A NOTE ON THE PRECEDENCE-CONSTRAINED CLASS SEQUENCING PROBLEM

JOSÉ R. CORREA*, SAMUEL FIORINI†, AND NICOLÁS STIER-MOSES*

* Operations Research Center, Massachusetts Institute of Technology
77 Massachusetts Ave., Cambridge MA 02139, USA
Email: {jcorrea,nstier}@mit.edu

† GERAD – HEC Montréal, boîte 5521, 3000 chemin de la Côte-Sainte-Catherine,
Montréal (Québec) H3T 2A7, Canada
Email: Samuel.Fiorini@gerad.ca

Abstract. We give a short proof of a result of Tovey [5] on the inapproximability of a scheduling problem known as precedence constrained class sequencing. In addition we present an approximation algorithm with performance guarantee $(c+1)/2$, where c is the number of colors. This improves upon Tovey's c–approximation.

1. Introduction

In this note, we consider the *precedence constrained class sequencing problem* (PCCS), defined by Tovey as the following scheduling problem [5]. Consider an acyclic directed graph $D = (N, A)$ with n nodes, a set C consisting of c colors, and a surjective color function $\psi : N \to C$ that associates each node with a color. *Applying* a color $\omega \in C$ to D, means deleting all nodes in $\psi^{-1}(\omega)$ that have no predecessor with a different color. A sequence of colors $\sigma = \sigma_1\sigma_2 \cdots \sigma_\ell \in C^\ell$ *clears* D if applying colors $\sigma_1, \sigma_2, \ldots, \sigma_\ell$ to D results in the empty graph. The goal is to find the shortest possible sequence σ that clears the graph. In the remainder of the paper, $\mathrm{OPT}(D, C, \psi)$ denotes the length of an optimal sequence.

The PCCS problem appeared earlier in the Operations Research literature under the name "station routing problem" [2, 4] and in the Computer Science literature as the "loop fusion problem" [3]. Typically, nodes represent tasks to be performed, arcs represent precedence constraints between those tasks, and colors represent flexible machines that will perform the tasks. A natural objective is to find a schedule that minimizes the number of setup operations; in many applications the cost incurred when several tasks are performed on the same machine is negligible. This objective corresponds to that of PCCS.

Lofgren, McGinnis and Tovey proved that PCCS is an NP-complete problem [2]. Later, Darte proved that the problem is still NP-complete even if the number of colors c is a constant greater than or equal to three [3], and not part of the input. Using a self-improvability property of PCCS, Lofgren et al. also proved the following result.

Key words and phrases. Precedence-Constrained Scheduling, Hardness of Approximation, Approximation Algorithms.

Note that, here and henceforth, all approximation algorithms are assumed to run in polynomial time.

Theorem 1 ([2, Theorem 4.3])**.** *If there exists an approximation algorithm for PCCS with a constant guarantee, then the problem has a polynomial time approximation scheme* (PTAS)*.*

In the same paper, the authors conjectured that no constant factor approximation algorithm exists, unless P = NP. Very recently, Tovey proved this conjecture [5]. However, his proof is rather long (about 7 pages in this journal) and involved. The main purpose of this note is to present a simple proof of the same result. In addition, we present an approximation algorithm for the problem that has performance guarantee $(c+1)/2$. This algorithm can be turned into a $\sqrt{n/2}$–approximation algorithm, improving upon Tovey's $\sqrt{n}$–approximation, which implicitly contains a c–approximation. In case the number of colors is a constant, we give an algorithm with a performance guarantee of $c/2+\epsilon$.

2. Results

Let us start with the short proof of Tovey's result, which is based on a reduction from Vertex Cover. Given an undirected graph $G=(V,E)$, a vertex cover W of G is a set of vertices such that every edge in E is adjacent to at least one vertex in W. Berman and Fujito proved that, in graphs of maximum degree 3, finding a minimum vertex cover is NP-hard and that there is no better than 7/6–approximation algorithm, unless P = NP [1].

Theorem 2 ([5, Theorem 7])**.** *No approximation algorithm for PCCS with a constant guarantee exists, unless P = NP.*

Proof. In view of Theorem 1, it suffices to show that there is some $\epsilon > 0$ such that PCCS is not approximable within $1+\epsilon$ in polynomial time, provided P $\neq$ NP. We prove that if this is not the case, Vertex Cover for graphs with maximum degree 3 has a PTAS, which contradicts the result in [1]. Therefore, let us assume that PCCS has a PTAS.

Consider a simple and undirected graph G without isolated nodes and of maximum degree 3. Let τ denote the minimum size of a vertex cover of G. Since G does not have isolated vertices, it has at least $|V|/2$ edges. Therefore, as G has maximum degree 3, we have $|V| \leq 6\tau$. Starting from G, we define an instance (D,C,ψ) of PCCS, where $D=(N,A)$, as follows. Let $C=V$ be the set of colors, and for each vertex $v \in V$, create two nodes v' and v'' in N, both with color v, i.e., $\psi(v')=\psi(v'')=v$. For each edge vw in G, create two arcs (v',w'') and (w',v'') in A.

