

FORTRAN SUBROUTINES FOR NETWORK FLOW OPTIMIZATION USING AN INTERIOR POINT ALGORITHM

J. PATRÍCIO, L.F. PORTUGAL, M.G.C. RESENDE, G.VEIGA, AND J.J. JÚDICE

ABSTRACT. We describe FORTRAN subroutines for network flow optimization using an interior point network flow algorithm. We provide FORTRAN and C language drivers, as well as C language functions that, together with the subroutines, make up PDNET (Portugal, Resende, Veiga, and Júdice, 2000). The algorithm is described in detail and its implementation is outlined. Usage of the package is described and some computational experiments are reported. Source code for the software can be downloaded at <http://www.research.att.com/~mgcr/pdnet>.

1. INTRODUCTION

Given a directed graph $G = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} is a set of m vertices and \mathcal{E} a set of n edges, let (i, j) denote a directed edge from vertex i to vertex j . The minimum cost network flow problem can be formulated as

$$\min \sum_{(i,j) \in \mathcal{E}} c_{ij} x_{ij}$$

subject to:

$$(1) \quad \sum_{(i,j) \in \mathcal{E}} x_{ij} - \sum_{(j,i) \in \mathcal{E}} x_{ji} = b_i, \quad i \in \mathcal{V}$$

$$(2) \quad l_{ij} \leq x_{ij} \leq u_{ij}, \quad (i, j) \in \mathcal{E},$$

where x_{ij} denotes the flow on edge (i, j) and c_{ij} is the cost of passing one unit of flow on edge (i, j) . For each vertex $i \in \mathcal{V}$, b_i denotes the flow produced or consumed at vertex i . If $b_i > 0$, vertex b_i is a source. If $b_i < 0$, vertex i is a sink. Otherwise ($b_i = 0$), vertex i is a transshipment vertex. For each edge $(i, j) \in \mathcal{E}$, l_{ij} (u_{ij}) denote the lower (upper) bound on flow on edge (i, j) . Most often, the problem data are assumed to be integer.

In matrix notation, the above network flow problem can be formulated as a primal linear program of the form

$$(3) \quad \min \{c^\top x \mid Ax = b; x + s = u; x, s \geq 0\},$$

where A is the $m \times n$ *vertex-edge incidence matrix* of the graph $G = (\mathcal{V}, \mathcal{E})$, i.e. for each edge (i, j) in \mathcal{E} there is an associated column in matrix A with exactly two nonzero entries: an entry 1 in row i and an entry -1 in row j , b , x , and u are

Date: March 13, 2004.

AT&T Labs Research Technical Report TD-5X2SLN.

defined as above, and s is an n -dimensional vector of upper bound slacks. The dual of (3) is

$$(4) \quad \max \{b^\top y - u^\top w \mid A^\top y - w + z = c; z, w \geq 0\},$$

where y is the m -dimensional vector of dual variables and w and z are n -dimensional vectors of dual slacks.

If graph G is disconnected and has p connected components, there are exactly p redundant flow conservation constraints, which are sometimes removed from the problem formulation. Without loss of generality, we rule out trivially infeasible problems by assuming

$$\sum_{j \in \mathcal{V}^k} b_j = 0, \quad k = 1, \dots, p,$$

where \mathcal{V}^k is the set of vertices for the k -th component of G .

If it is required that the flow x_{ij} be integer, (2) is replaced with

$$l_{ij} \leq x_{ij} \leq u_{ij}, \quad x_{ij} \text{ integer}, \quad (i, j) \in \mathcal{E}.$$

In the description of the algorithm, we assume without loss of generality, that $l_{ij} = 0$ for all $(i, j) \in \mathcal{E}$ and that $c \neq 0$. A simple change of variables is done in the subroutines to transform the original problem into an equivalent one with $l_{ij} = 0$ for all $(i, j) \in \mathcal{E}$. The flow is transformed back to the original problem upon termination. The case where $c = 0$ is a simple feasibility problem, and is handled by solving a maximum flow problem [1].

Before concluding this introduction, we present some notation and outline the remainder of the paper. We denote the i -th column of A by A_i , the i -th row of A by A_i and a submatrix of A formed by columns with indices in set S by A_S . Let $x \in \mathbb{R}^n$. We denote by X the $n \times n$ diagonal matrix having the elements of x in the diagonal. The Euclidean or 2-norm is denoted by $\|\cdot\|$.

This paper describes Fortran subroutines used in an implementation of PDNET, an interior point network flow method introduced in Portugal, Resende, Veiga, and Júdice [12]. The paper is organized as follows. In Section 2 we review the truncated infeasible-primal feasible-dual interior point method for linear programming. The implementation of this algorithm to handle network flow problems is described in Section 3. Section 4 describes the subroutines and their usage. Computational results, comparing PDNET with the network optimizer in CPLEX 8.1, are reported in Section 5. Concluding remarks are made in Section 6.

2. TRUNCATED PRIMAL-INFEASIBLE DUAL-FEASIBLE INTERIOR POINT ALGORITHM

In this section, we recall the interior point algorithm implemented in PDNET. Let

$$\begin{aligned} \mathcal{Q}_+ &= \{(x, y, s, w, z) \in \mathbb{R}^{n+m+n+n+n} : x > 0, s > 0, w > 0, z > 0\}, \\ \mathcal{S}_+ &= \{(x, y, s, w, z) \in \mathcal{Q}_+ : x + s = u, A^\top y - w + z = c\}. \end{aligned}$$

The truncated primal-infeasible dual-feasible interior point (TPIDF) algorithm [12] starts with any solution $(x^0, y^0, s^0, w^0, z^0) \in \mathcal{S}_+$. At iteration k , the Newton direction $(\Delta x^k, \Delta y^k, \Delta s^k, \Delta w^k, \Delta z^k)$ is obtained as the solution of the linear system of

equations

$$\begin{aligned}
 A\Delta x^k &= b - Ax^k + r^k, \\
 \Delta x^k + \Delta s^k &= 0, \\
 A^\top \Delta y^k - \Delta w^k + \Delta z^k &= 0, \\
 Z^k \Delta x^k + X^k \Delta z^k &= \mu_k e - X^k Z^k e, \\
 W^k \Delta s^k + S^k \Delta w^k &= \mu_k e - W^k S^k e,
 \end{aligned}
 \tag{5}$$

where r^k is such that

$$\|r^k\| \leq \beta_0 \|Ax^k - b\|, \quad 0 \leq \beta_0 < \beta_1.
 \tag{6}$$

Primal and dual steps are taken in the direction $(\Delta x^k, \Delta y^k, \Delta s^k, \Delta w^k, \Delta z^k)$ to compute new iterates according to

$$\begin{aligned}
 x^{k+1} &= x^k + \alpha_p \Delta x^k, \\
 s^{k+1} &= s^k + \alpha_p \Delta s^k, \\
 y^{k+1} &= y^k + \alpha_d \Delta y^k, \\
 w^{k+1} &= w^k + \alpha_d \Delta w^k, \\
 z^{k+1} &= z^k + \alpha_d \Delta z^k,
 \end{aligned}
 \tag{7}$$

where α_p and α_d are stepsizes in the primal and dual spaces, respectively.

