Probabilistic analysis of the generalised assignment problem

Martin Dyer*

School of Computer Studies, University of Leeds, Leeds LS2 9YT, UK

Alan Frieze**

Department of Mathematics, Carnegie-Mellon University, Pittsburgh, PA 15213, USA

We analyse the *generalised assignment problem* under the assumption that all coefficients are drawn uniformly and independently from $[0, 1]$. We show that such problems can be solved exactly with high probability, in a well-defined sense. The results are closely related to earlier work of Lueker, Goldberg and Marchetti-Spaccamela and ourselves.

Key words: Probabilistic analysis, generalised assignment.

1. Introduction

We are concerned here with the *generalised assignment problem*:

$$\text{Maximise} \quad z = \sum_{i=1}^{m} \sum_{j=1}^{n} c_{ij} x_{ij}$$

$$\text{subject to} \quad \sum_{j=1}^{n} a_{ij} x_{ij} \leq b_i = \beta_i n, \quad i = 1, 2, \ldots, m,$$

$$\sum_{i=1}^{m} x_{ij} = 1, \quad j = 1, 2, \ldots, n,$$

$$\text{and all} \quad x_{ij} \in \{0, 1\}. \tag{1}$$

This problem has applications in machine scheduling and elsewhere. See, for example [1, 5, 8]. Here we give a *probabilistic analysis* of (1), under the assumption that all a_{ij}, c_{ij} are independently and uniformly distributed in $[0, 1]$. We examine the asymptotic behaviour as $n \to \infty$, while m and the β_i remain fixed. (The assumption that m is constant while $n \to \infty$ is realistic, for example, in the machine-scheduling application.) We assume $m > 1$, since otherwise the problem is obviously trivial.

We have chosen the maximising form for (1). All our results apply equally well to the minimising form, since the constraints imply that the objective function

$$z = n - \sum_{i=1}^{m} \sum_{j=1}^{n} (1 - c_{ij})x_{ij}$$

and $(1 - c_{ij})$ has precisely the same distribution as c_{ij}.

The work presented here is closely related to, and in part rests on, our earlier paper [2] on the multidimensional knapsack problem. (This was itself an extension of work of Lueker [6] and Goldberg and Marchetti-Spaccamela [3].) As in that paper, the main technique here is to relate the solution of the integer program (1) to its linear programming (LP) relaxation:

(1) with $x_{ij} \in \{0, 1\}$ replaced by $x_{ij} \geqslant 0$.   (2)

As in [2], we use strongly the fact that the only conditioning on the data imposed by the solution of (2) arises from satisfying the primal and dual feasibility conditions (which are simply linear inequalities). We make use of well-known ideas from LP theory throughout the paper without comment. We use similarly standard methods from probability theory, in particular certain probability inequalities. See [7] for a useful survey of many such inequalities.

However, although our methods are similar to those of [2], extending a probabilistic analysis to an apparently similar problem is not straightforward, since changed structure produces different conditioning. Nevertheless, we will show that in this case the difficulties can be successfully negotiated. There is one proviso. For technical reasons, we must assume the following condition on the β_i holds. Consider the vector β as a point in $\mathbb{R}^m$. Let the set $\mathcal{F}$ be (the convex set of) those β for which (2) remains feasible with probability tending to 1 as $n \to \infty$. We clearly need $\beta \in \mathcal{F}$ in order that the feasibility of (1) does not dominate the analysis. We assume something stronger, however. Let $\mathcal{F}(\eta)$ be the set of β such that (2) remains feasible (with probability tending to 1) when the β_i are (all) reduced to $\beta_i - \eta$, for given $\eta > 0$. If $\beta \in \mathcal{F}(\eta)$, we will call (1) (or (2)) η-*feasible*. We require η-feasibility for some $\eta > 0$. It is clear that int $\mathcal{F} = \bigcup_{\eta > 0} \mathcal{F}(\eta)$. Thus, in a certain sense, our analysis covers "almost all" β's of interest.

From an algorithmic viewpoint, we do not require conditions for η-feasibility (for some $\eta > 0$). However, various simple sufficient conditions can be derived. For example,

$$\eta = \min_{1 \leqslant i \leqslant m} \beta_i - \frac{1}{m(m+1)} > 0,$$

is such a condition. To see this, consider the solution to (2) given by setting $x_{ij} = 1$ in column j for the i which gives the minimum a_{ij} in that column. Since this minimum is $1/(m+1)$ in expectation (the minimum of m uniform $[0, 1]$'s), and will occur in each row with probability $1/m$, the sufficiency of the conditions follows easily from the law of large numbers.

