Graph Based Vector Assignment Schemes

Leah Epstein* Tamir Tassa †

Abstract

We consider off-line vector assignment problems. The goal is to assign input vectors to machines so that a given target function is minimized. The target function usually gives some measure of the quality of distribution of input vectors among machines. We deal with an asymmetric setting where the cost functions per machine may be different for different machines, and design a PTAS for a wide class of target functions. This is done by combining a graph-based technique and a new technique of preprocessing the input vectors.

Keywords: scheduling, optimization, approximation schemes, layered graphs.

1 Introduction

A general framework for Vector Assignment Problems (VAPs) was presented in [8]. In a VAP one is given a set of vectors,

$$I = \{\mathbf{x}^i \in (\mathcal{R}^+)^d : 1 \le i \le n\} \tag{1}$$

and aims at finding an assignment of those vectors to m machines,

$$A : \{1, \ldots, n\} \to \{1, \ldots, m\} \tag{2}$$

such that the value of some target function is minimized. Typical target functions take the form

$$F(A) = f(g(\mathbf{l}^1), \ldots, g(\mathbf{l}^m)) \tag{3}$$

where

- $\mathbf{l}^k$, $1 \le k \le m$, are the corresponding load vectors on each of the machines,

$$\mathbf{l}^k = \sum_{A(i)=k} \mathbf{x}^i \quad , \quad 1 \le k \le m \; ; \tag{4}$$

- $g : (\mathcal{R}^+)^d \to \mathcal{R}^+$ is a function that evaluates the cost per machine; and
- $f : (\mathcal{R}^+)^m \to \mathcal{R}^+$ is a function that evaluates the final cost over all machines.

*School of Computer Science, The Interdisciplinary Center, Herzliya, Israel. `lea@idc.ac.il`. Research supported in part by the Israel Science Foundation, (grant no. 250/01)

†Department of Applied Mathematics, Tel-Aviv University, Ramat Aviv, Tel Aviv, Israel. `tassa@post.tau.ac.il`.

Such problems are known to be strongly NP-hard. Hence, polynomial time approximation schemes (PTAS) are sought. Such schemes produce, in polynomial time, a solution (i.e., an assignment (2)) whose cost is larger than that of an optimal solution by a factor of no more than $(1 + \text{Const} \cdot \varepsilon)$, where $\varepsilon > 0$ is an arbitrary parameter. Namely, if Φ^o is the optimal cost and $\varepsilon > 0$ is a given parameter, the scheme produces a solution A that satisfies

$$F(A) \leq (1 + \text{Const} \cdot \varepsilon) \cdot \Phi^o \;, \tag{5}$$

where the constant is independent of the input data (n, m and $\mathbf{x}^i$, $1 \leq i \leq n$) and ε.

This general framework encompasses most of the problems known in the art. In [8] we review some of the problems that are covered by this framework, *e.g.*, the makespan problem [9, 11], the ℓ_p minimization problem [1, 2], the extensible bin packing problem [5, 6] and the vector scheduling problem [4]. The reader is referred to that paper for a discussion of related problems and additional corresponding references. By considering the general setting (3), instead of the many private cases that were studied previously, and by carefully analyzing the necessary properties of the cost functions f and g, we were able to obtain a PTAS for a wide class of cost functions, using standard dynamical programming and linear programming methods. A novel idea introduced there enabled even a further generalization of our results: by replacing the small input vectors in I with other vectors of larger magnitude such that the impact on the value of the target function would be small, we were able to design a PTAS even for the case where the inner cost function g is not monotone (a similar preprocessing step was used in [1]; however, the procedure there was much simpler since it was implemented for scalars). A typical family of non-monotonic cost functions is that of the Sobolev norms; those norms take into account also an ℓ_p measure of the differences of the vector components. As described in [8], some applications necessitate such a non-monotonic measure.

However, the results of [8] strongly relied upon a symmetry assumption, i.e., the value of the target function is invariant under interchanges between the machines. As this is not always the case, it is desirable to extend the framework (3) to

$$F(A) = f(g^1(\mathbf{l}^1), \ldots, g^m(\mathbf{l}^m)) \;. \tag{6}$$

Namely, the per-machine cost evaluator, $g^k(\cdot)$, may be different for each machine (for example, it may depend on various machine parameters, such as capacity, speed or priority). Also, the outer cost function f need not be symmetric. We present here a graph-based approach that enables us to design a PTAS in the non-symmetric setting as well. Graph based techniques were used in the literature in [10, 3, 7]. All these papers, however, deal with the one-dimensional case. We generalize the method to deal with vectors.

Notation agreements. Throughout this paper we adopt the following conventions:

- Small case letters denote scalars; bold face small case letters denote vectors.
- A superscript of a vector denotes the index of the vector; a subscript of a vector indicates a component in that vector. E.g., $\mathbf{l}^k_j$ denotes the jth component of the vector $\mathbf{l}^k$.
- If $\gamma(k)$ is any expression that depends on k, then $f(\gamma(k))_{1\leq k\leq m}$ stands for $f(\gamma(1), \ldots, \gamma(m))$.
- If $\mathbf{x}$ and $\mathbf{y}$ are vectors and c is a scalar, then $\mathbf{x} \leq \mathbf{y}$ and $\mathbf{x} \leq c$ mean that the inequality holds component-wise.

2 The Cost Functions

Herein we list the assumptions that we make on the outer cost function $f(\cdot)$ and the inner cost function $g(\cdot)$.