The solution of the linear system (5) is obtained in two steps. First, we compute the Δy^k component of the direction as the solution of the system of normal equations

$$\begin{aligned}
 A\Theta^k A^\top \Delta y^k &= -A\Theta^k (\mu_k (X^k)^{-1} e - \mu_k (S^k)^{-1} e - c + A^\top y^k) + \\
 &\quad (b - Ax^k),
 \end{aligned}
 \tag{8}$$

where Θ_k is given by

$$\Theta^k = (Z^k (X^k)^{-1} + W^k (S^k)^{-1})^{-1}.
 \tag{9}$$

The remaining components of the direction are then recovered by

$$\begin{aligned}
 \Delta x^k &= \Theta^k A^\top \Delta y^k + \Theta^k (\mu_k (X^k)^{-1} e - \mu_k (S^k)^{-1} e - c + A^\top y^k), \\
 \Delta s^k &= -\Delta x^k, \\
 \Delta z^k &= -z^k + \mu_k (X^k)^{-1} e - Z^k (X^k)^{-1} \Delta x^k, \\
 \Delta w^k &= -w^k + \mu_k (S^k)^{-1} e - W^k (S^k)^{-1} \Delta s^k.
 \end{aligned}
 \tag{10}$$

Parameter μ_k is computed as

$$\mu_k = \beta_1 \frac{(x^k)^\top z^k + (w^k)^\top s^k}{2n}
 \tag{11}$$

with $\beta_1 = 0.1$ and the stepsizes are computed as

$$\begin{aligned}
 \alpha_p &= \varrho_p \max\{\alpha : x^k + \alpha \Delta x^k \geq 0, s^k + \alpha \Delta s^k \geq 0\}, \\
 \alpha_d &= \varrho_d \max\{\alpha : w^k + \alpha \Delta w^k \geq 0, z^k + \alpha \Delta z^k \geq 0\},
 \end{aligned}
 \tag{12}$$

where $\varrho_p = \varrho_d = 0.9995$.

```

procedure pcg( $A, \Theta^k, \bar{b}, \epsilon_{cg}, \Delta y$ )
1   $\Delta y_0 := \Delta y;$ 
2   $r_0 := \bar{b} - A\Theta^k A^\top \Delta y_0;$ 
3   $z_0 := M^{-1}r_0;$ 
4   $p_0 := z_0;$ 
5   $i := 0;$ 
6  do stopping criterion not satisfied  $\rightarrow$ 
7       $q_i := A\Theta^k A^\top p_i;$ 
8       $\alpha_i := z_i^\top r_i / p_i^\top q_i;$ 
9       $\Delta y_{i+1} := \Delta y_i + \alpha_i p_i;$ 
10      $r_{i+1} := r_i - \alpha_i q_i;$ 
11      $z_{i+1} := M^{-1}r_{i+1};$ 
12      $\beta_i := z_{i+1}^\top r_{i+1} / z_i^\top r_i;$ 
13      $p_{i+1} := z_{i+1} + \beta_i p_i;$ 
14      $i := i + 1$ 
15 od;
16  $\Delta y := \Delta y_i$ 
end pcg;

```

FIGURE 1. The preconditioned conjugate gradient algorithm

3. IMPLEMENTATION

We discuss how the truncated primal-infeasible dual-feasible algorithm can be used for solving network flow problems. For ease of discussion, we assume, without loss of generality, that the graph G is connected. However, disconnected graphs are handled by PDNET.

3.1. Computing the Newton direction. Since the exact solution of (8) can be computationally expensive, a preconditioned conjugate gradient (PCG) algorithm is used to compute approximately an interior point search direction at each iteration. The PCG solves the linear system

$$(13) \quad M^{-1}(A\Theta^k A^\top)\Delta y^k = M^{-1}\bar{b},$$

where M is a positive definite matrix and Θ^k is given by (9), and

$$\bar{b} = -A\Theta^k(\mu_k(X^k)^{-1}e - \mu_k(S^k)^{-1}e - c + A^\top y^k) + (b - Ax^k).$$

The aim is to make the preconditioned matrix

$$(14) \quad M^{-1}(A\Theta^k A^\top)$$

less ill-conditioned than $A\Theta^k A^\top$, and improve the efficiency of the conjugate gradient algorithm by reducing the number of iterations it takes to find a satisfiable direction.

Pseudo-code for the preconditioned conjugate gradient algorithm implemented in PDNET is presented in Figure 1. The matrix-vector multiplications in line 7 are of the form $A\Theta^k A^\top p_i$, and can be carried out without forming $A\Theta^k A^\top$ explicitly. PDNET uses as its initial direction Δy_0 the direction Δy produced in the previous call to the conjugate gradient algorithm, i.e. during the previous interior point iteration. The first time pcg is called, $\Delta y_0 = (0, \dots, 0)$.

The preconditioned residual is computed in lines 3 and 11 when the system of linear equations

$$(15) \quad Mz_{i+1} = r_{i+1},$$

is solved. PDNET uses primal-dual variants of the diagonal and spanning tree preconditioners described in [15, 16].

The diagonal preconditioner, $M = \text{diag}(A\Theta^k A^\top)$, can be constructed in $O(n)$ operations, and makes the computation of the preconditioned residual of the conjugate gradient possible with $O(m)$ divisions. This preconditioner has been shown to be effective during the initial interior point iterations [11, 14, 15, 16, 18].

In the spanning tree preconditioner [16], one identifies a maximal spanning tree of the graph G , using as weights the diagonal elements of the current scaling matrix,

$$w = \Theta^k e,$$

where e is a unit n -vector. An exact maximal spanning tree is computed with the Fibonacci heap variant of Prim's algorithm [13], as described in [1]. At the k -th interior point iteration, let $\mathcal{T}^k = \{t_1, \dots, t_q\}$ be the indices of the edges of the maximal spanning tree. The spanning tree preconditioner is

$$M = A_{\mathcal{T}^k} \Theta_{\mathcal{T}^k}^k A_{\mathcal{T}^k}^\top,$$

where

$$\Theta_{\mathcal{T}^k}^k = \text{diag}(\Theta_{t_1}^k, \dots, \Theta_{t_q}^k).$$

For simplicity of notation, we include in $A_{\mathcal{T}^k}$ the linear dependent rows corresponding to the redundant flow conservation constraints. At each conjugate gradient iteration, the preconditioned residual system

$$Mz_{i+1} = r_{i+1}$$

is solved with the variables corresponding to redundant constraints set to zero. Since $A_{\mathcal{T}^k}$ can be ordered into a block diagonal form with triangular diagonal blocks, then the preconditioned residuals can be computed in $O(m)$ operations.

