Approximating Low-Congestion Routing and Column-Restricted Packing Problems

Alok Baveja [a,1] and Aravind Srinivasan [b,2]

[a] *School of Business, Rutgers University, Camden, NJ 08102, USA.*

[b] *Bell Laboratories, Lucent Technologies, 600–700 Mountain Ave., Murray Hill, NJ 07974-0636, USA.*

Abstract

We contribute to a body of research asserting that the fractional and integral optima of *column-sparse* integer programs are "nearby". This yields improved approximation algorithms for some generalizations of the knapsack problem, with applications to low-congestion routing in networks, file replication in distributed databases, and other packing problems.

Keywords: Algorithms; approximation algorithms; packing; integer programming; routing.

1 Introduction

Let Z_+ denote the set of non-negative integers, v^T the transpose of a (column) vector v, and $[k] \doteq \{1, 2, \ldots, k\}$. A key family of packing problems that includes classical NP-hard problems such as knapsack, independent sets in graphs, matchings in hypergraphs etc., has been introduced in [5]. These problems, named *column-restricted packing integer programs* (CPIPs) in [5], are integer programs of the form "maximize $w^T x$ subject to: $Ax \leq b$ and $x_j \in \{0, 1, \ldots, d_j\}$ for each j" ($\forall j,\ d_j \in Z_+$); all entries of the matrix A and

[1] Work supported in part by Grants 2-02007 and 2-02206, through Rutgers University's Office of Research and Sponsored Programs. E-mail: `baveja@crab.rutgers.edu`

[2] Part of this work was done while at the School of Computing, National University of Singapore, Singapore 119260, and was supported in part by National University of Singapore Academic Research Fund Grant RP970607. E-mail: `srin@research.bell-labs.com`

of the vectors b and w are non-negative. Also, all nonzero entries in any given column of A are the *same*, hence "column-restricted". Suppose, e.g., that we have files $F_1, \ldots, F_n$ with F_j having size ρ_j, and m servers, each having some *capacity*. If F_j is selected, it is to be placed on a specified subset S_j of the m servers; the total load on any given server should not exceed its capacity. Given a benefit for each F_j, the problem of selecting a subset of the files that maximizes the total benefit subject to the above constraints, is a CPIP with $d_j = 1$ for all j. Since CPIPs are NP-hard, there is much interest in developing approximation algorithms for them. The best general provable approximations to-date for such packing problems start by considering the linear programming (LP) relaxation where each x_j is relaxed to be a *real* lying in $[0, d_j]$; the objective function value y^* of an optimal solution of this LP upper-bounds the optimal objective function value of the CPIP. The crucial step then is to show how to "round" the LP solution well. It will be immediate that our rounding runs in polynomial time; as is common, our focus will be on lower-bounding the quality of the final solution.

A related generalization of knapsack, called *low-congestion routing problems* (LCRPs) here, has been considered in [1]. The variables are $\{z_{u,v} : u \in [k], v \in [\ell_u]\}$ where the ℓ_u are some given integers. The objective is to maximize $\sum_{u\in[k]} w_u(\sum_{v\in[\ell_u]} z_{u,v})$ subject to: (i) a system of m linear constraints $Az \leq b$, where z is the vector of the $z_{u,v}$'s (arranged in any fixed order), (ii) $\forall u, \ \sum_v z_{u,v} \leq 1$ and (iii) $\forall (u,v), \ z_{u,v} \in \{0,1\}$. (The columns of A are naturally indexed by pairs (u,v).) We also have: **(C1)** $\forall u, \ w_u \geq 0$; **(C2)** all entries in A and b are non-negative, and, importantly, **(C3)** for each u, all nonzero entries in the columns of A that are associated with $z_{u,1}, z_{u,2}, \ldots, z_{u,\ell_u}$ are the same. (If $\ell_u = 1$ for all u, then this is a CPIP with $d_j = 1$ for all j.) The LP relaxation for an LCRP is to relax the integrality constraints (iii) to "$\forall (u,v), \ z_{u,v} \in [0,1]$".

