing time and weight for a job v_i' are $1/k^i$ and 0, respectively. Similarly, the processing time and weight for a job v_i'' are 0 and k^i, respectively.

S has the following precedence constraints: For each edge $\{v_i, v_j\} \in E$, the precedence constraints $v_i' \to v_j''$ and $v_j' \to v_i''$. Finally, we add $v_i' \to v_j''$ for every i, j with $i < j$. See Figure 4 for a small example.

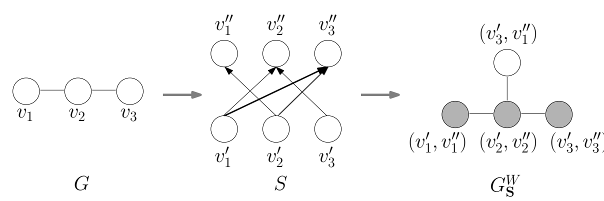

Figure 4. The transformation of a graph G.

Now consider the graph $G_{\mathbf{S}}^W$. It has at most n^2 vertices. The n vertices corresponding to the incomparable pairs (v_i', v_i'') have weight 1. All other vertices have weight at most $1/k$, which by the choice of k is very small. The total weight of these light vertices is no more than n^2/k.

Moreover, the subgraph induced by the vertices with weight 1 is isomorphic to G. To see this, recall that there is an edge between the vertices (v_i', v_i'') and (v_j', v_j'') of $G_{\mathbf{S}}^W$ if and only if both precedence constraints $v_i' \to v_j''$ and $v_j' \to v_i''$ are present in S. This in turn is the case if and only if $(v_i, v_j) \in G$.

Using the connection between S and $G_{\mathbf{S}}^W$ provided by Theorem 5.4 and the close relation between $G_{\mathbf{S}}^W$ and G, it is easy to see that an r-approximation algorithm for the optimum variable cost of $1|prec|\sum w_jC_j$ would imply an approximation algorithm for Vertex Cover with approximation ratio $r(1+n^2/k) < (r+\epsilon)$. $\square$

Acknowledgements

We are grateful to Andreas Schulz for many helpful discussions during his visit at IDSIA and to Uriel Feige for directing us to paper [16].

The first author is supported by Nuffield Foundation Grant NAL32608. The second and third author are supported by Swiss National Science Foundation project 200021-104017/1, "Power Aware Computing", and by the Swiss National Science Foundation project 200020-109854, "Approximation Algorithms for Machine scheduling Through Theory and Experiments II".

References

[1] C. Ambühl and M. Mastrolilli. Single machine precedence constrained scheduling is a vertex cover problem. In *Proceedings of the 14th Annual European Symposium on Algorithms (ESA)*, volume 4168 of *Lecture Notes in Computer Science*, pages 28–39. Springer, 2006.

[2] C. Ambühl, M. Mastrolilli, N. Mutsanas, and O. Svensson. Scheduling with precedence constraints of low fractional dimension. In M. Fischetti and D. P. Williamson, editors, *IPCO*, volume 4513 of *Lecture Notes in Computer Science*, pages 130–144. Springer, 2007.

[3] C. Ambühl, M. Mastrolilli, and O. Svensson. Approximating precedence-constrained single machine scheduling by coloring. In *Proceedings of the APPROX + RANDOM*, volume LNCS 4110, pages 15–26. Springer, 2006.

[4] C. Ambühl, M. Mastrolilli, and O. Svensson. Inapproximability results for sparsest cut, optimal linear arrangement, and precedence constrained scheduling. http://www.idsia.ch/~monaldo/publications.html, 2007.

[5] S. Arora, S. Rao, and U. V. Vazirani. Expander flows, geometric embeddings and graph partitioning. In *STOC*, pages 222–231, 2004.

[6] S. Chawla, R. Krauthgamer, R. Kumar, Y. Rabani, and D. Sivakumar. On the hardness of approximating multicut and sparsest-cut. *Computational Complexity*, 15(2):94–114, 2006.

[7] C. Chekuri and R. Motwani. Precedence constrained scheduling to minimize sum of weighted completion times on a single machine. *Discrete Applied Mathematics*, 98(1-2):29–38, 1999.

[8] F. A. Chudak and D. S. Hochbaum. A half-integral linear programming relaxation for scheduling precedence-constrained jobs on a single machine. *Operations Research Letters*, 25:199–204, 1999.

[9] J. R. Correa and A. S. Schulz. Single machine scheduling with precedence constraints. *Mathematics of Operations Research*, 30(4):1005–1021, 2005.

[10] N. R. Devanur, S. Khot, R. Saket, and N. K. Vishnoi. Integrality gaps for sparsest cut and minimum linear arrangement problems. In *STOC*, pages 537–546, 2006.

[11] U. Feige. Relations between average case complexity and approximation complexity. In *STOC*, pages 534–543, 2002.

[12] 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 Math., The Weizmann Institute of Science, 2004.

[13] U. Feige and J. R. Lee. An improved approximation ratio for the minimum linear arrangement problem. *Inf. Process. Lett.*, 101(1):26–29, 2007.

[14] R. Graham, E. Lawler, J. K. Lenstra, and A. H. G. Rinnooy Kan. Optimization and approximation in deterministic sequencing and scheduling: A survey. In *Annals of Discrete Mathematics*, volume 5, pages 287–326. North–Holland, 1979.

[15] L. A. Hall, A. S. Schulz, D. B. Shmoys, and J. Wein. Scheduling to minimize average completion time: off-line and on-line algorithms. *Mathematics of Operations Research*, 22:513–544, 1997.

[16] S. Khot. Ruling out PTAS for graph min-bisection, densest subgraph and bipartite clique. In *FOCS*, pages 136–145, 2004.

[17] 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 *FOCS*, pages 53–62, 2005.

[18] E. L. Lawler. Sequencing jobs to minimize total weighted completion time subject to precedence constraints. *Annals of Discrete Mathematics*, 2:75–90, 1978.

[19] F. T. Leighton and S. Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *J. ACM*, 46(6):787–832, 1999.

[20] F. Margot, M. Queyranne, and Y. Wang. Decompositions, network flows and a precedence constrained single machine scheduling problem. *Operations Research*, 51(6):981–992, 2003.

[21] N. N. Pisaruk. A fully combinatorial 2-approximation algorithm for precedence-constrained scheduling a single machine to minimize average weighted completion time. *Discrete Applied Mathematics*, 131(3):655–663, 2003.

[22] S. Rao and A. W. Richa. New approximation techniques for some linear ordering problems. *SIAM J. Comput.*, 34(2):388–404, 2004.

[23] A. S. Schulz. Scheduling to minimize total weighted completion time: Performance guarantees of LP-based heuristics and lower bounds. In *Proceedings of the 5th Conference on Integer Programming and Combinatorial Optimization (IPCO)*, pages 301–315, 1996.

[24] P. Schuurman and G. J. Woeginger. Polynomial time approximation algorithms for machine scheduling: ten open problems. *Journal of Scheduling*, 2(5):203–213, 1999.

[25] L. Trevisan. Inapproximability of combinatorial optimization problems. *Electronic Colloquium on Computational Complexity (ECCC)*, (065), 2004.

[26] W. T. Trotter. *Combinatorics and Partially Ordered Sets: Dimension Theory*. Johns Hopkins Series in the Mathematical Sciences. The Johns Hopkins University Press, 1992.