The results we prove are specifically as follows. The first relates the objective values of (1) and (2). (Natural logarithms are used throughout.)

Theorem 1. *Let* (1) *be* η-*feasible for some* $\eta > 0$, *and* z_{IP}, z_{LP} *be the optimal values for* (1), (2) *respectively. Then there exist constants* $0 < p(m) < 1$ *and* $\alpha(m, \eta) > 0$ *such that, for large n,*

$$\Pr(z_{LP} - z_{IP} > t\alpha(\log n)^2/n) < p', \quad t = 1, 2, \ldots, \lfloor n/(\log n)^2 \rfloor,$$

and

$$E(z_{LP} - z_{IP}) < (2\alpha/(1 - p)^2)(\log n)/n.$$

Moreover, if $\mathcal{N}$ is the umber of variables whose values differ in the two optimal solutions, then there exists $\sigma(m, \eta)$ such that

$$E(\mathcal{N}) < \sigma \log n.$$

The second shows that, with any specified probability of error, (1) can be solved in polynomial time.

Theorem 2. *Let* (1) *be* η-*feasible for some* $\eta > 0$, *and* $\tau > 0$ *a probability. Then there exists a constant* $d(m, \eta, \tau)$, *such that a proven optimal solution can be obtained by a branch-and-bound search algorithm in time* $O(n^d)$ *with probability at least* $(1 - \tau)$.

(Unfortunately, our estimate for d grows quite rapidly as $\tau \to 0$.)

The plan of the paper is as follows. In Section 2 we give some preliminary results on the linear programming (LP) relaxation of (1). In Section 3 we prove Theorem 1 and in Section 4 we prove Theorem 2. Finally Section 6 gives some brief concluding remarks.

2. Linear programming preliminaries

We show here some simple properties of the LP problem (2) and its dual:

$$\text{Minimise} \quad \sum_{i=1}^{m} b_i u_i + \sum_{j=1}^{n} v_j$$

$$\text{subject to} \quad a_{ij} u_i + v_j \geq c_{ij} \quad \forall i, j, \tag{3}$$

$$u_i \geq 0 \quad \forall i.$$

Lemma 3 below has been given by Benders and van Nunen [1], with a different proof and a slightly weaker conclusion.

Lemma 3. *Let $\{x_{ij}^{B}\}$ be any basic feasible solution to the system*

$$\text{subject to} \quad \sum_{j=1}^{n} a_{ij} x_{ij} \leqslant b_i, \quad i = 1, 2, \ldots, m,$$

$$\sum_{i=1}^{m} x_{ij} = 1, \quad j = 1, 2, \ldots, n, \tag{4}$$

$$\text{and all} \quad x_{ij} \geqslant 0.$$

and let $F^{B} = \{(i, j): 0 < x_{ij}^{B} < 1\}$, $E^{B} = \{j: \exists (i, j) \in F^{B}\}$. Then (a) $|E^{B}| \leqslant m$, *and* (b) $|F^{B}| \leqslant 2m$.

Proof. Any basic feasible solution to (4) is an optimal solution for some objective function

$$\text{Maximise} \quad z = \sum_{i=1}^{n} \sum_{j=1}^{n} \hat{c}_{ij} x_{ij}.$$

In such a solution, for each j, there is an $i^*(j)$ such that $x_{i^*j} > 0$. We use the constraints $\sum_{i=1}^{m} x_{ij} = 1$ to eliminate the x_{i^*j}. This will give an LP:

$$\text{Maximise} \quad z = \sum_{j=1}^{n} \sum_{i \neq i^*} c'_{ij} x_{ij}$$

$$\text{subject to} \quad \sum_{j=1}^{n} \sum_{i \neq i^*} a'_{ijk} x_{ij} \leqslant b'_k, \quad k = 1, 2, \ldots, m, \tag{5}$$

$$\sum_{i \neq i^*} x_{ij} \leqslant 1, \quad j = 1, 2, \ldots, n,$$

$$\text{and all} \quad x_{ij} \geqslant 0.$$

Now in the LP (5), all constraints $\sum_{i \neq i^*} x_{ij} \leqslant 1$ are non-binding at the optimum, and hence redundant. Thus the solution will be the same as that to