Rounding for LCRPs is closely related to the *unsplittable flow* problem (UFP) [4]. Suppose we are given a graph G, and a collection of vertex-pairs $\{(s_u, t_u) : u \in [k]\}$. Each pair (s_u, t_u) has a *demand* $\rho_u > 0$ and a weight $w_u > 0$; each edge f of G has a *capacity* $b_f > 0$. The UFP is to select a maximum-weight subset of the pairs and a *single* (s_u, t_u)-path for each chosen (s_u, t_u), so that for each edge f, the total demand of the pairs whose chosen path use f, is at most b_f. (The UFP generalizes the classical NP-hard *edge-disjoint paths* problem, where $\forall u, f, \ \rho_u = b_f = 1$.) A natural LP relaxation lets the demand for (s_u, t_u) be split among several (s_u, t_u)-paths; if m is the number of edges of G, then any *basic* optimal solution to the LP will have a set $\mathcal{S}_u$ of at most m many (s_u, t_u)-paths. Rounding this is easily written as an LCRP instance where $z_{u,v}$ is the indicator for choosing the vth path in $\mathcal{S}_u$. Conversely, rounding for LCRPs can be treated as rounding LP solutions for UFPs.

Suppose we are given a CPIP/LCRP; by scaling all entries of A and b, we will assume that $\max_{i,j} A_{i,j} \leq 1$. Some key parameters will be: (a) d, the maximum number of nonzero entries in any column of A; (b) $t = \max_j \sum_i A_{i,j}$, the maximum column sum of A: note, crucially, that $t \leq d$ since $\max_{i,j} A_{i,j} \leq 1$; and (c) $B = \min_i b_i$. We assume without loss of generality that $\min_{i,j:\ A_{i,j}\neq 0} b_i/A_{i,j} \geq 1$: if $A_{i,j} > b_i$, we can set the variable corresponding to the jth column to 0. It is hence generally assumed that $B \geq 1$. (Via the same scaling process as above, we will also take $\max_{i,j} A_{i,j} = 1$ unless stated otherwise; so, $t \geq 1$.)

Let y^* be the optimal solution value to the LP relaxation of a given CPIP/LCRP. It is shown in [5] that an integral solution of value at least $K_0 y^*/d^{1/\lfloor B \rfloor}$ can be efficiently found for CPIPs and for LCRPs. (If $B = 1$, there exists, e.g., an integral solution of value at least $K_1 y^*/(\sqrt{m}\log m)$ for LCRPs [5].) $K_0 > 0$ and $K_1 > 0$ denote constants here; see [5,1] for additional results. By building on [5,1], we improve these results to get a "$K_2 y^*/t^{1/\lfloor B \rfloor}$" result for both CPIPs and LCRPs. The main point to note is that d has essentially been replaced by t; for the file replication problem discussed above, this means, e.g., that "small" files can be replicated in many servers. For UFP instances, t denotes the maximum amount of capacity requested by any *virtual circuit* (connection path); our result shows that long paths can be tolerated if they belong to small demands ρ_i. From an optimization viewpoint, the result proves the intuitive idea that "small" $A_{i,j}$ should have only a limited effect on the optimal value.

It is shown in [2] that unless $NP \subseteq ZPP$, we cannot approximate the maximum independent set problem on graphs even to within a factor of $n^{1-\epsilon}$ in polynomial time, where n denotes the number of vertices and $\epsilon > 0$ is any constant. Since this problem can be formulated as a CPIP with $O(n^2)$ constraints, it is unlikely that good approximation guarantees can be presented for *all* CPIPs. So our focus here is on methodology and asymptotic improvements in the approximation guarantee.

2 Preliminaries

2.1 A common approach to LCRPs and CPIPs

We will use UFP terminology for LCRPs; e.g., we identify the m rows of the matrix A with the edge-set E of a network $G = (V, E)$. For LCRPs, we will use the randomized rounding approach of [1]. We first solve the LP relaxation of the given LCRP instance. Let $P_{i,j}$ denote the jth path in the set $\mathcal{S}_i$ of (s_i, t_i)-paths given by the LP; $P_{i,j}$ carries a flow of $\rho_i z^*_{i,j}$, where $z^*_{i,j} \in [0, 1]$. The rounding idea is to choose a suitable $\gamma > 1$, and then pick each $P_{i,j}$ *independently*, with probability $z^*_{i,j}/\gamma$; let $z_{i,j}$ be the indicator variable for this event. For each i: if at least one $P_{i,j}$ was picked, we choose one of them arbitrarily (if no $P_{i,j}$ was picked by our rounding, then no (s_i, t_i)-path is chosen).