A heuristic is used to select the preconditioner. The initial selection is the diagonal preconditioner, since it tends to outperform the other preconditioners during the initial interior point iterations. The number of conjugate gradients taken at each interior point iteration is monitored. If the number of conjugate gradient iterations exceeds $\sqrt{m}/4$, the current computation of the direction is discarded, and a new conjugate gradient computation is done with the spanning tree preconditioner. The diagonal preconditioner is not used again. The diagonal preconditioner is limited to at most 30 interior point iterations. If at iteration 30 the diagonal preconditioner is still in effect, at iteration 31 the spanning tree preconditioner is triggered. Also, as a safeguard, a hard limit of 1000 conjugate gradient iterations is imposed.

To determine when the approximate direction Δy_i produced by the conjugate gradient algorithm is satisfactory, one can compute the angle θ between

$$(A\Theta^k A^\top)\Delta y_i \text{ and } \bar{b}$$

and stop when $|1 - \cos \theta| < \epsilon_{cos}^k$, where ϵ_{cos}^k is the tolerance at interior point iteration k [8, 15]. PDNET initially uses $\epsilon_{cos}^0 = 10^{-3}$ and tightens the tolerance after each interior point iteration k , as follows:

$$\epsilon_{cos}^{k+1} = \epsilon_{cos}^k \times \Delta \epsilon_{cos},$$

where, in PDNET, $\Delta\epsilon_{cos} = 0.95$. The exact computation of

$$\cos \theta = \frac{|\bar{b}^\top (A\Theta^k A^\top)\Delta y_i|}{\|\bar{b}\| \cdot \|(A\Theta^k A^\top)\Delta y_i\|}$$

has the complexity of one conjugate gradient iteration and should not be carried out every conjugate gradient iteration. Since $(A\Theta^k A^\top)\Delta y_i$ is approximately equal to $\bar{b} - r_i$, where r_i is the estimate of the residual at the i -th conjugate gradient iteration, then

$$\cos \theta \approx \frac{|\bar{b}^\top (\bar{b} - r_i)|}{\|\bar{b}\| \cdot \|\bar{b} - r_i\|}.$$

Since, on network linear programs, the preconditioned conjugate gradient method finds good directions in few iterations, this estimate is quite accurate in practice. Since it is inexpensive, it is computed at each conjugate gradient iteration.

3.2. Stopping criteria for interior point method. In [15], two stopping criteria for the interior point method were used. The first, called the primal-basic (PB) stopping rule, uses the spanning tree computed for the tree preconditioner. If the network flow problem has a unique solution, the edges of the tree converge to the optimal basic sequence of the problem. Let \mathcal{T} be the index set of the edges of the tree, and define

$$\Omega^+ = \{i \in \{1, 2, \dots, n\} \setminus \mathcal{T} : x_i/z_i > s_i/w_i\}$$

to the index set of edges that are fixed to their upper bounds. If the solution $x_{\mathcal{T}}^*$ of the linear system

$$A_{\mathcal{T}}x_{\mathcal{T}}^* = b - \sum_{i \in \Omega^+} u_i A_i,$$

is such that $0 \leq x_{\mathcal{T}}^* \leq u$, then $x_{\mathcal{T}}^*$ is a feasible basic solution. Furthermore, if the data is integer, then $x_{\mathcal{T}}^*$ has only integer components. Optimality of $x_{\mathcal{T}}^*$ can be verified by computing a lower bound on the optimal objective function value. This can be done with a strategy introduced independently in [15] and [10, 17]. Denote by x_i^* the i -th component of $x_{\mathcal{T}}^*$, and let

$$\mathcal{F} = \{i \in \mathcal{T} : 0 < x_i^* < u_i\}.$$

A tentative optimal dual solution y^* (having a possibly better objective value than the current dual interior point solution y^k) can be found by orthogonally projecting y^k onto the supporting affine space of the dual face complementary to $x_{\mathcal{T}}^*$. In an attempt to preserve dual feasibility, we compute y^* as the solution of the least squares problem

$$\min_{y^* \in \mathbb{R}^m} \{\|y^* - y^k\| : A_{\mathcal{F}}^\top y^* = c_{\mathcal{F}}\}.$$

Resende and Veiga [15] describe a $O(m)$ operation procedure to compute this projection. A feasible dual solution (y^*, z^*, w^*) is built by adjusting the dual slacks. Let $\delta_i = c_i - A_{.i}^\top y^*$. Then,

$$w_i^* = \begin{cases} -\delta_i & \text{if } \delta_i < 0 \\ 0 & \text{otherwise} \end{cases} \quad z_i^* = \begin{cases} 0 & \text{if } \delta_i < 0 \\ \delta_i & \text{otherwise.} \end{cases}$$

If $c^\top x^* - b^\top y^* + u^\top w^* = 0$, then (x^*, s^*) and (y^*, w^*, z^*) are optimal primal and dual solutions, respectively. If the data is integer and $0 < c^\top x^* - b^\top y^* + u^\top w^* < 1$, (x^*, s^*) is a primal optimal (integer) solution.

To apply the second stopping procedure of [15], called the maximum flow (MF) stopping criterion, an indicator function to partition the edge set into active and inactive (fixed at upper or lower bounds) is needed. In PDNET, the indicator used is the so-called primal-dual indicator, studied by Gay [6] and El-Bakry, Tapia, and Zhang [5]. Let ξ be a small tolerance. Edge i is classified as inactive at its lower bound if

$$\frac{x_i}{z_i} < \xi \quad \text{and} \quad \frac{s_i}{w_i} > \xi^{-1}.$$

Edge i is classified as inactive at its upper bound if

$$\frac{x_i}{z_i} > \xi^{-1} \quad \text{and} \quad \frac{s_i}{w_i} < \xi.$$

The remaining edges are set active. In PDNET, ξ is initially set to 10^{-3} and this tolerance is tightened each time the MF test is triggered according to $\xi^{new} = \xi^{old} \times \Delta\xi$, where, in PDNET, $\Delta\xi = 0.95$.

We select a tentative optimal dual face \mathcal{F} as a maximum weighted spanning forest limited to the active edges as determined by the indicator. The edge weights used in PDNET are those of the scaling matrix Θ^k .