$$\text{Maximise} \quad z = \sum_{j=1}^{n} \sum_{i \neq i^*} c'_{ij} x_{ij}$$

$$\text{subject to} \quad \sum_{j=1}^{n} \sum_{i \neq i^*} a'_{ijk} x_{ij} \leqslant b'_k, \quad k = 1, 2, \ldots, m, \tag{6}$$

$$\text{and all} \quad x_{ij} \geqslant 0.$$

Let $F = \{(i, j): x_{ij} > 0$ in the optimal solution to (5)$\}$ and $E = \{j: \exists (i, j) \in F\}$. Clearly $|E| \leqslant |F|$, and since there is an optimal basic solution, $|F| \leqslant m$. Now $x_{i^*j} < 1$ if and only if $j \in E$. Thus $E^{B} = E$, proving (a). Also $F^{B} = F \cup \{(i^*, j): j \in E\}$, and since the union is disjoint $|F^{B}| = |F| + |E| \leqslant 2m$, proving (b). $\square$

Lemma 4. *Any basic feasible solution to* (3) *can be completely specified by giving* $E^B \subseteq \{1, 2, \ldots, n\}$ *with* $|E^B| \leq m$, *and* $F^B \subseteq \{(i, j): j \in E^B\}$ *with* $|F^B| \leq 2m$.

Proof. Consider the dual (3). We assume, without loss of generality, that it is non-degenerate. If the set F^B is as defined in Lemma 3, then the set of equations

$$a_{ij}u_i + v_j = c_j, \quad (i, j) \in F^B,$$

will uniquely determine the non-zero u_i. The remaining v_j can now be determined by $v_j = \max_i\{c_{ij} - a_{ij}u_i\}$. This completely specifies the dual basic feasible solution, and the primal solution can be recovered by complementary slackness. $\square$

The method of specifying a dual basic feasible solution implied by Lemma 4 is central to our arguments below. It has the following important consequence.

Corollary 5. (*Assuming non-degeneracy*) *there are less than* $(6m)^m n^m$ *basic feasible solutions to* (3).

Proof. There are clearly at most

$$\binom{n}{m}\binom{m^2}{2m} < (en/m)^m (em^2/2m)^{2m} = (e^3 mn/4)^m < (6mn)^m$$

ways of selecting E^B then F^B. $\square$

The non-degeneracy assumption holds with probability 1 in our model. Thus, as $n \to \infty$ with m fixed, we will have $O(n^m)$ (i.e., a polynomial number of) dual basic feasible solutions to consider. This should be contrasted with the m^n (i.e., an exponential number of) potential primal basic feasible solutions.

We now show that the η-feasibility condition discussed in Section 1 implies a bound on the optimal dual variables in (2). For ease of notation let us write (2) in "matrix form", i.e.,

LP(b): $z_{LP}(b) = \max z$

subject to $z = cx, \ Ax \leq b, \ Ex \leq e_n, \ x \geq 0$,

where in particular, e_n denotes an n-vector of 1's.

Lemma 6. *Let* $\eta > 0$. *If* (2) *is* η-*feasible and* u_i, $i = 1, \ldots, m$, *are the dual variables in its optimal solution, then*

$$\sum_{i=1}^{n} u_i \leq 1/\eta.$$

Proof. Let x be any feasible solution to $\text{LP}(b - n\eta e_m)$. Clearly $cx \geq 0$. The optimal dual variables for $\text{LP}(b)$ satisfy

$$uA + vE \geq c, \quad u \geq 0, \quad ub + ve_n \leq n, \tag{7}$$

the last inequality following from the duality theorem and the obvious upper bound on the primal objective function. Hence

$$n \geq ub + ve_n$$
$$\geq u(Ax + \eta n e_m) + ve_n$$
$$= uAx + \eta nue_m + vEx$$
$$= (uA + vE)x + \eta nue_m$$
$$\geq cx + \eta nue_m$$
$$\geq \eta nue_m.$$

Thus $ue_m \leq 1/\eta$, as required. $\square$

The last result in this section relates the optimal solution of (2) to "nearly optimal" solutions when the right hand sides are changed.