Definition 1

1. *A function $h : (\mathcal{R}^+)^n \to \mathcal{R}^+$ is monotone if*

$$h(\mathbf{x}) \le h(\mathbf{y}) \quad \forall \mathbf{x}, \mathbf{y} \in (\mathcal{R}^+)^n \text{ such that } \mathbf{x} \le \mathbf{y} \ . \tag{7}$$

2. *The function $h : (\mathcal{R}^+)^n \to \mathcal{R}^+$ dominates the function $\tilde{h} : (\mathcal{R}^+)^n \to \mathcal{R}^+$ if there exists a constant η such that*

$$\tilde{h}(\mathbf{x}) \le \eta h(\mathbf{x}) \quad \forall \mathbf{x} \in (\mathcal{R}^+)^n \ . \tag{8}$$

3. *The function $h : (\mathcal{R}^+)^n \to \mathcal{R}^+$ is Lipschitz continuous if there exists a constant M such that*

$$|h(\mathbf{x}) - h(\mathbf{y})| \le M \|\mathbf{x} - \mathbf{y}\|_\infty \quad \forall \mathbf{x}, \mathbf{y} \in (\mathcal{R}^+)^n \ . \tag{9}$$

Assumption 1 *The function $f : (\mathcal{R}^+)^m \to \mathcal{R}^+$ is:*

1. *monotone;*
2. *linear with respect to scalar multiplications, i.e., $f(c\mathbf{x}) = cf(\mathbf{x})$ for all $c \in \mathcal{R}^+$ and $\mathbf{x} \in (\mathcal{R}^+)^m$;*
3. *dominates the max norm with a domination factor η_f that is independent of m;*
4. *Lipschitz continuous with a constant M_f that is independent of m;*
5. *recursively computable (explained below).*

Assumption 2 *The functions $g^k : (\mathcal{R}^+)^d \to \mathcal{R}^+$ are:*

1. *dominating the ℓ_∞ and the ℓ_1 norms with a domination factor η_g that does not depend on m;*
2. *Lipschitz continuous with a constant M_g that does not depend on m.*

By assuming that f is recursively computable we mean that there exists a family of functions $\psi^k(\cdot, \cdot)$, $1 \le k \le m$, such that

$$f(g^1, \ldots, g^k, 0, \ldots, 0) = \psi^k \left(f(g^1, \ldots, g^{k-1}, 0, \ldots, 0), g^k \right) \tag{10}$$

(note that $f(0, \ldots, 0) = 0$ in view of Assumption 1-2). For example, if f is a weighted ℓ_p norm on $\mathcal{R}^m$, $1 \le p \le \infty$, with weights $(w_1, \ldots, w_m)$, then ψ^k is the ℓ_p norm on $\mathcal{R}^2$ with weights $(1, w_k)$.

Next, we see what functions comply with the above assumptions. Assumption 1 dictates a quite narrow class of outer cost functions. $f = \max$ is the most prominent member of that class (luckily, in many applications this is the only relevant choice of f). Other functions f for which our results apply are the ℓ_p norms taken on the t largest values in the argument vector, where $t = \min(m_0, m)$ for some constant m_0; e.g., the sum of the largest two components. Assumption 1 is not satisfied by any of the usual ℓ_p norms for $p < \infty$ because of the conjunction of conditions 3 and 4: no matter how we rescale an ℓ_p norm, $p < \infty$, one of the parameters η_f (condition 3) or M_f (condition 4) would depend on m.

As for g^k, basically any norm on $\mathcal{R}^d$ is allowed. The most interesting choices are the ℓ_p norms and the Sobolev norms, $\|\mathbf{l}\|_{1,p} := \|\mathbf{l}\|_p + \|\mathbf{\Delta l}\|_p$ where $\mathbf{\Delta l} \in \mathcal{R}^{d-1}$ and $\mathbf{\Delta l}_j = \mathbf{l}_{j+1} - \mathbf{l}_j$, $1 \le j \le d-1$. Another natural choice is the "extensible bin" cost function, $g^k(\mathbf{l}^k) = \|\max\{\mathbf{l}^k, \mathbf{c}^k\}\|$; here $\mathbf{c}^k$ is a constant vector reflecting the parameters of the kth machine and the outer norm may be any norm. In [8] we described a problem that arises in video transmission and broadcasting, the so called *line up problem*, where the above described choices for the inner cost functions are meaningful.

It is interesting to note that the set of functions that comply with either Assumption 1 or 2 is closed under positive linear combinations. For example, if f_1 and f_2 satisfy Assumption 1, so would $c_1 f_1 + c_2 f_2$ for all $c_1, c_2 > 0$.

3 A Graph Based Scheme

3.1 Preprocessing the vectors by means of truncation

Let I be the original instance of the VAP, (1). As in [8], we start by modifying I into another problem instance $\bar{I}$ where the vectors $\bar{\mathbf{x}}^i$ are defined by

$$\bar{\mathbf{x}}^i_j = \begin{cases} \mathbf{x}^i_j & \text{if } \mathbf{x}^i_j \ge \varepsilon \|\mathbf{x}^i\|_\infty \\ 0 & \text{otherwise} \end{cases} \qquad 1 \le i \le n\ ,\ 1 \le j \le d\ . \tag{11}$$

Lemma 1 *Let A be a solution to I and let $\bar{A}$ be the corresponding solution to $\bar{I}$. Then*

$$(1 - C_1\varepsilon)F(\bar{A}) \le F(A) \le (1 + C_1\varepsilon)F(\bar{A}) \quad \textit{where} \quad C_1 = M_g\eta_g\ . \tag{12}$$

