

VLSI routing and Lagrangian duality

Simon Sigurdhsson, <ssimon@student.chalmers.se>

Abstract This report discusses a solution to the first project of the Large scale optimization course (TMA521) given by the Applied mathematics department at Chalmers University of Technology. The problem considered is a routing problem in VLSI based on that discussed by Feo and Hochbaum (1986).

1 Introduction

Although the problem is defined and discussed by both Feo and Hochbaum (1986) and the project description, it will now be briefly presented to provide some context.

The problem consists of deciding whether it is possible to connect a number of components in the context of a two-layer board with horizontal wiring on one side and vertical on the other, and a specified number of connectors between these layers.

The problem is modelled mathematically as an ILP problem which is then Lagrangian relaxed resulting essentially in one cheapest route problem for each wanted connection.

2 Subgradient optimization

The first task of the project consists of implementing the subgradient optimization as MATLAB code. Routines that solve the lagrangian subproblems (the cheapest route problems) are given, and as such the code only has to reinterpret the output of those routines and perform the subgradient optimization.

Appendices A on page 7 shows the complete code listing for this task (sans the parts given in the problem description, *i.e.* `gsp.c`, `sph.c`, `visagrid.m` and the problem instance files), but a full description of the algorithm is given below.

2.1 The subgradient algorithm

Since the given function `gsp.c` solves the actual lagrangian subproblems, the implementation of the subgradient algorithm is very simple:

1. Call `gsp` to solve the subproblem, discard all paths that have a total cost of 1 or more, and transform the data into an x_{ijl} matrix. This is what the functions `okcom` and `getxij` (described in appendices A.1 to A.2 on page 10) do, respectively.
2. Calculate the dual value for the solution obtained in the current iteration,

$$h(\pi^t) = \sum_{i=1}^n \pi_i^t + \sum_{l=1}^k \left(x_{t_l s_l} - \sum_{i=1}^n \sum_{j=1}^n \pi_i^t x_{jil} \right).$$

3. Calculate the subgradient direction,

$$d_i^t = 1 - \sum_{l=1}^k \sum_{j=1}^n x_{jil}.$$

4. Calculate the step length,

$$s^t = \lambda^t \frac{h(\pi) - LBD}{\sum_{i=1}^n (d_i^t)^2}.$$

5. Update the dual variables by taking a step in the subgradient direction, *i.e.* set $\pi_i^{t+1} = \max\{0, \pi_i^t - s^t d_i^t\}$.
6. Finally, decrease lambda by setting $\lambda^{t+1} = 0.95\lambda^t$.
7. Repeat from step 1 unless the maximum number of iteration has been reached.

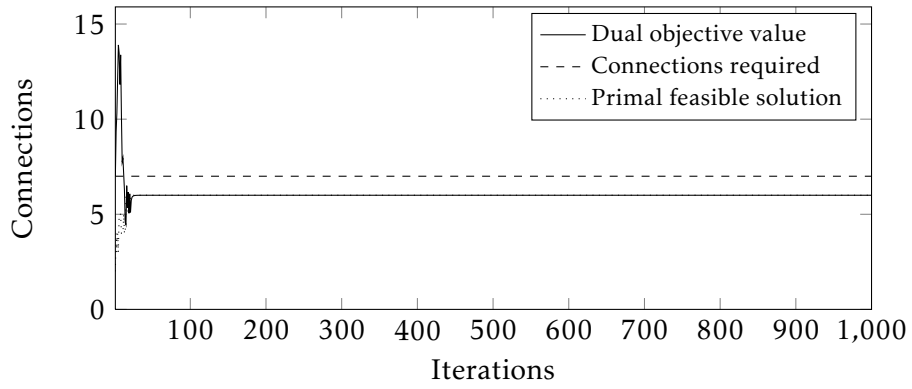
2.2 Results

As shown by figure 1 on the following page, the only problem instance that can be classified as either possible or impossible without using any primal feasibility heuristic is instance p6 (figure 1a on the next page), which has a dual objective value of approximately 6, meaning the optimal value of the problem (*i.e.* the maximum number of connections possible) is at most 6. Thus, the required number of connections (7) cannot be obtained.


The other two problems show dual objective values larger than the required number of connections, but this is no guarantee that the instance can have that many connections; the dual objective value is an optimistic bound.