Lemma 7. *Let y' be an optimum basic feasible solution to $\text{LP}(b')$ and $x = y' + \theta$ any feasible solution to $\text{LP}(b)$. Let $u \geq 0$ be the optimum dual variables for the inequalities in $\text{LP}(b')$ and $\bar{c} \geq 0$ the reduced costs associated with the nonbasic variables in y'. Then*

$$cx \leq z_{\text{LP}}(b) + \bar{c}\theta + u(b - Ax).$$

Proof. From the objective row of the basis associated with y' we have

$$cx = \bar{c}x + u(b' - Ax) + cy'.$$

Let $y \geq 0$ solve $\text{LP}(b)$, so

$$z_{\text{LP}}(b) = cy = \bar{c}y + u(b' - Ay) + cy'.$$

Thus

$$cx - cy = \bar{c}(y' + \theta) - \bar{c}y + u(Ay - Ax)$$
$$= \bar{c}\theta - \bar{c}y + u(Ay - Ax) \qquad (\bar{c}y' = 0)$$
$$\leq \bar{c}\theta + u(Ay - Ax) \qquad (\bar{c} \geq 0, y \geq 0)$$
$$\leq \bar{c}\theta + u(b - Ax) \qquad (Ay \leq b, u \geq 0). \qquad \square$$

3. Proof of Theorem 1

Assume $1 > \eta > 0$ is given, and that (1) is η-feasible. Let $K = \frac{9}{20}m^2\eta \log n$, $k = \lfloor 18m^2 \log n \rfloor$, $M = 2^{m+9}/\eta^2$ and $\varepsilon = 80 \, m^3 M \log n / n$.

Now, for $i = 1, \ldots, m$, let $b_i' = b_i - K$. Let x^B solve LP(b'). From Lemma 6 the dual variables u_i for LP(b') satisfy

$$\sum_{i=1}^{m} u_i \leqslant \frac{1}{\eta - K/n} < 2/\eta,$$

for all n large enough. So $u_i < 2/\eta$, for all i.

Let E^B, F^B be as defined in Section 2. For each $j = 1, \ldots, n$ define $i^*(j)$ as follows. For $j \notin E^B$, let $i^*(j)$ be the (unique) i such that $x_{ij}^B = 1$. For $j \in E^B$, choose $i^*(j)$ arbitrarily from $\{i: (i,j) \in F^B\}$. We define a "rounded" zero-one solution ξ^B by

$$\xi_{ij}^B = 1 \quad \text{if and only if} \quad (i,j) = (i^*, j).$$

Now let

$$\Delta_i^B = \begin{cases} b_i - \sum\limits_{j=1}^{n} a_{ij} \xi_{ij}^B, & u_i \neq 0, \\ K, & \text{otherwise.} \end{cases}$$

Then, since $0 \leqslant a_{ij} \leqslant 1$,

$$K - m \leqslant \Delta_i^B \leqslant K + m. \tag{8}$$

We now wish to "fill the deficits" Δ_i^B by changing only the values (in ξ^B) of variables with small reduced costs. Then we will use Lemma 7 to show that the solution ξ so constructed is close to optimal. We change the ξ^B by making cyclic changes within chosen blocks of m columns. To this end, we make the following definitions. Let s be any integer. We define an *s-good* block to be a set of m columns j, as follows. For ease of notation, let us assume the block comprises columns $1, 2, \ldots, m$, and that subscript arithmetic is modulo m. Let us write $\tilde{c}_{ij} = c_{ij} - u_i a_{ij}$. Then the block must be such that, for $1, 2, \ldots, m$:

   (a) $\xi_{ii}^B = 1$. (The changes we intend to make are $\xi_{ii} = 0$, ξ_{i-1}, $i = 1$.)
   (b) $a_{ii} \leqslant \frac{1}{20}\eta$, $a_{i,i+1} \in [a_{ii}, a_{ii} + \frac{1}{20}\eta]$.
   (c) $c_{ji} \leqslant \frac{1}{2}, j \neq i, i-1, \frac{3}{5} < c_{ii} \leqslant \frac{4}{5}$, and

$$c_{i-1,i} \in [\tilde{c}_{ii} + u_{i-1} a_{i-1,i} - s\varepsilon, \tilde{c}_{ii} + u_{i-1} a_{i-1,i} - (s-1)\varepsilon].$$

These imply, by straightforward calculations, the following conditions for column i in the block,

   (i) $\tilde{c}_{ki} \leqslant c_{ki} \leqslant \frac{1}{2}$ for $k \neq i$, $i-1$. (The LP solution must not pick x_{ki}.)
   (ii) $\tilde{c}_{ii} > \frac{3}{5} - \frac{1}{10} = \frac{1}{2}$. (The LP solution must pick x_{ii}.)
   (iii) $(s-1)\varepsilon \leqslant \tilde{c}_{i-1,i} - \tilde{c}_{i,i} \leqslant s\varepsilon$. (The LP solution picks x_{ii}, but $x_{i-1,i}$ is "nearly as good". Observe also that conditions (i)–(iii) imply that $\bar{c}_{ij} = \tilde{c}_{ij} - \tilde{c}_{ii}$ for all i, j.)
   (iv) $a_{i,i+1} - a_{ii}$ is uniform in $[0, \frac{1}{20}\eta]$ independently, even conditional on a_{ii}. (We need this since we wish to argue relative to the optimal basis for LP(b'), and a_{ii} is then conditioned by feasibility of the solution.)