We next present a rounding scheme for CPIPs, which closely follows the above; this is motivated by a simple connection between CPIPs and LCRPs shown in [1]. Suppose we have an optimal solution $\{x_1^*, x_2^*, \ldots, x_n^*\}$ for the LP relaxation of a given CPIP instance. Let $\gamma > 1$ be some parameter as above. Rounding the given fractional CPIP solution can be done by applying the above-seen rounding approach for LCRPs, as follows. (The LCRP that we construct now will in fact have $\ell_u = 1$ for all u.) Let $a_j = \sum_{i=1}^{j-1} \lceil x_j^*/\gamma \rceil$. We will construct a_{n+1} *independent* binary random variables $z_{1,1}, z_{2,1}, \ldots, z_{a_{n+1},1}$, as follows. Consider any j, $1 \le j \le n$. (Informally, the variables $z_{a_j+1,1}, z_{a_j+2,1}, \ldots, z_{a_{j+1},1}$ are associated with x_j^*.) For $1 \le r \le (\lceil x_j^*/\gamma \rceil - 1)$, $\Pr[z_{a_j+r,1} = 1] = 1$; also, $\Pr[z_{a_{j+1},1} = 1] = x_j^*/\gamma - \left\lceil x_j^*/\gamma \right\rceil + 1$. Given a rounding that satisfies $Az \le b$, a feasible solution to the CPIP is naturally obtained by setting $x_j = \sum_{r=1}^{\lceil x_j^*/\gamma \rceil} z_{a_j+r,1}$ for each j. Thus we have a simple reduction from CPIPs to LCRPs. The values x_j^*/γ may be very large (say, exponential in the size of the CPIP instance). However, it will be seen in the constructive version of Lemma 1 in Section 2.2, that all the $z_{u,v}$ with $\Pr[z_{u,v} = 1] = 1$ will get rounded to one. Thus, the real rounding issue is only for the random variables $z_{a_j,1}$; hence, the running time for CPIPs remains polynomial. Also, it is easy to see that for each j, $\mathbf{E}[\sum_{r=1}^{\lceil x_j^*/\gamma \rceil} z_{a_j+r,1}] = x_j^*/\gamma$. Thus, as with the LCRP rounding approach, we have the property **(P1)**: "$\forall f \in E,\ \mathbf{E}[(Az)_f] \le b_f/\gamma$" here also.

Having presented this reduction from CPIPs to LCRPs, we view the problem from now on as rounding the $z_{u,v}$'s for an LCRP instance, in order to satisfy $Az \le b$. We will only use LCRP (UFP) terminology. We next present two more properties (P2) and (P3), which hold for the LCRP rounding scheme given in the first paragraph of Section 2.1, as well as for our above rounding scheme for LCRP instances derived from CPIP instances. Define $\psi_u^* \in [0,1]$ by $\psi_u^* = \sum_{v \in [\ell_u]} \Pr[z_{u,v} = 1]$. It is easy to check that we have the following: **(P2)**: "$y^* = \gamma \sum_u \psi_u^*$". We claim that **(P3)**: "$\forall u,\ \sum_{1 \le v < v' \le \ell_u} \Pr[z_{u,v} = 1] \cdot \Pr[z_{u,v'} = 1] \le \psi_u^*/(2\gamma)$" also holds. This is trivially true for our LCRP instances derived from CPIP instances as above, since they have $\ell_u = 1$ for all u. For general LCRP instances, note from the first paragraph of Section 2.1 that $\psi_u^* \le 1/\gamma$. Next, since $\sum_{v \in [\ell_u]} \Pr[z_{u,v} = 1] = \psi_u^*$, it can be checked that $\sum_{1 \le v < v' \le [\ell_u]} \Pr[z_{u,v} = 1] \cdot \Pr[z_{u,v'} = 1] \le (\psi_u^*)^2/2$, which in turn is at most $\psi_u^*/(2\gamma)$ since $\psi_u^* \le 1/\gamma$.