3 A feasibility heuristic

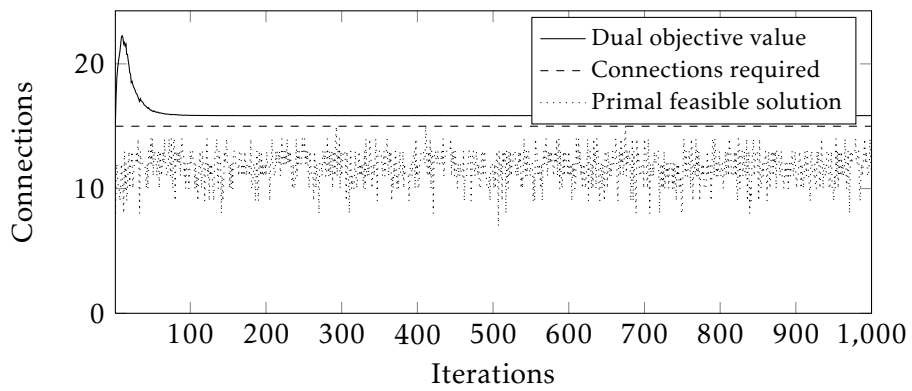
Implementing a heuristic that uses the problem formulation to reduce a dual solution to a feasible one might seem like a daunting task. However, thinking of the



(a) Instance p6.



(b) Instance p10.



(c) Instance p11.

Figure 1: The dual objective value and primal feasible solution for three problem instances.

subproblems as cheapest route problems, they only differ from feasible solutions of the original problem in that the paths are not required to be vertex-disjoint. Thus, removing or re-routing paths that cross in a dual solution will result in a feasible solution to the original problem, even though it may be a bad one that doesn't connect all pairs.

The heuristic implemented in this task uses that very idea. Starting with a dual solution, knowing that it is a collection of paths, the heuristic calculates the number of paths passing each node of the problem (*i.e.* $n_i = \sum_{l=1}^k z_{il}$, extending the notation used on page 4 of the project description). The algorithm then proceeds as follows:

1. Calculate n_i .
2. If $n_i < 2, \forall i$, abort the heuristic (since we evidently have a feasible solution).
3. Select a path p passing through any node with $n_i \geq 2$.
4. Let p_s and p_e be the first and last nodes of p , respectively. Remove p from the dual solution.
5. Find a cheapest path p' from p_s to p_e , with costs

$$c_i = \begin{cases} \pi_i, & \text{no path through node } i \\ \infty, & \text{otherwise.} \end{cases}$$

6. If $\text{cost}(p') < \infty$, add p' to the dual solution.
7. Repeat from step 1.

The algorithm will always terminate since it keeps removing paths from the overpopulated nodes either by finding an alternative route or by simply discarding the path. As such, the heuristic may (although this should be rare) terminate with a (still feasible) solution containing no paths at all. Generally it should be able to return at least one path, however.

The algorithm has a worst-case complexity $O(kn')$, where n' is the number of overpopulated nodes of the dual solution. Since the number of overpopulated nodes is less than the total number of nodes in the problem, we can rewrite the complexity as $O(k|\mathcal{V}|) = O(kn)$. The number of connections must also be less than $n/2$ due to the problem structure, and as such the complexity is $O(n^2)$, which is perfectly reasonable for a feasibility heuristic.

3.1 Results


As shown by figure 1 on the preceding page, the heuristic performs well for instance p6, finding a feasible solution that matches the dual objective value quite early in the subgradient algorithm. Instance p10 seems more difficult, and in that case the

heuristic fails to match the dual value and in fact provides no additional information (*i.e.* there is no way of telling if the actual optimum is above or below k). In instance p11, however, the heuristic finds a couple of feasible solutions which connect k pairs, which means we can conclude that the optimum value is at least k .


To summarize, the combination of Lagrangian relaxation, the subgradient algorithm and a feasibility heuristic indicates that p6 has no solution to the wiring problem, p11 has a solution to the wiring problem, and that p10 may have a solution to the wiring problem. Figure 2 on the next page shows the best feasible solution produced by the heuristic for each of these instances, and one can verify that all these solutions are in fact feasible by checking for overused nodes, and one can also see that they have 6, 14 and 15 paths respectively (where k is 7, 15 and 15 for the three problem instances).

References


Feo, T. A. and D. S. Hochbaum (Nov. 1986). 'Lagrangian Relaxation for Testing Infeasibility in VLSI Routing'. In: 34.6, pp. 819–831.



(a) Instance p6.



(b) Instance p10.



(c) Instance p11.

Figure 2: The best primal feasible solution for three problem instances.