(v) $[\tilde{c}_{ii} + u_{i-1}a_{i-1,i} - s\varepsilon, \tilde{c}_{ii} + u_{i-1}a_{i-1,i} - (s-1)\varepsilon] \subseteq [0, 1]$ for $s \le 1/(2\varepsilon)$. (We require that the probability that $c_{i-1,i}$ falls in this range depends only on its width, and not its location.)

Now the probability that a given column j can be the ith column in an s-good block is clearly a priori at least

$$(\tfrac{1}{2})^{m-2} \times \tfrac{1}{5} \times (\tfrac{1}{20}\eta)^2 \times \varepsilon > \frac{\varepsilon}{M}, \tag{9}$$

independently of i or j.

Suppose now that, for *any* dual feasible basis D, we try to construct s-good blocks. The determination of the u_i conditions only m columns, by Lemma 4. We now scan the remaining $n - m$ columns from left to right, attempting to form good blocks one column at a time. There is one problem. The circularity of our conditions for the columns of the blocks implies that we cannot test column 1 in the block until column m is known. For this reason we have to proceed slightly more carefully, as follows. We denote the elements of the columns as in the above conditions, but the reader should observe that we are in fact examining a general column j for potential addition to the current incomplete block. We scan to select the first column, checking conditions (a)–(c). However, instead of checking condition (b) on $a_{i-1,i}$ (i.e., a_{m1}) (since this is currently impossible) we merely check the necessary condition that $a_{m1} \le \tfrac{1}{10}\eta$. Note that we also check condition (c) involving a_{m1}, but the choice of our ranges implies that this does not further condition a_{m1}. It merely locates the point (c_{m1}, a_{m1}) in a small parallelogram whose projection is uniform on the a_{m1}-axis. Having chosen column 1 we scan to add columns $2, \ldots, m-1$, in each case checking conditions (a)–(c). There is no difficulty since the required quantities are all known at the time of checking. We now select column m. First we check all conditions except that in (b) involving a_{m1}. We will check condition (c) involving a_{mm}, but this does not condition it further than $a_{mm} \le \tfrac{1}{20}\eta$, by the same argument used for a_{m1}. Let us call a "block" which has passed all tests up to this point a *candidate* block. Now the point $\pi = (a_{mm}, a_{m1})$ is uniformly distributed in the rectangle $\Pi = [0, \tfrac{1}{20}\eta] \times [0, \tfrac{1}{10}\eta]$. We now test condition (b) on a_{m1}. This requires that the point π lies in a parallelogram whose area is half that of Π. Thus the probability that a candidate block is s-good is $\tfrac{1}{2}$. If the candidate block is not good, we reject it completely and start afresh. This may appear drastic, but to do otherwise introduces extra conditioning, and the above line of argument fails.

Note that, as we form candidate blocks, the probability of a given column being successfully added is still (independently) at least ε/M, since we check some subset of (a)–(c). However, some candidate blocks will not be good. Now the expected number of successful additions is $(n-m)\varepsilon/M > 4mk$ for large n. Thus

$$\text{Pr(Under } 3k \text{ candidates formed)} = \text{Pr(Under } 3km \text{ successful additions)}$$

$$< e^{-mk/8} < n^{-2m^3},$$

for large n, using Hoeffding's inequality [4, 7]. Then, since each candidate is good with probability $\frac{1}{2}$,

$$\Pr(\text{Under } k \text{ } s\text{-good blocks formed}) < e^{-k/18} + n^{-2m^3} < 2n^{-m^2}$$

for large n, again using Hoeffding's inequality.