As in the PB indicator, we project the current dual interior solution y^k orthogonally onto \mathcal{F} . Once the projected dual solution y^* is at hand, we attempt to find a feasible flow x^* complementary to y^* . A refined tentative optimal face is selected by redefining the set of active edges as

$$\tilde{\mathcal{F}} = \{i \in \{1, 2, \dots, n\} : |c_i - A_i^\top y^*| < \epsilon_r\},$$

where ϵ_r is a small tolerance ($\epsilon_r = 10^{-8}$ in PDNET). The method attempts to build a primal feasible solution, x^* , complementary to the tentative dual optimal solution by setting the inactive edges to lower or upper bounds, i.e., for $i \in \{1, 2, \dots, n\} \setminus \tilde{\mathcal{F}}$,

$$x_i^* = \begin{cases} 0 & \text{if } i \in \Omega^- = \{j \in \{1, 2, \dots, n\} \setminus \tilde{\mathcal{F}} : c_j - A_j^\top y^* > 0\} \\ u_i & \text{if } i \in \Omega^+ = \{j \in \{1, 2, \dots, n\} \setminus \tilde{\mathcal{F}} : c_j - A_j^\top y^* < 0\}. \end{cases}$$

By considering only the active edges, a *restricted network* is built. Flow on this network must satisfy

$$(16) \quad \begin{aligned} A_{\tilde{\mathcal{F}}} x_{\tilde{\mathcal{F}}} &= \tilde{b} = b - \sum_{i \in \Omega^+} u_i A_i, \\ 0 &\leq x_i \leq u_i, \quad i \in \tilde{\mathcal{F}}. \end{aligned}$$

Clearly, from the flow balance constraints (16), if a feasible flow $x_{\tilde{\mathcal{F}}}^*$ for the restricted network exists, it defines, along with $x_{\Omega^+}^*$ and $x_{\Omega^-}^*$, a primal feasible solution complementary to y^* . A feasible flow for the restricted network can be determined by solving a maximum flow problem on the *augmented network* defined by the underlying graph $\tilde{G} = (\tilde{\mathcal{V}}, \tilde{\mathcal{E}})$, where

$$\tilde{\mathcal{V}} = \{\sigma\} \cup \{\pi\} \cup \mathcal{V} \quad \text{and} \quad \tilde{\mathcal{E}} = \Sigma \cup \Pi \cup \tilde{\mathcal{F}}.$$

In addition, for each edge $(i, j) \in \tilde{\mathcal{F}}$ there is an associated capacity u_{ij} . Let

$$\mathcal{V}^+ = \{i \in \mathcal{V} : \tilde{b}_i > 0\} \quad \text{and} \quad \mathcal{V}^- = \{i \in \mathcal{V} : \tilde{b}_i < 0\}.$$

The additional edges are such that $\Sigma = \{(\sigma, i) : i \in \mathcal{V}^+\}$, with associated capacity \tilde{b}_i for each edge (σ, i) , and $\Pi = \{(i, \pi) : i \in \mathcal{V}^-\}$, with associated capacity $-\tilde{b}_i$ for each edge (i, π) . It can be shown that if $\mathcal{M}_{\sigma, \pi}$ is the maximum flow value from σ to π , and \tilde{x} is a maximal flow on the augmented network, then $\mathcal{M}_{\sigma, \pi} = \sum_{i \in \mathcal{V}^+} \tilde{b}_i$ if

and only if $\tilde{x}_{\mathcal{F}}$ is a feasible flow for the restricted network [15]. Therefore, finding a feasible flow for the restricted network involves the solution of a maximum flow problem. Furthermore, if the data is integer, this feasible flow is integer, as we can select a maximum flow algorithm that provides an integer solution.

Since this stopping criterion involves the solution of a maximum flow problem, it should not be triggered until the interior point algorithm is near the optimal solution. The criterion is triggered at iteration k , when $\mu_k < \epsilon_\mu$ occurs for first time. The choice $\epsilon_\mu = 1$ used in PDNET is appropriate for the set of test problems considered here. In a more general purpose implementation, a scale invariant criterion is desirable. All subsequent iterations test this stopping rule. In PDNET, the implementation of Goldfarb and Grigoriadis [7] of Dinic's algorithm is used to solve the maximum flow problems.

3.3. Other implementation issues. To conclude this section, we make some remarks on other important implementation issues of the primal-infeasible, dual-feasible algorithm, namely the starting solution, the adjustment of parameter μ_k , and the primal and dual stepsizes.

Recall that the algorithm starts with any solution $\{x^0, s^0, y^0, w^0, z^0\}$ satisfying

$$(17) \quad x^0 > 0, s^0 > 0, w^0 > 0, z^0 > 0,$$

$$(18) \quad x^0 + s^0 = u,$$

and

$$(19) \quad A^\top y^0 - w^0 + z^0 = c,$$

but does not have to satisfy $Ax^0 = b$. Additionally, it is desirable that the initial point also satisfy the remaining equations that define the central path (5), i.e.

$$(20) \quad X^0 z^0 = \mu e \quad \text{and} \quad S^0 w^0 = \mu e,$$

for $\mu > 0$.

For $(i, j) \in \mathcal{E}$, let

$$\begin{aligned} x_{ij}^0 &= \nu_{ij} u_{ij}, \\ s_{ij}^0 &= (1 - \nu_{ij}) u_{ij}, \\ z_{ij}^0 &= \mu / (\nu_{ij} u_{ij}), \\ w_{ij}^0 &= \mu / ((1 - \nu_{ij}) u_{ij}), \end{aligned}$$

be the starting solution, where $0 < \nu_{ij} < 1$ and $\mu > 0$. It is easy to verify that this starting solution satisfies (17–18) as well as (20).

Condition (19) is satisfied if, for $(i, j) \in \mathcal{E}$, ν_{ij} is chosen as

$$\nu_{ij} = \begin{cases} \frac{1}{2} + \frac{\mu}{\vartheta_{ij} u_{ij}} - \sqrt{\frac{1}{4} + \left(\frac{\mu}{\vartheta_{ij} u_{ij}}\right)^2} & \text{if } \vartheta_{ij} > 0, \\ \frac{1}{2} + \frac{\mu}{\vartheta_{ij} u_{ij}} + \sqrt{\frac{1}{4} + \left(\frac{\mu}{\vartheta_{ij} u_{ij}}\right)^2} & \text{if } \vartheta_{ij} < 0, \\ \frac{1}{2} & \text{if } \vartheta_{ij} = 0, \end{cases}$$

where, for some initial guess y^0 of the dual vector y ,

$$\vartheta_{ij} = -y_i^0 + y_j^0 + c_{ij}.$$

In PDNET, we set the initial guess

$$y^0 = \frac{\max\{|c_{ij}| : (i, j) \in \mathcal{E}\}}{\max\{|b_i| : i \in \mathcal{V}\}} b$$

and parameter

$$\mu = 0.2 \max\{\vartheta_{ij} u_{ij} \mid (i, j) \in \mathcal{E}\}.$$

The primal-dual parameter has an initial value $\mu_0 = \beta_1 \mu$, where in PDNET $\beta_1 = 0.1$. Subsequently, for iterations $k \geq 1$, μ_k is computed as in (11).

