Lehrstuhl für Operations Research
Fakultät für Mathematik
Technische Universität München

Problem Set 3

Relaxed Decision Procedures and a FPTAS

Approximation Algorithms (MA5517)

Dr. Jannik Matuschke | M.Sc. Marcus Kaiser

This problem set will be discussed in the tutorials on November 13th/14th, 2018.

Problem 3.1 (Relaxed Decision Procedure)

Consider a combinatorial minimization problem with non-negative, integer-valued objective function. For an instance $\mathcal{I}$, let $X_\mathcal{I}$ be the set of feasible solutions and $f_\mathcal{I} : X_\mathcal{I} \to \mathbb{N}$ be the objective function. Assume that the encoding length of the optimal value is polynomial in the encoding length of the instance $\mathcal{I}$.

A ρ-relaxed decision procedure is an algorithm that, given an instance $\mathcal{I}$ and a number $U \geq 0$, either finds a feasible solution $x \in X_\mathcal{I}$ with $f_\mathcal{I}(x) \leq \rho U$ or correctly states that no feasible solution $x \in X_\mathcal{I}$ with $f_\mathcal{I}(x) \leq U$ exists. Note that the output of the algorithm is not uniquely specified for $U < \min_{x \in X_\mathcal{I}} f_\mathcal{I}(x) \leq \rho U$.

Show that, for a given combinatorial minimization problem with integer-valued objective function, there exists a polynomial-time ρ-relaxed decision procedure if and only if there exists a ρ-approximation algorithm.

Problem 3.2 (Related Parallel Machine Scheduling)

In this problem, we consider a variant of the problem of scheduling on parallel machines so as to minimize the makespan, that is the maximum load of a machine. There are n jobs of lengths $p_1 \geq p_2 \geq \ldots \geq p_n > 0$ which have to be scheduled on m machines with speeds $s_1 \geq s_2 \geq \ldots \geq s_m > 0$. Processing a job $j \in [n]$ on machine $i \in [m]$ takes p_j/s_i units of time. This problem is called makespan minimization for scheduling on related machines.

Consider the following variant of the list scheduling algorithm. Given a deadline D, anytime a machine $i \in [m]$ becomes idle before time D, it processes the longest job $j \in [n]$ that has not yet been processed and satisfies $p_j/s_i \leq D$. If no such job is available or machine i becomes idle at time D or later, it stops processing. In the case that not all jobs are processed by this procedure, the algorithm states that no schedule of length D exists.

Prove that this algorithm is a polynomial-time 2-relaxed decision procedure and conclude that there is a 2-approximation algorithm.

Problem 3.3 (Unrelated Parallel Machine Scheduling)

We are given a scheduling problem with n jobs to be assigned to m machines. If job $j \in [n]$ is scheduled on machine $i \in [m]$, then it requires $p_{ij} \geq 0$ units of processing time. The goal is to find a schedule of minimum makespan. This problem is called makespan minimization for scheduling on unrelated parallel machines.

For a parameter $D \geq 0$, consider the following system of linear inequalities.

$$\begin{aligned}
\sum_{i=1}^{m} x_{ij} &= 1 && \text{for all } j \in [n] \\
\sum_{j=1}^{n} p_{ij} x_{ij} &\leq D && \text{for all } i \in [m] \\
x_{ij} &\geq 0 && \text{for all } i \in [n], j \in [m] : p_{ij} \leq D \\
x_{ij} &= 0 && \text{for all } i \in [n], j \in [m] : p_{ij} > D
\end{aligned}$$

i) Prove that there is a feasible solution to the system of linear inequalities, if the length of an optimal schedule is no greater than D.

ii)* For a feasible solution x to the system of linear inequalities, define the bipartite graph $G_x := (\mathcal{M} \cup \mathcal{J}, E_x)$ on machine vertices $\mathcal{M} := \{M_1, \ldots, M_m\}$ and job vertices $\mathcal{J} := \{J_1, \ldots, J_n\}$ with edge set

$$E_x := \big\{\{M_i, J_j\} : i \in [m], j \in [n], x_{ij} > 0\big\}.$$

Prove that if x is a basic feasible solution, each connected component of G_x has no more edges than vertices.

iii) Let x be a basic feasible solution. For all $i \in [m]$ and $j \in [n]$ with $x_{ij} = 1$, assign job j to machine i. Show that this partial assignment can be extended to a schedule by assigning at most one additional job to each machine. Argue that this results in a schedule of length at most $2D$.

iv) Show that there is a 2-approximation algorithm for makespan minimization for scheduling on unrelated parallel machines.

Problem 3.4 (Constrained Shortest Path Problem)

Suppose we are given a directed acyclic graph $G = (V, A)$ with specified source node $s \in V$ and sink node $t \in V$. Each arc $a \in A$ has an associated length $\ell_a \geq 0$ and cost $c_a \geq 0$. Further, we are given a budget $C \geq 0$. The goal is to find a shortest s-t-path with total cost no more than C. We can assume that all the numbers in the input are integers.

i) For some $L \in \mathbb{N}$, consider the expansion of G w.r.t. length $G^L := \big(V^L, A^L\big)$ defined by

$$V^L := V \times \{0, \ldots, L\} \qquad \text{and} \qquad A^L := \Big\{\Big((v, \ell), (w, \ell + \ell_a)\Big) : a = (v, w) \in A, \ell \in \{0, \ldots, L - \ell_a\}\Big\}.$$

Further, lift the costs to G^L by setting $c\big((v, \ell), (w, \ell + \ell_a)\big) := c_a$ for all $a = (v, w) \in A, \ell \in \{0, \ldots, L - \ell_a\}$.

Prove that there is a s-t-path of length L and cost at most C in G if and only if there is a $(s, 0)$-(t, L)-path of cost at most C in the graph G^L.

ii) For fixed $\varepsilon > 0$, consider the following relaxed decision procedure.

Given $L > 0$, set $\mu := \frac{L\varepsilon}{|V|}$ and $\ell_a^\mu := \big\lfloor \frac{\ell_a}{\mu} \big\rfloor$ for all $a \in A$.
Find a shortest s-t-path P in G w.r.t. ℓ^μ of total cost at most C. If P exists and $\ell^\mu(P) \leq \frac{|V|}{\varepsilon}$, return P. Otherwise, state that there is no s-t-path in G of total length at most L w.r.t. ℓ and total cost at most C.

Prove that this is a $(1 + \varepsilon)$-relaxed decision procedure for the problem of finding a shortest s-t-path w.r.t. ℓ of total cost at most C and give a fully polynomial-time approximation scheme.