We focus on rounding the $z_{u,v}$'s to achieve $Az \le b$, given (P1), (P2), and (P3).

2.2 More on the rounding

For each i, let Z_i be the indicator variable for having chosen at least one path $P_{i,j}$. For $f \in E$, let E_f be the "bad" event that $R_f \doteq \sum_{(i,j):\ f \in P_{i,j}} \rho_i z_{i,j}$ is more than the capacity b_f of f. Let $\overline{E_f}$ denote the complement of E_f; thus, $\bigwedge_{f \in E} \overline{E_f}$ is the event that all the edges' capacities are respected. The idea would be to analyze the above rounding process for a suitable γ and to show that we can simultaneously have $\bigwedge_{f \in E} \overline{E_f}$ **and** keep the objective function $\sum_{i \in [k]} w_i Z_i$ at least as large as y^*/ν for some $\nu > 1$. As in [5], it will be useful to partition the set of pairs $\{(s_u, t_u) : u \in [k]\}$ into groups, and process each group separately. Suppose we are focusing only on a group $\{(s_u, t_u) : u \in S\}$, for some $S \subseteq [k]$. For any $i \in S$, any j, and any edge f with $f \in P_{i,j}$, we define $R_f(S, i, j)$ to be the total flow on edge f from the pairs $\{(s_u, t_u) : u \in S\}$, *excluding* path $P_{i,j}$; $R_f(S, i, j) = \sum_{r,s} \rho_r z_{r,s}$, where the sum is over all $(r, s) \neq (i, j)$ such that $r \in S$ and $f \in P_{r,s}$. Lemma 3.4 of [1] will be helpful.

Lemma 1 ([1]) *Fix any $S \subseteq [k]$. For any $X \subseteq \{(i,j) : i \in S\}$ and any set $\{\alpha_{i,j} \in \{0,1\} : (i,j) \in X\}$, let $\mathcal{A} \equiv (\forall (i,j) \in X,\ z_{i,j} = \alpha_{i,j})$. Suppose $\Pr[\bigwedge_{f \in E} \overline{E_f} \mid \mathcal{A}] > 0$. Then, for any $i \in S$, $\mathbf{E}[Z_i \mid (\mathcal{A} \wedge (\bigwedge_{f \in E} \overline{E_f}))]$ is at least*

$$(\sum_j \Pr[z_{i,j} = 1 \mid \mathcal{A}](1 - \sum_{f \in P_{i,j}} \Pr[R_f(S, i, j) > (b_f - \rho_i) \mid \mathcal{A}])) -$$
$$\sum_{j < j'} \Pr[(z_{i,j} = z_{i,j'} = 1) \mid \mathcal{A}].$$

We now present a constructive version of Lemma 1 due to [1], for the special case where $\rho_i = 1$ for all $i \in S$. We suppress some usages of S to avoid messy notation (e.g., we let $R_f(i,j)$ denote $R_f(S,i,j)$). Since $\rho_i = 1$ for all $i \in S$, $R_f(i,j)$ is an integer; so "$R_f(i,j) > b_f - \rho_i$" $\equiv$ "$R_f(i,j) \geq \lfloor b_f \rfloor$". Next, let e denote the base of the natural logarithm, and suppose $Y_1, \ldots, Y_\ell$ are *independent* random variables, each taking values in $[0,1]$. Let $Y = \sum_i Y_i$, with $\mathbf{E}[Y] \leq \mu$. Then for any $\xi \geq 0$, a classical bound of [3] shows that

$$\Pr[Y \geq \mu(1+\xi)] \leq \mathbf{E}[(1+\xi)^Y]/(1+\xi)^{\mu(1+\xi)} \leq (e^\xi/(1+\xi)^{1+\xi})^\mu. \quad (1)$$