Proof. Let $\mathbf{l}^k$ and $\bar{\mathbf{l}}^k$, $1 \le k \le m$, denote the load vectors in A and $\bar{A}$ respectively. In view of (11),

$$\bar{\mathbf{l}}^k \le \mathbf{l}^k \le \bar{\mathbf{l}}^k + \varepsilon \sum_{A(i)=k} \|\mathbf{x}^i\|_\infty \tag{13}$$

Since $\|\mathbf{x}^i\|_\infty = \|\bar{\mathbf{x}}^i\|_\infty \le \|\bar{\mathbf{x}}^i\|_1$ we conclude that $\sum_{A(i)=k} \|\mathbf{x}^i\|_\infty \le \|\bar{\mathbf{l}}^k\|_1$. Recalling Assumption 2-1 we get that

$$\sum_{A(i)=k} \|\mathbf{x}^i\|_\infty \le \eta_g g^k(\bar{\mathbf{l}}^k)\ . \tag{14}$$

Therefore, by (13) and (14),

$$\bar{\mathbf{l}}^k \le \mathbf{l}^k \le \bar{\mathbf{l}}^k + \varepsilon\eta_g g^k(\bar{\mathbf{l}}^k)\ . \tag{15}$$

Next, by the uniform Lipschitz continuity of g^k we conclude that

$$(1 - C_1\varepsilon)g^k(\bar{\mathbf{l}}^k) \le g^k(\mathbf{l}^k) \le (1 + C_1\varepsilon)g^k(\bar{\mathbf{l}}^k) \qquad \text{where } C_1 = M_g\eta_g\ . \tag{16}$$

Finally, we invoke the monotonicity of f and its linear dependence on scalar multiplications to arrive at (12). $\square$

We assume henceforth that the input vectors have been subjected to the truncation procedure (11). To avoid cumbersome notations we shall keep denoting the truncated vectors by $\mathbf{x}^i$ and their collection by I.

3.2 Large and small vectors

Let Φ^o denote the optimal cost, let A^o be an optimal solution, $F(A^o) = \Phi^o$, and let $\mathbf{l}^k$, $1 \le k \le m$, be the load vectors in that solution. Then, in view of Assumption 1-3 and Assumption 2-1,

$$\mathbf{l}^k \le \eta_f \eta_g \Phi^o \quad 1 \le k \le m \, . \tag{17}$$

Consequently, we conclude that all input vectors satisfy the same bound,

$$\mathbf{x}^i \le \eta_f \eta_g \Phi^o \quad 1 \le i \le n \, . \tag{18}$$

Hence, we get the following lower bound for the optimal cost:

$$\Phi^o \ge \Phi := \frac{\max_{1 \le i \le n} \|\mathbf{x}^i\|_\infty}{\eta_f \eta_g} \, . \tag{19}$$

This lower bound induces a decomposition of the set of input vectors (1) into two subsets of large and small vectors as follows:

$$\mathcal{L} = \{\mathbf{x}^i \ : \|\mathbf{x}^i\|_\infty \ge \Phi\varepsilon^{2d+1}, \ 1 \le i \le n\} \, , \tag{20}$$

$$\mathcal{S} = \{\mathbf{x}^i \ : \|\mathbf{x}^i\|_\infty < \Phi\varepsilon^{2d+1}, \ 1 \le i \le n\} \, . \tag{21}$$

We present below a technique to replace $\mathcal{S}$ with another set of vectors $\tilde{\mathcal{S}} = \{\mathbf{z}^1, \ldots, \mathbf{z}^{\tilde{\nu}}\}$ where

$$\tilde{\nu} = |\tilde{\mathcal{S}}| \le \nu = |\mathcal{S}| \quad \text{and} \quad \|\mathbf{z}^i\|_\infty = \Phi\varepsilon^{2d+1} \ \ 1 \le i \le \tilde{\nu} \, . \tag{22}$$

In other words, all vectors in $\tilde{\mathcal{S}}$ are large.

Let $\mathbf{x} \in \mathcal{S}$. Then, in view of the truncation procedure (11),

$$\varepsilon \le \frac{\mathbf{x}_j}{\|\mathbf{x}\|_\infty} \le 1 \qquad \forall \mathbf{x}_j > 0 \, , \ 1 \le j \le d \, . \tag{23}$$

Next, we define a geometric mesh on the interval $[\varepsilon, 1]$:

$$\xi_0 = \varepsilon \ \ ; \quad \xi_i = (1+\varepsilon)\xi_{i-1} \ \ , \quad 1 \le i \le q \ \ ; \quad q := \left\lceil \frac{-\lg \varepsilon}{\lg(1+\varepsilon)} \right\rceil \, . \tag{24}$$

In view of the above, every nonzero component of $\mathbf{x}/\|\mathbf{x}\|_\infty$ lies in an interval $[\xi_{i-1}, \xi_i]$ for some $1 \le i \le q$. Next, we define

$$\mathbf{x}' = \|\mathbf{x}\|_\infty \mathcal{H}\left(\frac{\mathbf{x}}{\|\mathbf{x}\|_\infty}\right) \, , \tag{25}$$

where the operator $\mathcal{H}$ retains components that are 0 or 1 and replaces every other component by the left end point of the interval $[\xi_{i-1}, \xi_i]$ where it lies. Hence, the vector $\mathbf{x}'$ may be in one of

$$s = (q+2)^d - 1 \tag{26}$$

linear subspaces of dimension 1 in $\mathcal{R}^d$; we denote those subspaces by W^σ, $1 \le \sigma \le s$. In view of the above, we define the set

$$\mathcal{S}' = \{\mathbf{x}' \ : \ \ \mathbf{x} \in \mathcal{S}\} \, . \tag{27}$$