The stepsize parameters ϱ_p and ϱ_d are both set to 0.995 throughout the iterations, slightly more conservative than as suggested by [9].

4. FORTRAN SUBROUTINES

The current implementation PDNET consists of a collection of core FORTRAN 77 [2] subroutines with additional subsidiary modules written in both FORTRAN 77 and ANSI C. The software distribution associated with this paper provides fully functional utilities by including FORTRAN 77 and ANSI C reference implementations for the main program, providing default parameter settings and supplying additional routines for data input and output functions. In this section, we describe the usage of the PDNET framework with comments on user extensions. Specific instructions for building and installing PDNET from the source code are provided with the software distribution. Table 1 lists all modules provided in the PDNET software distribution.

We adopted several design and programming style guidelines in implementing the PDNET routines in FORTRAN 77. We required that all arrays and variables be passed as subroutine arguments, avoiding the use of COMMON blocks. In addition, there is a single externally allocated work array used to store all data structures used inside the core subroutines. The resulting code is expected to compile and run without modification under most software development and computing environments. For ANSI C modules, external functions always have an alternate FORTRAN callable interface, with all arguments passed by reference.

TABLE 1. Source Code Files

file name	description
cdriver.c	ANSI C main program
dblas1.f	FORTRAN 77 Level 1 BLAS reference implementation
debug.h	ANSI C debug macros used in PDNET
dnsubs.f	FORTRAN 77 implementation of Dinic's Algorithm
fdriver.f	FORTRAN 77 main program
pddflt.f	FORTRAN 77 function setting parameter defaults
pderr.c	ANSI C error functions
pdmxfl.c	ANSI C interface for invoking maximum flow computation
pdnet.f	PDNET core subroutines
pdnet.h	ANSI C macros used in PDNET
pdout.f	PDNET report generator
pdrdmc.c	ANSI C functions for reading network data

4.1. PDNET core subroutines. The PDNET core subroutines, available in file `pdnet.f`, are invoked via a single interface provided by function `pdnet()`. Following the specifications listed in Table 2, the calling program must provide data via the input arguments and allocate the internal and output arrays appropriately as described in Subsection 4.3 in files `fdriver.f` and `cdriver.c`.

We provide reference implementations of the PDNET main program which also serve as guides for developing custom applications invoking the PDNET core subroutines. In Subsection 4.2, we discuss in detail the input and output routines used in the reference implementations. Subsection 4.4 discusses the setting of parameters in PDNET. In addition, the core subroutines call an external function for maximum flow computation, which is provided in a subsidiary module, with its interface discussed in Subsection 4.5

TABLE 2. Arguments for `pdnet()`

variable	size	type	status	description
<code>b</code>	<code>nn</code>	integer	input	netflow values for each node
<code>c</code>	<code>na</code>	integer	input	cost for each arc
<code>dinfo</code>	<code>34</code>	double precision	input	parameters and statistics
<code>endn</code>	<code>na</code>	integer	input	end node for each arc
<code>info</code>	<code>23</code>	integer	input	parameters and statistics
<code>l</code>	<code>na</code>	integer	input	lower bound of flow for each arc
<code>maxwrk</code>	<code>-</code>	integer	input	maximum size of workspace
<code>na</code>	<code>-</code>	integer	input	number of arcs
<code>nn</code>	<code>-</code>	integer	input	number of nodes
<code>optflo</code>	<code>na</code>	integer	output	optimal flow for each arc
<code>ptrwkd</code>	<code>maxwrk</code>	double precision	internal	Pointer to work space
<code>ptrwki</code>	<code>-</code>	integer	internal	alias pointer to work space
<code>strn</code>	<code>na</code>	integer	input	start node for each arc
<code>u</code>	<code>na</code>	integer	input	upper bound of flow for each arc

4.2. Data input. Programs invoking the PDNET subroutines must supply an instance for the minimum-cost flow problem. As presented in Table 3, the data structure includes the description of the underlying graph, node netflow values, arc costs, capacities and lower bounds. All node and arc attributes are integer valued.

TABLE 3. Network Data Structure

variable	dimension	type	description
<code>nn</code>	<code>-</code>	integer	number of nodes
<code>na</code>	<code>-</code>	integer	number of arcs
<code>strn</code>	<code>na</code>	integer	start node for each arc
<code>endn</code>	<code>na</code>	integer	end node for each arc
<code>b</code>	<code>nn</code>	integer	flow values for each node
<code>c</code>	<code>na</code>	integer	cost for each arc
<code>l</code>	<code>na</code>	integer	lower bound of flow for each arc
<code>u</code>	<code>na</code>	integer	upper bound of flow for each arc

The reference implementations of the main program PDNET reads network data from an input file by invoking functions available in the ANSI C module `pdrdmc.c`.

These functions build PDNET input data structures from data files in the DIMACS format [3] with instances of minimum cost flows, maximum flow, and assignment problems. As illustrated in the PDNET modules `cdriver.c` and `fdriver.f`, we provide specialized interfaces used in FORTRAN 77 programs and ANSI C programs using dynamic memory allocation.

4.3. Memory allocation. In PDNET, total memory utilization is carefully managed by assigning workspace pointers to each individual array passed as argument to internal PDNET functions. This process is mostly transparent to the users, except for the initial allocation of the workspace as a double precision array `ptrwkd`, with an integer alias `ptrwki` used to refer to internal integer arrays. In addition to the workspace array, input and output arrays also passed as arguments to function `pdnet()` must be allocated by the calling procedure.

In the FORTRAN 77 reference implementation of the PDNET main program, available in file `fdriver.f`, we initially set the maximum problem and workspace sizes by specifying the FORTRAN compile-time parameters in Table 4. As presented in Tables 3 and 5, the external arrays passed to the PDNET core routines must be statically dimensioned using the FORTRAN parameters for the maximum number of nodes and arcs. The workspace array `ptrwkd` is similarly allocated, with an integer alias `ptrwki` assigned via an `equivalence` statement.

TABLE 4. FORTRAN Parameters

parameter	description
<code>maxa</code>	maximum number of arcs
<code>maxn</code>	maximum number of nodes
<code>maxwrk</code>	maximum size for double precision workspace array

TABLE 5. External Arrays

variable	dimension	type	description
<code>dinfo</code>	34	double precision	parameters and statistics
<code>info</code>	23	integer	parameters and statistics
<code>optflo</code>	<code>na</code>	integer	optimal flow for each arc
<code>ptrwkd</code>	<code>maxwrk</code>	double precision	pointer to work space
<code>ptrwki</code>	-	integer	alias pointer to work space

For implementations using dynamic memory allocation, arrays can be dimensioned to exact sizes at run time. In our ANSI C reference implementation available in module `cdriver.c`, we use a capability present in the network data reading module described in Subsection 4.2. By invoking function `pdrdmm`, memory is allocated for the network data structure arrays in Table 3. Other external arrays in Table 5 are explicitly allocated in the main program. Note that the pointer to work space array needs an integer alias `ptrwki` passed as an argument to the core subroutines.