Define $\mu_{f,i,j} \doteq \mathbf{E}[R_f(i,j)]$, and define $\delta_{f,i,j}$ by $\mu_{f,i,j}(1+\delta_{f,i,j}) = \lfloor b_f \rfloor$. Note that $R_f(i,j)$ is a sum of independent random variables, each lying in $[0,1]$; so, if $\delta_{f,i,j} \geq 0$, then

$$\begin{aligned} \Pr[R_f(i,j) > b_f - \rho_i] &= \Pr[R_f(i,j) \geq \lfloor b_f \rfloor] \\ &\leq \mathbf{E}[(1+\delta_{f,i,j})^{R_f(i,j)}]/(1+\delta_{f,i,j})^{\lfloor b_f \rfloor}. \end{aligned} \quad (2)$$

Constructive version of Lemma 1. Let $X = \emptyset$ in Lemma 1; hence $\mathcal{A}$ is the tautology. Since the random variables $z_{i,j}$ are independent, the event $\mathcal{E}$: "$\forall i, j$: if $\Pr[z_{i,j} = 1] = 1$, then $z_{i,j}$ gets rounded to one; else, $z_{i,j}$ gets rounded to zero" happens with nonzero probability. Note next from (P1) that if $\gamma \geq 1$, then the event $\mathcal{E}$ implies the event $\bigwedge_{f \in E} \overline{E_f}$. So, since $\gamma > 1$, "$\Pr[\bigwedge_{f \in E} \overline{E_f}] > 0$" will always hold. Suppose $\delta_{f,i,j} \geq 0$ for all $i \in S$, f, and j. Thus, by Lemma 1, (2), and the linearity of expectation, there exists a feasible solution with $\sum_{i \in S} w_i Z_i$ at least as high as

$$\sum_{i \in S} w_i \cdot [(\sum_j \Pr[z_{i,j} = 1](1 - \sum_{f \in P_{i,j}} \frac{\mathbf{E}[(1 + \delta_{f,i,j})^{R_f(i,j)}]}{(1 + \delta_{f,i,j})^{\lfloor b_f \rfloor}})) - L_i], \qquad (3)$$

where L_i is shorthand for $\sum_{j < j'} \Pr[(z_{i,j} = z_{i,j'} = 1)]$.

To efficiently construct a feasible solution of value at least that given by (3), the algorithm of [1] rounds all $z_{i,j}$ with $\Pr[z_{i,j} = 0] = 1$ to zero; all $z_{i,j}$ with $\Pr[z_{i,j} = 1] = 1$ are then rounded to one. After this, a sequential method to round the variables suitably (based on a potential function) is shown in [1]; the reader is referred to [1] for the details.

3 Results

We will always take $\gamma \geq e$ unless specified otherwise; this will let us satisfy the condition "$\delta_{f,i,j} \geq 0$" of (2), as shown now. Note that for any f, i, j, $\mu_{f,i,j} = \mathbf{E}[R_f(i,j)] \leq b_f/\gamma$, by (P1). Thus, for all f, i, j, $\lfloor b_f \rfloor / \mu_{f,i,j} \geq e \lfloor b_f \rfloor / b_f \geq e/2 > 1$, since $b_f \geq 1$. Hence, $\delta_{f,i,j} \geq 0$. Next, for any $S \subseteq [k]$, let $W(S) \doteq \gamma \sum_{i \in S} w_i \psi_i^*$; since each ψ_i^* is inversely proportional to γ, $W(S)$ does not depend on the value of γ. By a path-selection, we mean a choice of one (s_u, t_u)-path for some of the given pairs (s_u, t_u).

Lemma 2 **(i)** *Suppose $B \geq 1$ and that for some $S \subseteq [k]$, we have $\rho_i = 1$ for all $i \in S$. Suppose also that $\forall i \in S\ \forall j\ \forall f \in P_{i,j}$, $\mathbf{E}[(1 + \delta_{f,i,j})^{R_f(i,j)}]/(1 + \delta_{f,i,j})^{\lfloor b_f \rfloor} \leq p$, for some $\gamma \geq e$ and some p. Then, we can efficiently construct a path-selection that violates no edge-capacity, such that $\sum_{i \in S} w_i Z_i \geq [(1 - dp - 1/(2\gamma))/\gamma] \cdot W(S)$.* **(ii)** *Suppose $B \geq 1$ and $\exists S \subseteq [k]\ \exists \lambda \geq 1\ \forall i \in S,\ \rho_i \geq 1/\lambda$. For any given $B^* \in [1, B]$, let $\gamma = \gamma(t, B^*, \lambda) = e(et\lambda(\lfloor B^* \rfloor + 1))^{1/\lfloor B^* \rfloor}$. We can efficiently construct a path-selection for $\{(s_u, t_u) : u \in S\}$ with $\sum_{i \in S} w_i Z_i \geq g(t, B^*, \lambda) \cdot W(S)$, where $g(t, B^*, \lambda) = (\frac{\lfloor B^* \rfloor}{\lfloor B^* \rfloor + 1} - \frac{1}{2\gamma(t, B^*, \lambda)})/(\lambda \cdot \gamma(t, B^*, \lambda))$.*

