Information Processing Letters 115 (2015) 570–575

Contents lists available at ScienceDirect

Information Processing Letters

www.elsevier.com/locate/ipl

Packing items into several bins facilitates approximating the separable assignment problem

Marco Bender [a,1], Clemens Thielen [b,*], Stephan Westphal [a]

[a] Clausthal University of Technology, Institute for Applied Stochastics and Operations Research, Erzstr. 1, D-38678 Clausthal-Zellerfeld, Germany

[b] University of Kaiserslautern, Department of Mathematics, Paul-Ehrlich-Str. 14, D-67663 Kaiserslautern, Germany

ARTICLE INFO

Article history:
Received 23 January 2014
Received in revised form 28 November 2014
Accepted 5 February 2015
Available online 11 February 2015
Communicated by B. Doerr

Keywords:
Separable assignment problem
Approximation algorithms
Randomized rounding

ABSTRACT

We consider a variant of the separable assignment problem (SAP). In the classic version of SAP, we are given a set of bins and a set of items to pack into the bins together with a profit $p_{i,j}$ for assigning item i to bin j. Each bin j has a separate packing constraint, i.e., only certain subsets of the items fit into bin j. The objective is to find an assignment of a subset of the items to the bins such that the packing constraints of all bins are satisfied, no item is assigned to more than one bin, and the total profit is maximized. As an important special case, this problem contains the maximum generalized assignment problem (GAP).

It is known that, given a β-approximation algorithm for the single bin subproblem (i.e., the problem of finding the most profitable packing for a single bin), it is possible to obtain a $((1-\frac{1}{e})\beta)$-approximation for SAP using randomized rounding. If the single bin subproblem admits an FPTAS, one can obtain a $(1-\frac{1}{e})$-approximation. This is best possible in the sense that there exist special cases of SAP which do not admit polynomial-time approximation algorithms with an approximation factor better than $(1-\frac{1}{e})$ unless $\mathsf{NP} \subseteq \mathsf{DTIME}(n^{\mathcal{O}(\log\log n)})$.

In this paper, we consider the case of SAP where each item may be assigned at most $k \geq 1$ times (but at most once to each bin) and present a $((1-\frac{1}{e^k})\beta)$-approximation algorithm for this case under the assumption that the single bin subproblem admits a β-approximation algorithm. If the single bin subproblem admits an FPTAS, we obtain a $(1-\frac{1}{e^k})$-approximation, which shows that, for $k \geq 2$, the problem admits approximation algorithms that beat the upper bound of $(1-\frac{1}{e})$ known for other special cases of SAP.

1. Introduction

In the *separable assignment problem* (SAP), we are given a set $\mathcal{B} = \{1, \dots, m\}$ of *bins* and a set $\mathcal{I} = \{1, \dots, n\}$ of *items* to pack into the bins. For each item i and bin j, there is a profit $p_{i,j}$ that is obtained when assigning item i to bin j. Moreover, for every bin $j \in \mathcal{B}$, there is a separate packing constraint. This means that only certain subsets of the items fit into the bin, but if a set $S \subseteq \mathcal{I}$ of items is a feasible packing for bin j, then every subset $S' \subseteq S$ is also a feasible for bin j. The objective is to find an assignment of a subset of the items to the bins such that a feasible packing is obtained for each bin, no item is assigned to more than one bin, and the total profit is maximized.

* Corresponding author. Tel.: +49 631 205 4590; fax: +49 631 205 4737.
E-mail addresses: marco.bender@tu-clausthal.de (M. Bender), thielen@mathematik.uni-kl.de (C. Thielen), stephan.westphal@tu-clausthal.de (S. Westphal).

[1] Partially supported by the German Research Foundation (DFG), grant GRK 1703/1 for the Research Training Group "Resource Efficiency in Interorganizational Networks".

http://dx.doi.org/10.1016/j.ipl.2015.02.001
0020-0190/

As an important special case, SAP contains the *maximum generalized assignment problem* (GAP), in which there is a size $s_{i,j}$ for each item i in each bin j and the feasible packings of bin j are all subsets of the items with total size at most a given bin capacity B_j.

For each bin, the *single bin subproblem* of SAP is defined as the problem of finding a feasible packing with maximum profit for the bin. Fleischer et al. [1] provided a $((1-\frac{1}{e})\beta)$-approximation for SAP under the assumption that the single bin subproblem admits a β-approximation algorithm. In particular, this yields a $(1-\frac{1}{e}-\epsilon)$-approximation whenever the single bin subproblem admits a PTAS. If the single bin subproblem admits an FPTAS, Fleischer et al. [1] showed that one can get rid of the ϵ and obtain a $(1-\frac{1}{e})$-approximation. Moreover, they showed that a special case of SAP called the *capacitated distributed caching problem* (CapDC) cannot be approximated in polynomial time with an approximation factor better than $(1-\frac{1}{e})$ unless $\mathsf{NP} \subseteq \mathsf{DTIME}(n^{\mathcal{O}(\log\log n)})$. In CapDC, the bins are cache locations and a capacity A_j is given for each cache location $j \in \mathcal{B}$. There are l different files $t \in \mathcal{T}$ that can be stored at each cache location, each with a size a_t. The items are requests for certain files, i.e., each request $i \in \mathcal{I}$ is associated with some file $t_i \in \mathcal{T}$. There is a cost $c_{i,j}$ incurred for connecting request i to cache location j and a reward R_i that is obtained when request i is connected to any cache location. Hence, the profit obtained from connecting request i to cache location j is $p_{i,j} = R_i - c_{i,j}$. A cache location j can service a set $S \subseteq \mathcal{I}$ of requests if the total size of the files associated with the requests in S is no more than the capacity A_j of the cache location, i.e., if $\sum_{t \in \{t_i | i \in S\}} a_t \leq A_j$.