Next, we define for each type $1 \le \sigma \le s$

$$\mathbf{w}^\sigma = \sum \{\mathbf{x}' \ : \ \mathbf{x}' \in \mathcal{S}' \cap W^\sigma\} \qquad 1 \le \sigma \le s \; ; \tag{28}$$

namely, $\mathbf{w}^\sigma$ aggregates all vectors $\mathbf{x}'$ of type σ. We now slice this vector into large identical "slices", where each of those slices and their number are given by:

$$\tilde{\mathbf{w}}^\sigma = \frac{\mathbf{w}^\sigma}{\|\mathbf{w}^\sigma\|_\infty} \cdot \Phi\varepsilon^{2d+1} \quad \text{and} \quad \kappa_\sigma = \left\lceil \frac{\|\mathbf{w}^\sigma\|_\infty}{\Phi\varepsilon^{2d+1}} \right\rceil \; . \tag{29}$$

Finally, we define the set $\tilde{\mathcal{S}}$ as follows:

$$\tilde{\mathcal{S}} = \cup_{\sigma=1}^{s} \{\mathbf{z}^{\sigma,k} = \tilde{\mathbf{w}}^\sigma \ : \ 1 \le k \le \kappa_\sigma\} \; . \tag{30}$$

Namely, the new set $\tilde{\mathcal{S}}$ includes for each type σ the "slice"-vector $\tilde{\mathbf{w}}^\sigma$, (29), repeated κ_σ times. As implied by (29), all vectors in $\tilde{\mathcal{S}}$ have a max norm of $\Phi\varepsilon^{2d+1}$, in accord with (22). Also, the number of vectors in $\tilde{\mathcal{S}}$, $\tilde{\nu} = \sum_{\sigma=1}^{s} \kappa_\sigma$, is obviously no more than ν as the construction of the new vectors implies that $\kappa_\sigma \le |\mathcal{S}' \cap W^\sigma|$ (recall that $\|\mathbf{x}'\|_\infty < \Phi\varepsilon^{2d+1}$ for all $\mathbf{x}' \in \mathcal{S}'$).

So we have modified the original problem instance I, having n input vectors $\mathcal{L} \cup \mathcal{S}$, into an intermediate problem instance $I' = \mathcal{L} \cup \mathcal{S}'$, see (27), and then to a new problem instance,

$$\tilde{I} = \mathcal{L} \cup \tilde{\mathcal{S}} \; , \tag{31}$$

see (28)-(30), that has $\tilde{n} = n - \nu + \tilde{\nu}$ input vectors. The following theorem states that those problem instances are close in the sense that interests us (this theorem is a modified version of Theorem 3 in [8]).

Theorem 1 *For each solution $A \in \{1, \ldots, m\}^{\{1,\ldots,n\}}$ of I there exists a solution $\tilde{A} \in \{1, \ldots, m\}^{\{1,\ldots,\tilde{n}\}}$ of $\tilde{I}$ such that*

$$(1 - C_1\varepsilon) \cdot \left(F(\tilde{A}) - C_2\Phi\varepsilon\right) \le F(A) \le (1 + C_1\varepsilon) \cdot \left(F(\tilde{A}) + C_2\Phi\varepsilon\right) \; , \tag{32}$$

where C_1 is given in (12) and

$$C_2 = M_f M_g \; . \tag{33}$$

Conversely, for each solution $\tilde{A} \in \{1, \ldots, m\}^{\{1,\ldots,\tilde{n}\}}$ of $\tilde{I}$ there exists a solution $A \in \{1, \ldots, m\}^{\{1,\ldots,n\}}$ of I that satisfies (32).

Proof. Let A be a solution of I and A' be its counterpart solution of I'. Let $\mathbf{l}^k$ and $\mathbf{l}'^k$, $1 \le k \le m$, denote the load vectors in A and A', respectively. By (25), $1 \le \mathbf{l}^k / \mathbf{l}'^k \le 1 + \varepsilon$. Hence, by Assumption 2-1,

$$\|\mathbf{l}^k - \mathbf{l}'^k\|_\infty \le \varepsilon \eta_g g^k(\mathbf{l}'^k) \; . \tag{34}$$

Therefore, by the uniform Lipschitz continuity of g^k,

$$(1 - C_1\varepsilon) g^k(\mathbf{l}'^k) \le g^k(\mathbf{l}^k) \le (1 + C_1\varepsilon) g^k(\mathbf{l}'^k) \quad 1 \le k \le m \tag{35}$$

where C_1 is as in (12). Applying f on (35) and using Assumptions 1-1 and 1-2, we get that

$$(1 - C_1\varepsilon) F(A') \le F(A) \le (1 + C_1\varepsilon) F(A') \; . \tag{36}$$

Next, given a solution A' of I', we construct a solution $\tilde{A}$ of $\tilde{I}$ such that

$$F(\tilde{A}) - C_2\Phi\varepsilon \leq F(A') \leq F(\tilde{A}) + C_2\Phi\varepsilon \ , \tag{37}$$

with C_2 as in (33). Showing this will enable us to construct for any solution A of I a solution $\tilde{A}$ of $\tilde{I}$ for which, in view of (36) and (37), (32) holds. Then, in order to complete the proof, we shall show how from a given solution $\tilde{A}$ of $\tilde{I}$, we are able to construct a solution A' of I' for which (37) holds.