Proof. (i) As argued in the first paragraph of Section 3, $\delta_{f,i,j} \geq 0$ for all $i \in S$, f, and j. Thus, the constructive version of Lemma 1 lets us efficiently construct a feasible path-selection, with $\sum_{i \in S} w_i Z_i$ being at least the value (3).

By definition of d, each $P_{i,j}$ is composed of at most d edges. So we can efficiently construct a feasible path-selection, with $\sum_{i\in S} w_i Z_i \geq \sum_{i\in S} w_i \cdot [(\sum_j \Pr[z_{i,j} = 1]\cdot(1-dp)) - \sum_{j<j'} \Pr[z_{i,j}=1]\cdot\Pr[z_{i,j'}=1]]$. Recalling that $\sum_j \Pr[z_{i,j}=1] = \psi_i^*$ and using (P3), part (i) is proved.

For (ii), we borrow an idea from [5]. Scale up all demands ρ_i, $i \in S$, to 1; divide the current fractional solution (i.e., all the $z_{i,j}^*$, $i \in S$) by λ to get a feasible fractional solution to this problem. It is easy to check that this problem induced by S has "d" being at most $t\lambda$, since $\rho_i \geq 1/\lambda$ for all $i \in S$. Note next from (1) that for any f, i, j, $\mathbf{E}[(1+\delta_{f,i,j})^{R_f(i,j)}]/(1+\delta_{f,i,j})^{\lfloor b_f \rfloor} \leq e^{\lfloor b_f \rfloor - b_f/\gamma}/(\gamma(\lfloor b_f \rfloor/b_f))^{\lfloor b_f \rfloor} \leq (b_f/\lfloor b_f \rfloor)^{\lfloor b_f \rfloor}\cdot(e/\gamma)^{\lfloor b_f \rfloor} \leq e\cdot(e/\gamma)^{\lfloor B^* \rfloor}$. (The bound $(b_f/\lfloor b_f \rfloor)^{\lfloor b_f \rfloor} \leq e$ holds, since $b_f \leq 1+\lfloor b_f \rfloor$ and since for $y \geq 1$, $(1+1/y)^y \leq e$.) Now set $d \leq t\lambda$ and $\gamma = \gamma(t, B^*, \lambda)$, and use part (i). $\square$

Lemma 3 *Suppose, for some h, any given LCRP instance with $B \geq 1$, any $B^* \in [1, B]$, and any $\lambda > 1$, we can efficiently construct a path-selection for $\{(s_u, t_u) : u \in S'\}$ such that $\sum_{u\in S'} w_u Z_u \geq h(t, B^*, \lambda)\cdot W(S')$; S' here denotes $\{u \in [k] : \rho_u > 1/\lambda\}$. Given any LCRP instance and any $(\alpha, \beta, \lambda, \kappa)$ such that $\alpha \geq 1$, $\beta \in (0,1)$, $\lambda > 1$, $\alpha\kappa \geq 1/B$, and $\kappa\lambda \leq (1-\beta)(\lambda-1)^2$, let $S'' = \{u \in [k] : \rho_u \leq 1/(\alpha\lambda)\}$. We can efficiently construct a path-selection for $\{(s_u, t_u) : u \in S''\}$ such that $\sum_{u\in S''} w_u Z_u \geq (\beta \min_{i\geq 1} h(\alpha t\lambda^i, \alpha\kappa i B, \lambda))\cdot W(S'')$.*