Consider any vertex cover W of G with cardinality t. Any sequence of the form $\sigma=\alpha\overline{\alpha}\alpha$ clears D, where α (resp. $\overline{\alpha}$) is any enumeration of W (resp. $V \setminus W$). As the length of σ is $t+|V|$, we have $\mathrm{OPT}(D,C,\psi) \leq \tau+|V|$. Now, consider a clearing sequence σ of length ℓ, and let W be the set of colors appearing at least twice in σ. Clearly, the set W is a vertex cover of G with cardinality at most $\ell-|V|$. It follows that we can find in polynomial time a vertex cover of G with cardinality at most

$$(1+\epsilon)(\tau+|V|)-|V|=(1+\epsilon)\tau+\epsilon|V| \leq (1+\epsilon)\tau+6\epsilon\tau=(1+7\epsilon)\tau,$$

for any fixed $\epsilon > 0$. Therefore, Vertex Cover for graphs with maximum degree 3 has a PTAS, a contradiction. □

We now give an approximation algorithm with performance guarantee $(c+1)/2$. Assume we have an instance (D, C, ψ), where $C = \{1, \ldots, c\}$. Consider the sequences $\mu = \mu_1 \, \mu_2 \cdots \mu_c$, and $\nu = \nu_1 \, \nu_2 \cdots \nu_c$, defined by $\mu_i = i$ and $\nu_i = c+1-i$. Our algorithm is as follows:

(i) Apply μ as many times as needed to clear D.
(ii) Apply ν as many times as needed to clear D.
(iii) Choose the shortest sequence of the two.

Theorem 3. *The algorithm we just described is a $(c+1)/2$–approximation algorithm for PCCS. Furthermore, it implies the existence of a $\sqrt{n/2}$–approximation algorithm. Finally, it also implies a $(c/2+\epsilon)$–approximation algorithm when c is a constant, for any fixed $\epsilon > 0$.*

Proof. Consider an optimal clearing sequence σ of length $\text{OPT} = \text{OPT}(D, C, \psi)$. Let $A = \{1 \leq i < \text{OPT} : \sigma_i < \sigma_{i+1}\}$ and $B = \{1 \leq i < \text{OPT} : \sigma_i > \sigma_{i+1}\}$. Clearly $|A| + |B| = \text{OPT} - 1$, and therefore, either A or B has cardinality at least $\lfloor \text{OPT}/2 \rfloor$. This implies that the sequence returned by the algorithm repeats μ or ν at most $\text{OPT} - \lfloor \text{OPT}/2 \rfloor = \lceil \text{OPT}/2 \rceil$ times, and thus, it has length bounded by

$$c \left\lceil \frac{OPT}{2} \right\rceil \leq c\,\frac{\text{OPT}+1}{2} \leq \frac{c}{2}\,\text{OPT} + \frac{c}{2} \leq \frac{c+1}{2}\,\text{OPT}\,. \tag{1}$$

Here, the last inequality follows because $\text{OPT} \geq c$. This proves the first claim of the theorem.

To see that this implies the existance of a $\sqrt{n/2}$–approximation algorithm, we follow Tovey's approach. If $c \leq \sqrt{2n} - 1$, then we use the approximation algorithm above. The result follows directly from (1). Otherwise, we have $c \geq \lfloor \sqrt{2n} \rfloor$. We then determine in polynomial time whether $\text{OPT} = c$ or $\text{OPT} \geq c+1$. In the first case, we output any clearing sequence of length c. In the second case, we output any clearing sequence of length n. The sequence output by the algorithm is optimal or has a length bounded by $n \leq \sqrt{\frac{n}{2}}\,(\lfloor \sqrt{2n} \rfloor + 1) \leq \sqrt{\frac{n}{2}}(c+1) \leq \sqrt{\frac{n}{2}}\,\text{OPT}$.

Finally, note that if the number of colors is fixed a priori, we may assume that $c \leq \epsilon\,\text{OPT}$ for a given $\epsilon > 0$. Otherwise, $\text{OPT} \leq c/\epsilon = O(1)$, meaning that we can try all possible sequences using an exhaustive search procedure. Plugging $c \leq \epsilon\,\text{OPT}$ into (1), the claim follows. □

Let us note that an algorithm that applies a random permutation of $\{1, \ldots, c\}$ as many times as needed to clear D has an expected guarantee of at most $(c+1)/2$. This may be useful in an online setting, common in scheduling applications, for which only partial information is available a priori.

Finally, an interesting open question is to determine the approximability of PCCS when the number of colors is fixed. There is not even enough evidence to rule out the existence of a PTAS for this problem.

References

[1] P. Berman and T. Fujito. Approximating independent sets in degree 3 graphs. In *Proceedings of the 4th Workshop on Algorithms and Data Structures*, volume 955 of *Lecture Notes in Computer Science*, pages 449–460, 1995.

[2] C.B. Lofgren, L.F. McGinnis, and C.A. Tovey. Routing printed circuit cards through an assembly cell. *Operations Research*, 39(6):992–1004, 1991.

[3] A. Darte. On the complexity of loop fusion. *Parallel Computing*, 26(9):1175–1193, 2000.

[4] C.B. Lofgren. *Machine configuration of flexible printed circuit board assembly systems*. PhD thesis, School of ISyE, Georgia Institute of Technology, Atlanta (GA), USA, 1986.

[5] C.A. Tovey. Non-approximability of precedence-constrained sequencing to minimize setups. *Discrete Applied Mathematics*, 134:351–360, 2004.