To this end, we fix $1 \leq \sigma \leq s$ and define for every machine k the following vector:

$$\mathbf{y}^{\sigma,k} = \sum\{\mathbf{x}'^i \ : \ \mathbf{x}'^i \in \mathcal{S}' \cap W^\sigma \ , \ A'(i) = k\} \ ; \tag{38}$$

i.e., $\mathbf{y}^{\sigma,k}$ is the sum of small vectors of type σ in I' that are assigned to the kth machine. Recalling (29), $\tilde{\mathcal{S}}$ includes the vector $\tilde{\mathbf{w}}^\sigma$ repeated κ_σ times, where

$$\kappa_\sigma = \left\lceil \sum_{k=1}^{m} \frac{\|\mathbf{y}^{\sigma,k}\|_\infty}{\Phi\varepsilon^{2d+1}} \right\rceil \ . \tag{39}$$

We may now select for each k an integer $t_{\sigma,k}$ such that

$$\left| t_{\sigma,k} - \frac{\|\mathbf{y}^{\sigma,k}\|_\infty}{\Phi\varepsilon^{2d+1}} \right| \leq 1 \tag{40}$$

and

$$\sum_{k=1}^{m} t_{\sigma,k} = \kappa_\sigma \ . \tag{41}$$

The integers $t_{\sigma,k}$ can be found in the following manner: we define

$$t_{\sigma,k}^{low} = \lfloor \|\mathbf{y}^{\sigma,k}\|_\infty / \Phi\varepsilon^{2d+1} \rfloor \quad \text{and} \quad t_{\sigma,k}^{high} = \lceil \|\mathbf{y}^{\sigma,k}\|_\infty / \Phi\varepsilon^{2d+1} \rceil \ .$$

Clearly,

$$\sum_{k=1}^{m} t_{\sigma,k}^{low} \leq \kappa_\sigma \quad \text{and} \quad \sum_{k=1}^{m} t_{\sigma,k}^{high} \geq \kappa_\sigma \ .$$

Since $t_{\sigma,k}^{high} - t_{\sigma,k}^{low} \leq 1$ for all $1 \leq k \leq m$, there exists an integer number $0 \leq x \leq m$ such that

$$\sum_{k=1}^{m} t_{\sigma,k}^{low} = \kappa_\sigma - x \ .$$

Finally, we set

$$t_{\sigma,k} = \begin{cases} t_{\sigma,k}^{high} & 1 \leq k \leq x \\ & \\ t_{\sigma,k}^{low} & x < k \leq m \end{cases} \ .$$

With this, the solution $\tilde{A}$ is that which assigns to the kth machine, $1 \leq k \leq m$, $t_{\sigma,k}$ vectors $\tilde{\mathbf{w}}^\sigma$ for all $1 \leq \sigma \leq s$ (and coincides with A' for all large vectors in $\mathcal{L}$). In view of (40) and the definition of $\tilde{\mathbf{w}}^\sigma$, see (29),

$$\|t_{\sigma,k} \cdot \tilde{\mathbf{w}}^\sigma - \mathbf{y}^{\sigma,k}\|_\infty \leq \Phi\varepsilon^{2d+1} \ . \tag{42}$$

Therefore, summing (42) over $1 \le \sigma \le s$, we conclude that $\tilde{\mathbf{l}}^k$ and $\mathbf{l}'^k$ – the loads on the kth machine in $\tilde{A}$ and A' respectively – are close,

$$\|\tilde{\mathbf{l}}^k - \mathbf{l}'^k\|_\infty \le s\Phi\varepsilon^{2d+1} . \tag{43}$$

However, as (24) and (26) imply that

$$s \le \varepsilon^{-2d} \qquad \text{for all } \ 0 < \varepsilon \le 1 \tag{44}$$

we conclude by (43) and (44) that

$$\|\tilde{\mathbf{l}}^k - \mathbf{l}'^k\|_\infty \le \Phi\varepsilon . \tag{45}$$

Finally, the Lipschitz continuity of both g and f imply that (37) holds with C_2 as in (33).

Next, we show how to construct from a solution $\tilde{A}$ of $\tilde{I}$, a solution A' of I' for which (37) holds. The two assignments will coincide for the large vectors $\mathcal{L}$. As for the small vectors, let us fix one vector type $1 \le \sigma \le s$, where s is the number of types (26). $\tilde{\mathcal{S}}$ includes the vector $\tilde{\mathbf{w}}^\sigma$ repeated κ_σ times, (29)-(30). Let $t_{\sigma,k}$ be the number of those vectors that $\tilde{A}$ assigns to the kth machine. The counters $t_{\sigma,k}$ satisfy (41). We now assign the vectors $\mathbf{x}' \in \mathcal{S}' \cap W^\sigma$, see (28), to the m machines so that the ℓ_∞-norm of their sum in the kth machine is greater than $(t_{\sigma,k}-1)\Phi\varepsilon^{2d+1}$ but no more than $(t_{\sigma,k}+1)\Phi\varepsilon^{2d+1}$. In view of (28) and (29), it is easy to see that such an assignment exists: Assign the jobs one by one greedily, in order to obtain in the kth machine, $1 \le k \le m$, a load with an ℓ_∞-norm of at least $(t_{\sigma,k}-1)\Phi\varepsilon^{2d+1}$. Since the ℓ_∞-norm of the sum of all small jobs is at least $(\sum_{k=1}^m t_{\sigma,k} - 1)\Phi\varepsilon^{2d+1}$, this goal can be achieved. Also, as the size of each of those jobs is no more than $\Phi\varepsilon^{2d+1}$, we may perform this assignment in a manner that keeps the load in each machine below $t_{\sigma,k}\Phi\varepsilon^{2d+1}$. After achieving that goal in all machines, we assign the remaining jobs so that the total load in each machine is bounded by $(t_{\sigma,k}+1)\Phi\varepsilon^{2d+1}$. This is possible given the small size of the jobs and the size of their sum (at most $(\sum_{k=1}^m t_{\sigma,k})\Phi\varepsilon^{2d+1}$). Clearly, if we let $\mathbf{y}^{\sigma,k}$ denote the sum of vectors $\mathbf{x}'$ of type σ thus assigned to the kth machine, then $\mathbf{y}^{\sigma,k}$ satisfies (42). As we saw before, this implies that $\tilde{A}$ and A' satisfy (37). This completes the proof.□