Proof. For $i \geq 1$, define $S_i = \{u : \alpha^{-1}\lambda^{-(i+1)} < \rho_u \leq \alpha^{-1}\lambda^{-i}\}$. S'' gets partitioned into $S_1, S_2, \ldots, S_\ell$ for some finite ℓ; clearly, at most k of the S_i are nonempty. For each $i \geq 1$ such that $S_i \neq \emptyset$, consider the LCRP $\mathcal{P}_i$ where we only consider the pairs $\{(s_u, t_u) : u \in S_i\}$. Crucially, by building on an idea of [5], set the capacity of each edge f here to be $b_{f,i} = \kappa i B/\lambda^i + \beta \sum_{(r,s)} \rho_r z_{r,s}^*$; the sum is over all (r,s) such that $r \in S_i$ and $f \in P_{r,s}$. We get a feasible fractional solution for this problem by multiplying $z_{r,s}^*$, for all $r \in S_i$ and for all s, by β. Next, scale up all the demands and capacities by $\alpha\lambda^i$; t and B become $\alpha t\lambda^i$ and $\alpha\lambda^i \min_f b_{f,i} \geq \alpha\kappa i B$ respectively. Now, $\alpha\kappa i B \geq 1$ since $\alpha\kappa \geq 1/B$. So we can efficiently construct a path-selection of value at least $\beta\cdot h(\alpha t\lambda^i, \alpha\kappa i B, \lambda)\cdot W(S_i)$ for $\{(s_u, t_u) : u \in S_i\}$. We put together these solutions for all $i \in [\ell]$; we now argue that this is feasible, i.e., that no edge's capacity is violated. It would suffice to show that for each f, $\sum_i b_{i,f} \leq b_f$. To see this, note that $\sum_{i\geq 1} b_{i,f} = \sum_{i\geq 1}(\kappa i B/\lambda^i + \beta \sum_{(r,s): r\in S_i,\ f\in P_{r,s}} \rho_r z_{r,s}^*) \leq \beta b_f + \sum_{i\geq 1} \kappa i B/\lambda^i$. The last term is at most $\beta b_f + (1-\beta)B \leq b_f$ since: (a) for any $y \in (0,1)$, $\sum_{i\geq 1} i y^i = y/(1-y)^2$, and (b) $\kappa\lambda/(\lambda-1)^2 \leq 1-\beta$. The quality of the final solution is at least $\beta \sum_{i\geq 1} h(\alpha t\lambda^i, \alpha\kappa i B, \lambda) W(S_i) \geq \beta(\min_{i\geq 1} h(\alpha t\lambda^i, \alpha\kappa i B, \lambda))\cdot\sum_{i\geq 1} W(S_i)$, which equals $\beta \min_{i\geq 1} h(\alpha t\lambda^i, \alpha\kappa i B, \lambda)\cdot W(S'')$. $\square$

Theorem 4 *There are constants $K_2, K_3 > 0$ such that:* **(i)** *for any LCRP instance with $B \geq 1$, we can efficiently construct a path-selection of value at least $K_2 y^*/t^{1/\lfloor B \rfloor}$;* **(ii)** *for any $\epsilon \in (0, 1/2]$ and any LCRP instance with $B \geq (K_3/\epsilon^4)\ln(t/\epsilon)$, we can efficiently construct a path-selection of value at least $y^*(1-\epsilon)(1-1/(2e))/e$.*

Proof. (i) Suppose $t, B, \lambda \geq 1$. It is easy to see that $e \leq \gamma(t, B, \lambda) \leq 2e^2\lambda t^{1/\lfloor B \rfloor}$. So, **(Q)**: "$g(t, B, \lambda) \leq (1/2 - 1/(2e))/(2\lambda^2 e^2 t^{1/\lfloor B \rfloor})$" holds. Let $K_1', K_2' > 0$ denote some absolute constants; we set $B^* = B$ throughout. Take $\lambda = 3.75$, $S' = \{u : \rho_u > 1/\lambda\}$, and $S'' = ([k] - S')$. Lemma 2(ii) and (Q) show that we can efficiently handle S' with objective function value $\nu_1 \geq K_1' W(S')/t^{1/\lfloor B \rfloor}$; we have, e.g., that $K_1' \geq (1/4 - 1/(4e))/(3.75e)^2$. Next, letting $\alpha = \kappa = 1$ and $\beta = 1/2$, Lemma 3 and Lemma 2(ii) show that we can efficiently handle S'' with objective function value at least $\nu_2 \doteq (\min_{i \geq 1} g(t(3.75)^i, iB, 3.75)) \cdot W(S'')/2$. By (Q), $\nu_2 \geq K_2' W(S'')/t^{1/\lfloor B \rfloor}$, where, e.g., $K_2' \geq (1/2 - 1/(2e))/[4e^2(3.75)^3]$. Finally we choose one of these two routings, depending on which of ν_1 and ν_2 is bigger. Since $W(S') + W(S'') = y^*$ by (P2), $\max\{W(S'), W(S'')\} \geq y^*/2$; so, $\max\{\nu_1, \nu_2\} \geq K_2 y^*/t^{1/\lfloor B \rfloor}$.