In this paper, we consider the case of the separable assignment problem in which each item may be assigned at most $k \geq 1$ times (but no item may be assigned more than once to the same bin). For each fixed value of $k \geq 1$, we refer to this problem as k-SAP. This version of the separable assignment problem can be motivated, e.g., from the assignment of ads to magazines or financial assets to investors. Moreover, we study the generalization of k-SAP in which we are given a different number $k_i \geq 1$ for each item $i \in \mathcal{I}$ that specifies the maximum number of bins it may be assigned to.

It is easy to see that, in general, finding an optimal solution for an instance of k-SAP for any fixed $k \geq 1$ is no easier than finding an optimal solution for an instance of SAP: Given any instance of SAP, we simply add $k-1$ additional bins for which all subsets of the items are feasible and in which every item i yields profit one more than the maximum profit $\max_{j \in \mathcal{B}} p_{i,j}$ obtainable from item i in the original bins $\mathcal{B}$. It is then immediate that an optimal solution for the resulting instance of k-SAP will assign every item once to each of the additional bins and finding an optimal assignment for the rest of the bins is equivalent to finding an optimal solution to the given instance of SAP.

Conversely, k-SAP for any fixed $k \geq 1$ can be viewed as a special case of SAP by replacing each item by k identical copies and forbidding more than one copy of an item to be packed into the same bin. The upper bound of $(1-\frac{1}{e})$ known for CapDC, however, does *not* carry over to k-SAP for $k \geq 2$ as CapDC is not a special case of k-SAP for $k \geq 2$. The reason is that, in CapDC, one cannot enforce that at most one copy of the same item is assigned to each bin when duplicating items as above.

Even though, in general, finding an *optimal* solution for an instance of k-SAP is no easier than for SAP, intuition suggests that, for larger k, the different bins are less linked, which could facilitate finding a good *approximate* solution. We show that this intuition is correct by generalizing the methods used in [1] to obtain a polynomial-time approximation algorithm for k-SAP with approximation factor $(1-\frac{1}{e^k})\beta$ for each $k \geq 1$ under the assumption that the single bin subproblem admits a β-approximation algorithm. Whenever the single bin subproblem admits an FPTAS, our method yields a $(1-\frac{1}{e^k})$-approximation algorithm for k-SAP, which shows that k-SAP for $k \geq 2$ admits approximation algorithms that beat the upper bound of $(1-\frac{1}{e})$ that is known for CapDC. Moreover, we show that our method easily generalizes to the case that, for each item $i \in \mathcal{I}$, there is a different number $k_i \geq 1$ that specifies the maximum number of bins item i may be assigned to. Given a β-approximation for the single bin subproblem, we show that our method yields a $((1-\frac{1}{e^k})\beta)$-approximation for this case, where $k := \min_{i \in I} k_i$.

1.1. Related work

Even the special case of GAP is known to be APX-hard [2]. The best known approximation algorithm for GAP is due to Feige and Vondrák [3] and achieves an approximation factor of $(1-\frac{1}{e}+\delta)$ for a small constant $\delta > 0$. Previously, Chekuri and Khanna [2] observed that a $\frac{1}{2}$-approximation algorithm was implicit in [4] and Nutov et al. [5] presented a $(1-\frac{1}{e})$-approximation algorithm for the special case of GAP in which the profit of an item is independent of the bin it is assigned to. The special case of GAP in which both the size and the profit of each item are independent of the bin is known as the *multiple knapsack problem* and admits a PTAS [2]. For a survey on different variants of assignment problems studied in literature, we refer to [6].

For the general case of SAP, the best known approximation result is the $((1-\frac{1}{e})\beta)$-approximation obtained by Fleischer et al. [1] under the assumption that the single bin subproblem admits a β-approximation algorithm. If the single bin subproblem admits an FPTAS, their method yields a $(1-\frac{1}{e})$-approximation for SAP. A $((1-\frac{1}{e})\beta)$-approximation for SAP can also be obtained from the results of Calinescu et al. [7], who gave a $(1-\frac{1}{e})$-approximation for a general class of submodular maximization problems.

2. The integer programming formulation

We start by presenting the exponential-size integer programming formulation of k-SAP that was already used in the randomized rounding algorithm presented in [1] for the case $k=1$. For each bin $j \in \mathcal{B}$, we let $\mathcal{S}_j \subseteq 2^{\mathcal{I}}$ denote the set of feasible packings of bin j. The condition that all subsets of a feasible packing for bin j are also feasible for bin j than means that $(\mathcal{I}, \mathcal{S}_j)$ is an independence