3.3 The scheme

In view of the previous two subsections, we assume that the original set of input vectors I was subjected to the truncation procedure, along the lines of §3.1, and then modified into a problem instance $\tilde{I}$ where all vectors are large, using the procedure described in §3.2. For convenience, we shall keep denoting the number of input vectors in $\tilde{I}$ by n and the input vectors by $\mathbf{x}^i$, $1 \le i \le n$. Hence, all vectors in $\tilde{I}$ satisfy

$$\|\mathbf{x}^i\|_\infty \ge \Phi\varepsilon^{2d+1} \qquad 1 \le i \le n .$$

This, together with (11) on one hand and (19) on the other hand, yield the following lower and upper bounds:

$$\varepsilon^{2d+2} \le \frac{x_j^i}{\Phi} \le \eta_f\eta_g \qquad \text{for } \ 1 \le i \le n, \ 1 \le j \le d \ \text{ and } \ x_j^i \ne 0 . \tag{46}$$

Next, we define a geometric mesh on the interval given in (46):

$$\xi_0 = \varepsilon^{2d+2} \ \ ; \ \ \xi_i = (1+\varepsilon)\xi_{i-1} \ , \ \ 1 \le i \le q \ \ ; \ \ q := \left\lceil \frac{\lg(\eta_f\eta_g\varepsilon^{-2(d+1)})}{\lg(1+\varepsilon)} \right\rceil . \tag{47}$$

In view of the above, every nonzero component of $\mathbf{x}^i/\Phi$, $1 \le i \le n$, lies in an interval $[\xi_{i-1}, \xi_i]$ for some $1 \le i \le q$. We use this in order to define a new set of vectors,

$$\hat{I} = \left\{ \hat{\mathbf{x}}^i = \Phi \mathcal{H} \left(\frac{\mathbf{x}^i}{\Phi} \right) : \mathbf{x}^i \in \tilde{I} \right\} , \tag{48}$$

where the operator $\mathcal{H}$ replaces each nonzero component in the vector on which it operates by the left end point of the interval $[\xi_{i-1}, \xi_i]$ where it lies.

Theorem 2 *Let $\tilde{A}$ be a solution of $\tilde{I}$ and let $\hat{A}$ be the corresponding solution of $\hat{I}$. Then*

$$(1 - C_1\varepsilon) F(\hat{A}) \le F(\tilde{A}) \le (1 + C_1\varepsilon) F(\hat{A}) , \tag{49}$$

where C_1 is given in (12).

Proof. Let $\tilde{\mathbf{l}}^k$ and $\hat{\mathbf{l}}^k$, $1 \le k \le m$, be the load vectors in $\tilde{A}$ and $\hat{A}$. Then

$$\hat{\mathbf{l}}^k \le \tilde{\mathbf{l}}^k \le (1+\varepsilon)\hat{\mathbf{l}}^k \qquad 1 \le k \le m .$$

Hence, $\|\tilde{\mathbf{l}}^k - \hat{\mathbf{l}}^k\|_\infty \le \varepsilon \|\hat{\mathbf{l}}^k\|_\infty$; using Assumptions 2-2 and 2-1 we get that

$$|g^k(\tilde{\mathbf{l}}^k) - g^k(\hat{\mathbf{l}}^k)| \le \varepsilon M_g \eta_g g^k(\hat{\mathbf{l}}^k) \qquad 1 \le k \le m ,$$

or, equivalently,

$$(1 - C_1\varepsilon) g^k(\hat{\mathbf{l}}^k) \le g^k(\tilde{\mathbf{l}}^k) \le (1 + C_1\varepsilon) g^k(\hat{\mathbf{l}}^k) \quad 1 \le k \le m \quad \text{where} \quad C_1 = M_g \eta_g .$$

These inequalities, together with the monotonicity of f and its linearity with respect to scalar multiplications (Assumptions 1-1 and 1-2) imply (49). □

The vectors in $\hat{I}$ belong to the set

$$W = \mathcal{X}^d \quad \text{where} \;\; \mathcal{X} = \{0, \xi_0, \ldots, \xi_{q-1}\} . \tag{50}$$