4.4. Parameter setting. PDNET execution is controlled by a variety of parameters, selecting features of the underlying algorithm and the interface with the calling program. A subset of these parameters is exposed to the calling procedure via a

TABLE 6. List of runtime integer parameters assigned to vector `info`

component	variable	description
1	<code>bound</code>	used for building maximum spanning tree
2	<code>iprcnd</code>	preconditioner used
3	<code>iout</code>	output file unit (6: standard output)
4	<code>maxit</code>	maximum number of primal–dual iterations
5	<code>mxcgit</code>	maximum number of PCG iterations
6	<code>mxfv</code>	vertices on optimal flow network
7	<code>mxfe</code>	edges on optimal flow network
8	<code>nbound</code>	used for computing additional data structures
9	<code>nbuk</code>	number of buckets
10	<code>pduit</code>	number of primal–dual iterations performed
11	<code>root</code>	sets beginning of structure
12	<code>opt</code>	flags optimal flow
13	<code>pcgit</code>	number of PCG iterations performed
14	<code>pf</code>	iteration level
15	<code>ierror</code>	flags errors in input
16	<code>cgstop</code>	flags how PCG stopped
17	<code>mfstat</code>	maxflow status
18	<code>ioflag</code>	controls the output level of the summary
19	<code>prttyp</code>	printout level
20	<code>optgap</code>	duality gap optimality indicator flag
21	<code>optmf</code>	maximum flow optimality indicator flag
22	<code>optst</code>	spanning tree optimality indicator flag
23	<code>mr</code>	estimate of maximum number of iterations

PDNET core subroutine. The remaining parameters are set at compile time with default values assigned in inside module `pddf1t.f`.

The run time parameters listed in Tables 6 and 7 are set with FORTRAN 77 function `pdsprm()`. The integer parameters are assigned to components of vector `info` and double precision parameters are assigned to vector `dinfo`. In ANSI C reference implementation of the PDNET driver program, parameters can be specified as command line arguments. In addition, parameter values can be read from a specification file by invoking function `pdspec()` in the calling program.

4.5. Maximum flow computation. PDNET includes an interface to an external maximum flow computation module, called `pdmxfl`, which was coded in ANSI C, featuring the implementation of Goldfarb and Grigoriadis [7] of Dinic’s algorithm. This interface is used to check the maximum flow stopping criterion, as well as by the FORTRAN 77 module `pdckfs()` that is called after reading the network file to compute a maximum flow on the network, therefore checking for infeasibility.

4.6. Running PDNET. Module `fdriver` (`cdriver`) inputs the network structure and passes control to module `pdnet()`. This subroutine starts by assigning working vectors to pointers, with `pdpntr()`. The control parameters are then read from the vector `info`, with `pdginf()`. Correctness of the input is checked with `pdchkp()`, and data structure created with `pdmked()`. Internal parameters for methods are set with `pddf1t()`, and additional structures are created with `pddstr()`. Subroutines `pdtrfm()` and `pdprtb()` are called in order to shift the lower bounds to zero and to transform the data into double precision, respectively. Subroutines `pdpdat()`

TABLE 7. List of runtime double precision parameters assigned to vector `dinfo`

component	variable	description
1	<code>dobj</code>	dual objective function value
2	<code>dof</code>	dual objective function value on tree
3	<code>factor</code>	factor for Newton step
4	<code>fcttol</code>	factorization zero tolerance
5	<code>gap</code>	value of the gap
6	<code>grho0</code>	factor for rho0 update
7	<code>gtolcg</code>	factor for PCG tolerance
8	<code>gtolmf</code>	update factor for maxflow tolerances
9	<code>mfdopt</code>	maxflow value
10	<code>miu</code>	value of interior-point μ parameter
11	<code>mntlcg</code>	lower bound for PCG tolerance
12	<code>mrho0</code>	lower bound for rho0
13	<code>oldtol</code>	value of residual in previous PCG iteration
14	<code>pcgres</code>	value of residual in current PCG iteration
15	<code>pobj</code>	objective function value
16	<code>pof</code>	primal objective function value on tree
17	<code>rho0</code>	parameter from IQRD preconditioner
18	<code>s1fctr</code>	factor for miu
19	<code>stptol</code>	tolerance for dual slacks
20	<code>stpval</code>	value of largest slack on dual face
21	<code>tol1</code>	lower bound of tolerance for maxflow
22	<code>tol2</code>	upper bound of tolerance for maxflow
23	<code>tolcg</code>	tolerance for PCG stopping criterion
24	<code>tolcg0</code>	guess for tolerance for PCG stopping criterion
25	<code>tolslk</code>	zero for dual slacks
26	<code>huge</code>	largest real number
27	<code>tolpcg</code>	initial value for tolcg
28	<code>oldcos</code>	value of cosine in previous PCG iteration
29	<code>pcgcos</code>	value of cosine in current PCG iteration
30	<code>zero</code>	value for zero
31	<code>tolsw1</code>	maximum value for sw1
32	<code>tolsw2</code>	maximum value for sw2
33	<code>sw1</code>	iteration limit on PCG for diagonal to spanning tree switch
34	<code>sw2</code>	iteration limit on PCG for spanning tree to IQRD switch

and `pdckfs()` check the problem characteristics and inquire if the network has enough capacity to transport the proposed amount of commodity. The primal-dual main loop is then started and the right-hand-side of the search direction system is computed by `pcrhs()`. The maximum spanning tree is computed by `pdhprm()`, and its optimality is tested by `pdoptm()`. Under certain conditions ($\mu < 1$), the maxflow stopping criterion is invoked by subroutine `pdchmf()`. If required, preconditioner switch takes place, followed by a call to `pdpcgd()` to solve the Newton direction linear system using the chosen preconditioner. A summary of the iteration is printed by `pdsout()`. Primal and dual updates are made by `pdups1()` and stopping criteria check takes place, before returning to the start of the iteration loop.




FIGURE 2. Example of a minimum cost flow network.

```

p min 4 5
n 1 2
n 2 -2
n 3 -4
n 4 4
a 1 2 0 10 3
a 2 4 0 10 -7
a 4 3 0 10 1
a 3 1 0 10 -4
a 2 3 0 10 2

```

FIGURE 3. File `test.min` with the DIMACS format representation of the problem of Figure 2.

We now present a usage example. Let us consider the problem of finding the minimum cost flow on the network represented in Figure 2. In this figure, the integers close to the nodes represent the node's produced (positive value) or consumed (negative value) net flows, and the integers close to the arcs stand for their unit cost. Furthermore the capacity of each arc is bounded between 0 and 10. In Figure 3, we show the DIMACS format representation of this problem, stored in file `test.min`. Furthermore, Figures 4 and 5 show the beginning and the end of the printout given by PDNET, respectively.