(ii) For a sufficiently small constant $K_3' \in (0, 1)$, set $\lambda = 1 + K_3'\epsilon$, $\beta = 1 - K_3'\epsilon$, $\kappa = (K_3'\epsilon)^3/\lambda$, and $\alpha = \kappa^{-1}$. Divide all the demands and capacities by $\alpha\lambda$; t and B become $t' = t/(\alpha\lambda)$ and $B' = B/(\alpha\lambda)$ respectively. (We will choose K_3 large enough such that $B' \geq 1$.) Since $\rho_u \leq 1/(\alpha\lambda)$ for all u now, we have $S'' = [k]$ in the notation of Lemma 3. By Lemma 2(ii) and Lemma 3, we can efficiently construct a path-selection of value $(\beta \min_{i \geq 1} g(\alpha t' \lambda^i, \alpha\kappa i B', \lambda)) \cdot y^*$. For $i \geq 1$, Lemma 2(ii) shows that $g(\alpha t' \lambda^i, \alpha\kappa i B', \lambda)$, i.e., $g(t\lambda^{i-1}, \kappa i B \lambda^{-1}, \lambda)$, equals $(1 - 1/(\lfloor \kappa i B \lambda^{-1} \rfloor + 1) - 1/(2\gamma))/(\lambda\gamma)$, where $\gamma = e(et\lambda^i(\lfloor \kappa i B \lambda^{-1} \rfloor + 1))^{1/\lfloor \kappa i B \lambda^{-1} \rfloor}$. If K_3 is large enough and K_3' small enough (e.g., if $K_3' \leq 1/18$ and $K_3 \geq 10(K_3')^{-3}$), then $\beta \cdot g(\alpha t' \lambda^i, \alpha\kappa i B', \lambda) \geq (1-\epsilon)(1 - 1/(2e))/e$. □

Part (i) of Theorem 4 and the reduction from CPIPs to LCRPs, show that for CPIPs, we can efficiently compute feasible solutions of value at least $K_2 y^*/t^{1/\lfloor B \rfloor}$.

Acknowledgements. We thank Stavros Kolliopoulos and Cliff Stein for valuable discussions; some of the ideas from [5] have provided much of the motivation for our results. Thanks also to Rajan Batta, Chung-Piaw Teo and Sushil Verma for helpful discussions. We also thank a reviewer for her/his helpful comments.

References

[1] A. Baveja, A. Srinivasan, Approximation algorithms for disjoint paths and related routing and packing problems, to appear in Mathematics of Operations Research.

[2] J. Håstad, Clique is hard to approximate within $n^{1-\epsilon}$, Proc. IEEE Symposium on Foundations of Computer Science, 1996, pp. 627–636.

[3] W. Hoeffding, Probability inequalities for sums of bounded random variables, American Statistical Association Journal 58 (1953), 13–30.

[4] J. Kleinberg, Approximation algorithms for disjoint paths problems, Ph.D. Dissertation, Department of EECS, Massachusetts Institute of Technology, Cambridge, MA, 1996.

[5] S. G. Kolliopoulos, C. Stein, Approximating disjoint-path problems using greedy algorithms and packing integer programs, Proc. MPS Conference on Integer Programming and Combinatorial Optimization, 1998, pp. 153–168, Lecture Notes in Computer Science 1412, Springer-Verlag.