As the size of W is $s = (q+1)^d$, it may be ordered:

$$W = \{\mathbf{w}^1, \ldots, \mathbf{w}^s\} . \tag{51}$$

With this, the set of modified vectors $\hat{I}$ may be identified by a configuration vector

$$\mathbf{z} = (\mathbf{z}_1, \ldots, \mathbf{z}_s) \quad \text{where} \quad \mathbf{z}_i = \#\{\hat{\mathbf{x}} \in \hat{I} \;:\; \hat{\mathbf{x}} = \mathbf{w}^i\} \;,\; 1 \le i \le s . \tag{52}$$

Next, we may describe all possible assignments of vectors from $\hat{I}$ to the m machines using a layered graph $G = (V, E)$. To that end, assume that $\hat{A} : \hat{I} \to \{1, \ldots, m\}$ is such an assignment. We let $\hat{I}^k$ denote the subset of $\hat{I}$ consisting of those vectors that were assigned to one of the first k machines,

$$\hat{I}^k = \{\hat{\mathbf{x}} \in \hat{I} \;:\; \hat{A}(\hat{\mathbf{x}}) \le k\} \qquad 1 \le k \le m .$$

Furthermore, we define the corresponding *state vector*

$$\mathbf{z}^k = (\mathbf{z}^k_1, \ldots, \mathbf{z}^k_s) \quad 1 \le k \le m \quad \text{where} \quad \mathbf{z}^k_i = \#\{\hat{\mathbf{x}} \in \hat{I}^k \;:\; \hat{\mathbf{x}} = \mathbf{w}^i\} \;,\; 1 \le i \le s .$$

We note that

$$\emptyset = \hat{I}^0 \subseteq \hat{I}^1 \subseteq \ldots \subseteq \hat{I}^{m-1} \subseteq \hat{I}^m = \hat{I} \tag{53}$$

and

$$\mathbf{0} = \mathbf{z}^0 \leq \mathbf{z}^1 \leq \ldots \leq \mathbf{z}^{m-1} \leq \mathbf{z}^m = \mathbf{z} \, , \tag{54}$$

where $\mathbf{z}$ is given in (52). In addition, when $0 < k < m$, $\hat{I}^k$ may be any subset of $\hat{I}$ while $\mathbf{z}^k$ may be any vector in

$$Z = \{\mathbf{y} \, : \, \mathbf{0} \leq \mathbf{y} \leq \mathbf{z}\} \, . \tag{55}$$

With this, we define the graph $G = (V, E)$ as follows:

- The set of vertices consists of $m + 1$ layers, $V = \cup_{k=0}^{m} V^k$. If $v \in V$ is a vertex in the kth layer, V^k, then it represents one of the possible state vectors after assigning vectors to the first k machines. Hence $V^0 = \{\mathbf{0}\}$, $V^m = \{\mathbf{z}\}$ and the intermediate layers are $V^k = Z$, see (55), $0 < k < m$.

- The set of edges consists of m subsets:

$$E = \cup_{k=1}^{m} E^k \quad \text{where} \quad E^k = \{(\mathbf{u}, \mathbf{v}) \; : \; \mathbf{u} \in V^{k-1} \, , \, \mathbf{v} \in V^k \, , \, \mathbf{u} \leq \mathbf{v}\} \, . \tag{56}$$

  In other words, there is an edge connecting two vertices in adjacent layers, $\mathbf{u} \in V^{k-1}$ and $\mathbf{v} \in V^k$, if and only if there exists an assignment to the kth machine that would change the state vector from $\mathbf{u}$ to $\mathbf{v}$.

Note that all intermediate layers, V^k, $0 < k < m$, are composed of the same number of vertices, t, given by the number of sub-vectors that $\mathbf{z}$ has:

$$t = |Z| = \prod_{i=1}^{s} (\mathbf{z}_i + 1) \leq (n + 1)^s \, . \tag{57}$$

Next, we turn the graph into a weighted graph, using a weight function $w : E \to \mathcal{R}^+$ that computes the cost that the given edge implies on the corresponding machine: Let $e = (\mathbf{u}, \mathbf{v}) \in E^k$. Then the difference $\mathbf{v} - \mathbf{u}$ tells us how many vectors of each of the s types are assigned by this edge to the kth machine. The weight of this edge is therefore defined as

$$w(e) = g^k(T(\mathbf{v} - \mathbf{u})) \quad \text{where} \quad T(\mathbf{v} - \mathbf{u}) = \sum_{i=1}^{s} (\mathbf{v}_i - \mathbf{u}_i)\mathbf{w}^i \, , \tag{58}$$

$\mathbf{w}^i$ are as in (51). We continue to define a cost function on the vertices, $r : V \to \mathcal{R}^+$. The cost function is defined recursively according to the layer of the vertex, using Assumption 1-5:

$$r(v) = 0 \quad , \quad v \in V^0 \, ; \tag{59}$$

$$r(v) = \min \left\{ \psi^k(r(u), w(e)) \; : \; u \in V^{k-1}, \; e = (u, v) \in E^k \right\} \quad , \quad v \in V^k \tag{60}$$

(the functions ψ^k are as in (10)). This cost function coincides with the cost function of the VAP, (6). More specifically, if $v \in V^k$ and it represents a subset of vectors $\hat{I}^k \subseteq \hat{I}$, then $r(v)$ equals the value of an optimal assignment of the vectors in $\hat{I}^k$ to the first k machines. Hence, the cost of the end vertex, $r(v)$, $v \in V^m$, equals the value of an optimal solution of the VAP for $\hat{I}$.