5. COMPUTATIONAL RESULTS

A preliminary version of PDNET was tested extensively with results reported in [12]. In this section, we report on a limited experiment with the version of PDNET in the software distribution.

The experiments were done on a PC with an Intel Pentium II processor running at 300 MHz and 256 Mb of main memory. The operating system is SuSE Linux Professional 8.2 (kernel 2.4.20-4Gb). The code was compiled on the Intel Fortran and C compilers version 7.1 using the flags `-O3 -tpp7 -xM -ip`. CPU times in seconds were computed by calling the system routine `etime()`. The test problems are instances of the classes `mesh`, `grid` and `netgen_lo` of minimum-cost network flow problems, taken from the First DIMACS Implementation Challenge [4]. The

```

-----
PDNET 1.0      Release date: 2002-02-01
Primal-feasible dual-infeasible interior point network flow method

Problem Size and Memory Requirements:

      Number of nodes:      4  Maximum number of nodes:      10000
      Number of arcs:      5  Maximum number of arcs:      200000
      Size of work array:   222 Maximum work array:      4000000

-----
Runtime Parameters

      CG Preconditioner:    3      Output level:      1
      Maximum iterations:  1000   Maximum CG iterations: 1000
      Duality Gap Optimality: 1      Max. Flow Optimality: 1
      Span. Tree Optimality: 1

-----
Compile-time Parameters

      Output file unit:      6
      Spanning tree root:    1  Number of sort buckets:      5
      Factor for Newton step: 0.900000000E+00 Number of sort buckets: 0.70064923E-44

CONV>   itr:      1      mu: 0.17152671E+01

COST>   pobj: -.35209131E+02  dobj: -.69456608E+02  pd-gap: 0.42881677E+02

CG>   precnd: SPAN      cgstop: COS      pcgitr:      3
CG>   pcgtol: 0.16061519E+01  pcgres: 0.47184479E-15
CG>   cgtol: 0.10000000E-02  pcgcos: 0.22204460E-15

TREE> pb-feas: TRUE      pb-opt: FALSE
TREE>   pobj: -.32000000E+02  dobj: -.37683242E+02

MFLOW> mfstat: INACTIVE

CONV>   itr:      2      mu: 0.57549545E+00

COST>   pobj: -.25487418E+02  dobj: -.45470504E+02  pd-gap: 0.14387386E+02

CG>   precnd: SPAN      cgstop: COS      pcgitr:      2
CG>   pcgtol: 0.34821594E+00  pcgres: 0.51611112E-01
CG>   cgtol: 0.10000000E-02  pcgcos: 0.54634098E-04

TREE> pb-feas: TRUE      pb-opt: FALSE
TREE>   pobj: -.32000000E+02  dobj: -.35834791E+02

MFLOW> mxfv:      6      mxfe:      7
MFLOW> pdlotol: 0.95000000E-02  pdhitol: 0.10526316E+03  dobj: -.70000000E+02
MFLOW> mfstat: SUBOPTIMAL
-----

```

FIGURE 4. Startup and first two iterations of a typical PDNET session.

specifications of the `mesh` instances generated are presented in Table 8. The specifications used in the `GRIDGEN` generator to build the `grid` problems are displayed in Table 9. Finally, the instances of the test set `netgen_lo` were generated according to the guidelines stated in the computational study of Resende and Veiga [15], using the `NETGEN` generator following the specifications presented in Table 10. All these generators can be downloaded from the FTP site `dimacs.rutgers.edu`.

TABLE 8. Specifications for instances in class `mesh`

Problem	n	m	maxcap	maxcost	seed	x	y	xdeg	ydeg
1	256	4096	16384	4096	270001	36	7	9	5
2	2048	32768	16384	4096	270001	146	14	9	5
3	8192	131072	16384	4096	270001	409	20	9	5

```

-----
CONV>   itr:           8      miu: 0.54012752E-03

COST>   pobj: -.32000034E+02   dobj: -.32011037E+02   pd-gap: 0.13503188E-01

CG>   precnd: SPAN           cgstop: COS           pcgitr:           1
CG>   pcgtol: 0.16712951E-03   pcgres: 0.62266830E-04
CG>   cgtol: 0.97500000E+00   pcgcos: 0.10817880E-09

TREE>   pb-feas: TRUE         pb-opt: TRUE
TREE>   pof: -.32000000E+02   dof: -.32003820E+02

MFLOW>   mxfv:           6      mxfe:           7
MFLOW>   pdlotol: 0.69833730E-02   pdhitol: 0.14319728E+03   dobj: -.32000000E+02
MFLOW>   mfstat: OPTIMAL

-----
PDNET 1.0: Primal-feasible, dual-infeasible interior point network flow method
OPTIMAL primal flow

      from   to   arc   lower   upper   arc
      node  node cost  bound  bound  flow
-----
      1     1     2     3       0     10     8
      2     2     4    -7       0     10     6
      3     4     3     1       0     10    10
      4     3     1    -4       0     10     6
-----

```

FIGURE 5. Last iteration and results printout in typical PDNET session.

TABLE 9. Specifications for instances in class `grid`

Problem	nodes	grid size	sources	sinks	average degree	supply	arc costs		arc caps	
							min	max	min	max
1	101	10 × 10	50	50	16	100000	0	2000	20	2000
2	901	30 × 30	450	450	16	100000	0	2000	20	2000
3	2501	50 × 50	1250	1250	16	100000	0	2000	20	2000

TABLE 10. NETGEN specification file for class `netgen_lo`

<code>seed</code>	Random number seed:	27001
<code>problem</code>	Problem number (for output):	1
<code>nodes</code>	Number of nodes:	$m = 2^x$
<code>sources</code>	Number of sources:	2^{x-2}
<code>sinks</code>	Number of sinks:	2^{x-2}
<code>density</code>	Number of (requested) arcs:	2^{x+3}
<code>mincost</code>	Minimum arc cost:	0
<code>maxcost</code>	Maximum arc cost:	4096
<code>supply</code>	Total supply:	$2^{2(x-2)}$
<code>tsources</code>	Transshipment sources:	0
<code>tsinks</code>	Transshipment sinks:	0
<code>hicost</code>	Skeleton arcs with max cost:	100%
<code>capacitated</code>	Capacitated arcs:	100%
<code>mincap</code>	Minimum arc capacity:	1
<code>maxcap</code>	Maximum arc capacity:	16

Instances of increasing dimension were considered. The three `netgen_lo` instances presented were generated considering $x = 9$, $x = 13$ and $x = 15$ respectively.