Now, since there are less than n possible values of s, and only $O(n^m)$ dual feasible bases D to consider, we see that for large n,

$$\Pr(\exists \, D, s: \text{less than } k \text{ } s\text{-good blocks exist}) < n^{-m^2/2}.$$

For a given s-good block, let $\delta_i = a_{i,i-1} - a_{ii}$ for $i = 1, 2, \ldots, m$. These give the changes in the left hand sides of the constraints in (1) when the cyclic changes $\xi_{ii} = 0$, $\xi_{i-1,i} = 1$ are made to the variables in the block. If we have k s-good blocks we denote the corresponding values by δ_{ir} for $r = 1, 2, \ldots, k$. By construction the δ_{ir} are distributed uniformly and independently in $[0, \frac{1}{20}\eta]$. Let

$$Z_{ir} = \left(\frac{\delta_{ir} - \frac{1}{40}\eta}{\frac{1}{40}\eta} \right) \sqrt{3}.$$

Z_{ir} is uniform in $[-\sqrt{3}, +\sqrt{3}]$ with mean 0, variance 1. Thus

$$\Delta_i^B - \sum_{r=1}^{k} \delta_{ir} = \Delta_i^B - \left(\frac{\eta}{40\sqrt{3}} \right) \sum_{r=1}^{k} Z_{ir} - \frac{1}{40}k\eta = \Gamma_i - \frac{\eta}{40\sqrt{3}} \sum_{r=1}^{k} Z_{ir},$$

where $\Gamma_i = \Delta_i^B - (\frac{1}{40}k\eta) = K - (\frac{1}{40}k\eta) + O(1) = O(1)$, using (8) and the definitions of k, K. Let $\Phi_i = 40\sqrt{3}\Gamma_i/\eta$. Now we use the following lemma, which is a (modified) restatement of Lemma 3.4 of [2].

Lemma 8. *Let Z_{ir}, $i = 1, \ldots, m$, $r = 1, \ldots, k$ be bounded i.i.d. random variables with bounded density, expectation 0 and variance 1. Let $I_i = [\Phi_i - \theta, \Phi_i]$, where $\theta = (\frac{1}{2}\pi k)^{(m+1)/2m}/2^{k/m+1/2}$ and $\Phi_i = O(k^{1/4})$, be m intervals. Let $q_m = 1/(2(1 + (2/\sqrt{3})^m))$, and let S be an arbitrary subset of $\{1, 2, \ldots, k\}$ with $|S| = \lfloor \frac{1}{2}k \rfloor$. Then, for large k,*

$$\Pr\left(\exists S: \sum_{r \in S} Z_{ir} \in I_i, i = 1, \ldots, m \right) \geq q_m. \qquad \square$$

Applying this to our situation, we see that with probability at least q_m we can choose $\lfloor \frac{1}{2}k \rfloor$ s-good blocks to change which will leave the deficit in each left hand side in (1) at most $\eta\theta/(40\sqrt{3})$. Substituting for k, this is less than (say) n^{-10m}, for large n. Now apply Lemma 7. By using only s-good blocks, with probability at least q_m we can construct an integer solution ξ to (1) with objective value z_{IP} satisfying

$$0 \leq z_{LP} - z_{IP} \leq ks\varepsilon + 2mn^{-10m}/\eta < s\alpha(\log n)^2/n, \tag{10}$$

for large n, with $\alpha = 2^{m+20}m^5/\eta^2$, say. Now we try the s-good sets for $s = 1, 2, \ldots, t$, seeking to find s so that (10) is satisfied. (Note that t can be as large as $n/(\log n)^2$.) Let $p = 1 - q_m$. Each "trial" is independent (since it involves only the δ_{ir}). The probability that we fail t times is thus at most p^t, as required. This proves the first statement of Theorem 1. The second is derived from it as follows.

Let $t = \lceil 2 \log n / \log(1/p) \rceil$, say. Using a simple bound on expectation, and the fact that $z_{LP} - z_{IP} \le n$ always, we have

$$0 \le E(z_{LP} - z_{IP}) \le \frac{\alpha (\log n)^2}{n} \sum_{s=1}^{t} s p^{s-1} + n p^t$$

$$< \frac{\alpha}{(1-p)^2} \frac{(\log n)^2}{n} + \frac{1}{n}$$

$$< \frac{2\alpha}{(1-p)^2} \frac{(\log n)^2}{n}, \tag{11}$$

for large enough n. The final statement of Theorem 1 follows from this and Lemma 9 below. Let

$$C_1 = \{(i,j): \bar{c}_{ij} \in (0, 2 \log n / n]\}.$$

It is shown in Lemma 9 that

$$E(|C_1|) < 4m^2 \log n. \tag{12}$$

Now, there are at most $2m$ fractional basic variables by Lemma 3. Thus clearly

$$\mathcal{N} \le 2m + |C_1| + \frac{(z_{LP} - z_{IP})}{(2 \log n)/n}.$$

Taking expectations of both sides in this inequality, the theorem now follows using (11) and (12).   □

4. Proof of Theorem 2

Let $\eta > 0$ be given, and let $0 \le \tau \le 1$ be the required failure probability. Assume n is large enough that

$$\Pr((1) \text{ is } \eta\text{-feasible}) \ge 1 - \tfrac{1}{2}\tau.$$

Now consider the following algorithm.