The goal is to find the shortest path from V^0 to V^m that achieves this minimal cost. Namely, we look for a sequence of vertices $v^k \in V^k$, $0 \le k \le m$, such that

$$e^k := (v^{k-1}, v^k) \in E^k \qquad 1 \le k \le m \tag{61}$$

and

$$f(w(e^1), \ldots, w(e^m)) = r(v^m) \ . \tag{62}$$

We may apply a standard algorithm to find this minimal path within $\mathcal{O}(|V| + |E|)$ steps. As

$$|V| \le 2 + (m-1) \cdot (n+1)^s \tag{63}$$

and

$$|E| = \sum_{k=1}^{m} |E^k| \le m \cdot (n+1)^{2s} \tag{64}$$

the running time would be polynomial in n and m.

The shortest path thus found represents an assignment of the vectors of the modified set $\hat{I}$,

$$\hat{A} : \hat{I} = \{\hat{\mathbf{x}}_1, \ldots, \hat{\mathbf{x}}_n\} \to \{1, \ldots, m\} \ . \tag{65}$$

We need to translate this assignment into an assignment of the original vectors,

$$A : I = \{\mathbf{x}_1, \ldots, \mathbf{x}_n\} \to \{1, \ldots, m\} \ . \tag{66}$$

To that end, let us review all the problem modifications that we performed:

- First modification: I to $\bar{I}$, see (11) and Lemma 1.
- Second modification: $\bar{I}$ to $\tilde{I}$, see (28)-(31) and Theorem 1.
- Third modification: $\tilde{I}$ to $\hat{I}$, see (48) and Theorem 2.

In view of the above, we translate the solution that we found, $\hat{A}$, into a solution $\tilde{A}$ of $\tilde{I}$, then – along the lines of Theorem 1 – we translate it into a solution $\bar{A}$ of $\bar{I}$ and finally we take the corresponding solution A of I.

Theorem 3 *Let Φ^o be the optimal cost of the original problem instance I. Let A be the solution of I that is obtained using the above scheme. Then A satisfies (5) with a constant that depends only on η_g, M_g and M_f.*

Proof. Let $\hat{\Phi}^o$ denote the optimal cost of $\hat{I}$. Then, using the left inequalities in Lemma 1, Theorem 1 and Theorem 2, together with (19), we conclude that

$$\hat{\Phi}^o \le \frac{1}{1 - C_1\varepsilon} \cdot \left(\frac{1}{(1 - C_1\varepsilon)^2} + C_2\varepsilon \right) \Phi^o \le (1 + T\varepsilon)\Phi^o \ , \tag{67}$$

for an appropriate choice of T that depends only on C_1 and C_2. However, $\hat{\Phi}^o$ is no other than the cost of the shortest path that we found in the graph, namely, the cost of the solution $\hat{A}$ that we found for $\hat{I}$. As A is the solution of I that is obtained from $\hat{A}$, we may upper bound its cost using the right inequalities in Lemma 1, Theorem 1 and Theorem 2:

$$F(A) \leq (1 + C_1\varepsilon)^2 \cdot \left((1 + C_1\varepsilon)\hat{\Phi}^o + C_2\Phi^o\varepsilon\right) \ . \tag{68}$$

The Inequalities (67) and (68) imply that

$$F(A) \leq (1 + C_1\varepsilon)^2 \cdot ((1 + C_1\varepsilon)(1 + T\varepsilon) + C_2\varepsilon)\,\Phi^o \leq (1 + \text{Const} \cdot \varepsilon)\Phi^o \ ,$$

where the constant depends only on $C_1 = M_g\eta_g$ and $C_2 = M_f M_g$. $\square$

References

[1] N. Alon, Y. Azar, G. Woeginger, and T. Yadid. Approximation schemes for scheduling. In *Proc. 8th ACM-SIAM Symp. on Discrete Algorithms (SODA'97)*, pages 493–500, 1997.

[2] N. Alon, Y. Azar, G. Woeginger, and T. Yadid. Approximation schemes for scheduling on parallel machines. *Journal of Scheduling*, 1:1:55–66, 1998.

[3] Y. Azar and L. Epstein. Approximation schemes for covering and scheduling on related machines. In *1st Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX98)*, pages 39–47, 1998.

[4] C. Chekuri and S. Khanna. On multi-dimensional packing problems. In *Proceedings of the Tenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'99)*, pages 185–194, 1999.

[5] E. G. Coffman, Jr. and G. S. Lueker. Approximation algorithms for extensible bin packing. In *Proceedings of the Twelfth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'01)*, pages 586–588, 2001.

[6] P. Dell'Olmo, H. Kellerer, M. G. Speranza, and Zs. Tuza. A $13/12$ approximation algorithm for bin packing with extendable bins. *Information Processing Letters*, 65(5):229–233, 1998.

[7] L. Epstein and J. Sgall. Approximation schemes for scheduling on uniformly related and identical parallel machines. In *7th Annual European Symposium on Algorithms (ESA'99)*, pages 151–162, 1999.

[8] L. Epstein and T. Tassa. Vector assignment problems: A general framework. manuscript, 2002.

[9] R. L. Graham. Bounds for certain multiprocessor anomalies. *Bell System Technical Journal*, 45:1563–1581, 1966.

[10] D. Hochbaum and D. Shmoys. A polynomial approximation scheme for scheduling on uniform processors: using the dual approximation approach. *SIAM Journal on Computing*, 17(3):539–551, 1988.

[11] D. S. Hochbaum and D. B. Shmoys. Using dual approximation algorithms for scheduling problems: theoretical and practical results. *Journal of the ACM*, 34(1):144–162, 1987.