TABLE 11. Computational Experience with instances of the sets `mesh`, `grid` and `netgen_lo` of test problems

Set	Problem	$ \mathcal{V} $	$ \mathcal{E} $	CPLEX 8.1		PDNET	
				IT	CPU	IT	CPU
mesh	1	256	4096	8497	0.17	40	0.79
	2	2048	32768	110802	16.43	59	15.64
	3	8192	131072	554058	404.44	76	109.67
grid	1	101	1616	611	0.01	27	0.19
	2	901	14416	3492	0.20	35	3.30
	3	2501	40016	11933	1.37	44	19.98
netgen_lo	1	512	4102	6482	0.16	28	0.80
	2	8192	65709	171877	32.21	46	33.14
	3	32768	262896	888334	708.02	59	226.77

In Table 11, the number of iterations (IT) and the CPU time (CPU) of PDNET and CPLEX 8.1 are compared. The reported results show that PDNET tends to outperform CPLEX in the larger instances of the `mesh` and `netgen_lo` set problems, but fails to do so in `grid` set problems. For testset `netgen_lo` the difference is quite noticeable for the largest instance. We observed that in this problem about 90% of the CPU time and iteration count in CPLEX 8.1 was spent computing a feasible solution. The results in this experience and those reported in [12] show that this code is quite competitive with CPLEX and other efficient network flow codes for large-scale problems.

6. CONCLUDING REMARKS

In this paper, we describe a FORTRAN implementation of PDNET, a primal–infeasible dual–feasible interior point method for solving large–scale linear network flow problems. The subroutines are described, and directions for usage are given. A number of technical features of the implementation, which enable the user to control several aspects of the program execution, are also presented. Some computational experience with a number of test problems from the DIMACS collection is reported. These results illustrate the efficiency of PDNET for the solution of linear network flow problems.

Source code for the software is available for download at <http://www.research.att.com/~mgcr/pdnet>.

REFERENCES

- [1] N. K. AHUJA, T. L. MAGNANTI, AND J. B. ORLIN, *Network Flows*, Prentice Hall, Englewood Cliffs, NJ, 1993.
- [2] AMERICAN NATIONAL STANDARDS INSTITUTE, *American National Standard programming language FORTRAN: approved April 3, 1978*, American National Standards Institute, Inc.: ANSI X3.9-1978. Revision of ANSI X3.9-1966. World-Wide Web document., 1978. This is an online hypertext of the Fortran 77 Standard.
- [3] DIMACS, *The first DIMACS international algorithm implementation challenge: Problem definitions and specifications*. World-Wide Web document., 1991.
- [4] DIMACS, *The first DIMACS international algorithm implementation challenge: The benchmark experiments*, tech. rep., DIMACS, New Brunswick, NJ, 1991.

- [5] A. EL-BAKRY, R. TAPIA, AND Y. ZHANG, *A study on the use of indicators for identifying zero variables for interior-point methods*, SIAM Review, 36 (1994), pp. 45–72.
- [6] D. GAY, *Stopping tests that compute optimal solutions for interior-point linear programming algorithms*, tech. rep., AT&T Bell Laboratories, Murray Hill, NJ, 1989.
- [7] D. GOLDFARB AND M. GRIGORIADIS, *A computational comparison of the Dinic and network simplex methods for maximum flow*, Annals of Operations Research, 13 (1988), pp. 83–123.
- [8] N. KARMAKAR AND K. RAMAKRISHNAN, *Computational results of an interior point algorithm for large scale linear programming*, Mathematical Programming, 52 (1991), pp. 555–586.
- [9] K. MCSHANE, C. MONMA, AND D. SHANNO, *An implementation of a primal-dual interior point method for linear programming*, ORSA Journal on Computing, 1 (1989), pp. 70–83.
- [10] S. MEHROTRA AND Y. YE, *On finding the optimal facet of linear programs*, Mathematical Programming, (1993), pp. 497–516.
- [11] L. PORTUGAL, F. BASTOS, J. JÚDICE, J. PAIXÃO, AND T. TERLAKY, *An investigation of interior point algorithms for the linear transportation problem*, SIAM J. Sci. Computing, 17 (1996), pp. 1202–1223.
- [12] L. PORTUGAL, M. RESENDE, G. VEIGA, AND J. JÚDICE, *A truncated primal-infeasible dual-feasible network interior point method*, Networks, 35 (2000), pp. 91–108.
- [13] R. PRIM, *Shortest connection networks and some generalizations*, Bell System Technical Journal, 36 (1957), pp. 1389–1401.
- [14] M. RESENDE AND G. VEIGA, *Computing the projection in an interior point algorithm: An experimental comparison*, Investigación Operativa, 3 (1993), pp. 81–92.
- [15] ———, *An efficient implementation of a network interior point method*, in Network Flows and Matching: First DIMACS Implementation Challenge, D. S. Johnson and C. C. McGeoch, eds., vol. 12 of DIMACS Series in Discrete Mathematics and Theoretical Computer Science, American Mathematical Society, 1993, pp. 299–348.
- [16] ———, *An implementation of the dual affine scaling algorithm for minimum cost flow on bipartite uncapacitated networks*, SIAM Journal on Optimization, 3 (1993), pp. 516–537.
- [17] Y. YE, *On the finite convergence of interior-point algorithms for linear programming*, Mathematical Programming, 57 (1992), pp. 325–335.
- [18] Q.-J. YEH, *A reduced dual affine scaling algorithm for solving assignment and transportation problems*, PhD thesis, Columbia University, New York, NY, 1989.

(J. Patrício) DEPARTAMENTO DE MATEMÁTICA, ESCOLA SUPERIOR DE TECNOLOGIA DE TOMAR, 2300 TOMAR, PORTUGAL

E-mail address, J. Patrício: `Joao.Patricio@aim.estt.ipt.pt`

(L.F. Portugal) LUIS PORTUGAL PASSED AWAY ON JULY 21, 1998. AT THE TIME OF HIS DEATH, HE WAS AFFILIATED WITH DEPARTAMENTO DE CIÊNCIAS DA TERRA, UNIVERSIDADE DE COIMBRA, 3000 COIMBRA, PORTUGAL

(M.G.C. Resende) INTERNET AND NETWORK SYSTEMS RESEARCH CENTER, AT&T LABS RESEARCH, 180 PARK AVENUE, ROOM C241, FLORHAM PARK, NJ 07932 USA.

E-mail address, M.G.C. Resende: `mgcr@research.att.com`

(G.Veiga) HTF SOFTWARE, RIO DE JANEIRO, BRAZIL

E-mail address, G.Veiga: `gveiga@yahoo.com`

(J.J. Júdice) DEPARTAMENTO DE MATEMÁTICA, UNIVERSIDADE DE COIMBRA, 3000 COIMBRA, PORTUGAL

E-mail address, J.J. Júdice: `Joaquim.Judice@co.it.pt`