Algorithm OPT_τ.
 Step 1. $t \leftarrow \lceil \log(\tfrac{1}{2}\tau) / \log p \rceil$.
 Step 2. Solve LP(b) giving x^B. Modify this to ξ^B.
 Step 3. Let $T = \{(i,j): \xi^B_{ij} = 0, \ 0 \le \bar{c}_{ij} \le t\alpha(\log n)^2 / n\}$, where $\bar{c}_{ij}$ are the reduced costs for x^B.
 Step 4. Compute $\min_{S \subseteq T} c\xi_S$, for all S and ξ_S such that
   (a) there is at most one $(i,j) \in S$ for all j.
   (b) $\sum_{(i,j) \in S} \bar{c}_{ij} \le t\alpha(\log n)^2 / n$.
   (c) ξ_S is feasible in (1), where ξ_S is obtained from ξ^B by $\xi_{ij} \leftarrow 1$, $\xi_{kj} \leftarrow 0$ $(k \ne i)$, for each $(i,j) \in S$.

The minimising S at termination (if defined), gives an optimal solution $\xi_0 = \xi_S$ to (1). This follows directly, since the search accounts for all feasible integer solutions having objective value at most $z_{LP} + t\alpha(\log n)^2/n$. Failure because (1) is not η-feasible has probability at most $\frac{1}{2}\tau$. The probability that S is undefined is, by Theorem 1, less than p', which is at most $\frac{1}{2}\tau$ by choice of t. Thus it remains to analyse the running time of OPT_τ.

Fix any dual feasible basis D, and let

$$J_l = ((2(l-1)\log n)/n, (2l\log n)/n],$$

$$C_l = \{(i,j): \bar{c}_{ij} \in J_l\}.$$

(C_1 has already been used in Section 3.)

Lemma 9. $\Pr(|C_l| \geqslant 4m^2 \log n) \leqslant n^{-2m^2/3}$.

Proof. For any column j not determining the u_i, we have

$$E(\text{number of } (i,j) \in C_l) \leqslant E(\text{number of pairs } i, r: \tilde{c}_{ij} - \tilde{c}_{rj} \in J_l)$$

$$\leqslant m(m-1)\Pr(c_{1j} \in [\lambda, \lambda + (2\log n)/n])$$

$$(\text{for } \lambda = u_1 a_{1j} + c_{2j} - u_2 a_{2j})$$

$$\leqslant (2m^2 \log n)/n.$$

Hence, using Hoeffding's inequality,

$$\Pr(|C_l| > 4m^2 \log n) \leqslant e^{-2m^2 \log n/3} = n^{-2m^2/3}. \qquad \square$$

Since there $O(n^m)$ dual feasible bases, and we are interested in only $O(\log n)$ values of l, it follows from Lemma 9 that we may assume $|C_l| \leqslant 4m^2 \log n$ for all l in the optimal basis. We examine the search in OPT_τ relative to the sets C_l. To these we have only to add the (less than $2m$ by Lemma 3) fractional variables which have been rounded down. The argument is now similar to that in [2].

Thus let $R_r = \bigcup_{l=2^{r-1}}^{2^r-1} C_l$ $(r=1,2,\ldots)$. We clearly need only consider S for which

$$|S \cap R_r| \leqslant \left\lfloor \frac{(t\alpha(\log n)^2)/n}{((2^r-2)\log n)/n} \right\rfloor = \left\lfloor \frac{t\alpha\log n}{2^r-2} \right\rfloor = \rho_r,$$

say, if $r > 1$. Note that $\rho_r = 0$ for all large enough r. Thus, with probability $1 - o(1)$, we need to check at most N sets S, where

$$N = 2^{2m} \times 2^{4m^2 \log n} \times \prod_{r=2}^{\infty} N_r,$$

sets S, where

$$N_r = \sum_{r=0}^{\rho_r} \binom{\lfloor 4m^2 \log n 2^{r-1} \rfloor}{r}.$$

Now define $r_0 = \max\{r: t\alpha/(2^r - 2) \geq m^2 2^{r-1}\}$, and note that r_0 is independent of n. For all r, clearly $N_r \leq 2^{4m^2 \log n 2^{r-1}}$. However, for $r \geq r_0$, using well-known bounds on binomial coefficients, we have the better bound

$$N_r \leq 2 \left(\frac{e\,4m^2 \log n 2^{r-1}}{\rho_r} \right)^{\rho_r} \leq 2^{\gamma(r/2^r) \log n},$$

after a straightforward calculation, for large n and suitably chosen γ (independent of n and r). Hence

$$N \leq 2^{2m + 4m^2 \log n + \sum_{r=2}^{r_0} 2^{r+1} m^2 \log n + \sum_{r=r_0+1}^{\infty} \gamma(r/2^r) \log n} = n^{d-1},$$

for a suitable d. Since each set S takes $O(n)$ time to process, the theorem follows. $\square$

5. Extension to more general distributions

The preceding analysis can be extended quite easily to more general distributions. We can assume for example that the $c_{i,j}$ are drawn independently from one distribution F_c and the $a_{i,j}$ are drawn independently from another F_a. These distributions should have bounded supports $[0, L_c]$, $[0, L_a]$ respectively and positive densities throughout. Under these circumstances our results go through with only marginal changes. Thus the inequality (9) will have a right hand side ε/M' where M' depends on m, η and F_a. As in the case of the uniform distribution, we will be able to show the existence of sufficiently many s-good blocks with high probability. We can then use Lemma 8 to finish the proof of the theorem, since it handles sufficiently general random variables. The proof of Theorem 2 will not change much either, since Lemma 9 will require little modification. The same comments seem to be true for distributions which have positive density on $[0, \infty)$ and small enough tail probabilities, e.g., the exponential distribution.

If the distributions are discrete, then our analysis does not go through so easily. However, we imagine that, for $c_{i,j}$ uniform in $\{0, 1, \ldots, L_c\}$ and $a_{i,j}$ uniform in $\{0, 1, \ldots, L_a\}$ say, it may not be too difficult to modify our proof, provided L_c and L_a grow sufficiently fast with n.

Other types of distribution might pose more severe problems and are left to future research.

6. Final remarks

(1) For convenience in the development, we have stated all our results for large n, but obviously they can be modified to hold for all n (simply by enlarging the constants if necessary).

(2) We claim only that our results hold for constant m, but it would appear that generally they can be modified to remain true if m grows slowly enough with n.

However, the running time is (at least) $\Omega(C^{m^2 \log n})$, for some constant C. Thus the algorithm is not provably polynomial time for non-constant m.

(3) The approach we have developed here appears to be applicable to some related problems, for example the multiple-choice multidimensional knapsack problem.

References

[1] J.F. Benders and J.A.E.E. van Nunen, "A property of assignment type mixed integer linear programming problems," *Operations Research Letters* 2 (1983) 47–52.

[2] M.E. Dyer and A.M. Frieze, "Probabilistic analysis of the multidimensional knapsack problem," *Mathematics of Operations Research* (1989) 162–176.

[3] A.V. Goldberg and A. Marchetti-Spaccamela, "On finding the exact solution of a zero-one knapsack problem," in: *Proceedings of the 16th Annual ACM Symposium on the Theory of Computing* (ACM Press, New York, 1984) pp. 359–368.

[4] W. Hoeffding, "Probability inequalities for sums of bounded random variables," *Journal of the American Statistical Association* (1963) 13–30.

[5] K.O. Jörnsten and M. Nasberg, "A new Lagrangian relaxation approach to the generalised assignment problem," *European Journal of Operations Research* (1986) 313–323.

[6] G.S. Lueker, "On the average difference between the solutions to linear and integer knapsack problems," in: R.L. Disney and T.J. Ott, eds., *Applied Probability-Computer Science: The Interface, Vol. I, Progress in Computer Science* (Birkhäuser, Boston, MA, 1982) pp. 489–504.

[7] C.J.H. McDiarmid, "On the method of bounded differences," in: J. Siemons, ed., *Surveys in Combinatorics. London Mathematical Society Lecture Notes No. 141* (1989) pp. 148–188.

[8] G.T. Ross and R.M. Soland, "A branch-and-bound algorithm for the generalised assignment problem," *Mathematical Programming* 8 (1975) 91–103.