A Program code

```
1 % TMA521 - Large scale optimization
2 % Spring 2013
3 % Project 1, tasks 1 & 2
4 % Simon Sigurdhsson
5
6 % This file is the main file that solves (after editing according to
7 % instructions in comments) tasks 1 and 2 of the project.
8
9 % Files have been organized such that given code (gsp.c, sph.c and
10 % visagrid.m) resides in a subfolder "given", and the instance files
11 % are in the "instances" subfolder. Hence, we need to add these to
12 % the PATH in order to use these files.
13 % Additionally, all variables are cleared and all figures closed at
14 % the start of the program, to avoid confusion from earlier results.
15 addpath('given','instances')
16 clear all; close all; clc;
17
18 % Initialization of variables
19 % This is where all non-local variables of the program are defined.
20 % First, we have a couple of variables that control the program flow,
21 % deciding what instance we are solving and how many iterations we
22 % should run the subgradient solver for.
23 p11; % Problem to solve.
24 maxIter = 1000; % Maximum number of subgradient iterations.
25 % Then, we initialize some of the variables used in the subgradient
26 % algorithm, so that we don't have to reallocate them inside the
27 % loop (this is very inefficient). All these variables are set to
28 % 0 initially, which makes for a good starting value for pi and will
29 % be overwritten for all other variables, except for the list of
30 % upper bounds, which is set to Inf (so that all calculated upper
31 % bounds are smaller), and lambda which is set to 2 in accordance
32 % with the project description.
33 pi = zeros(maxIter, dimX*dimY*2); % Lagrangian multipliers for each time step
34 d = zeros(maxIter, dimX*dimY*2); % Subgradient direction for each time step
35 s = zeros(maxIter, 1); % Step length for each time step
36 UBDS = ones(maxIter, 1)*Inf; % Upper bound for each time step
37 LBDS = zeros(size(UBDS)); % Lower bound for each time step
38 lambda = 2; % Step length modifier
39 optCom = []; % Storage for best feasible solution (com)
40 optN1 = []; % Storage for best feasible solution (n1)
41 optPi = []; % Storage for best feasible solution (pi)
42
43 % The subgradient scheme (solver)
44 % This is where the actual work begins. The only stopping criterion
45 % used is the maximum number of iterations, and the code inside the
46 % loop pretty much follows the flow chart given in the project
47 % description.
48 for t=1:maxIter
49 % 1. Solve the Lagrangian subproblems
50 % The Lagrangian subproblems are solved as cheapest-path problems
51 % by the given function gsp. Since the output of gsp is hard to
52 % handle, a function getxij (see getxij.m) is used to transform
```

```

53 % the output to two 3-dimensional matrices describing the problem
54 % variables  $x_{\{ij\}}$  and  $x_{\{t\}}s_{\{l\}}$ . But first, the okcom
55 % function (see okcom.m) eliminates the paths with cost larger
56 % than 1, as described in the project description.
57 nl = gsp(dimX,dimY,pi(t,:)',k,com); % Solve the subproblems
58 [ok, oknl] = okcom(pi(t,:),k,com,nl); % Which paths are "ok"?
59 kok = find(ismember(com, ok, 'rows') == 1); % Which rows of com are "ok"?
60 [xij, xt1sl] = getxij(dimX,dimY,k,com,nl,kok);% Transform to  $x_{\{ij\}}$  matrix
61
62 % 1.1. Calculate the primal feasibility heuristic (only for task 2)
63 % The heuristic, described in heuristic.m, creates a feasible
64 % solution based on the current solution. It returns data of the
65 % same structure as gsp, so it too requires getxij to transform
66 % the output into similar matrices.
67 [hcom, hn1] = heuristic(dimX, dimY, pi(t,:), k, com, nl); % Find primal feasible so
68 [hxij, hxt1sl] = getxij(dimX,dimY,k,hcom,hn1,(1:length(hcom))');% Transform to  $x_{\{ij\}}$  ma
69
70 % 2. Calculate upper bound, h(pi)
71 % The upper bound, essentially the Lagrangian dual value, is
72 % calculated as described in the project description. Since the
73 % output of gsp has been transformed into a matrix, the actual
74 % calculation is very similar to the mathematical formula shown
75 % in the project description.
76 lpi = repmat(pi(t,:)', [1 dimX*dimY*2 k]); % Expand pi, for element-wise
77 h = sum(pi(t,:)) + sum(xt1sl - sum(sum(lpi.*xij, 1), 2), 3);% Calculate h(pi) according t
78
79 % 2.1. Calculate the lower bound (only for task 2)
80 % The lower bound is given by the primal feasible solution found
81 % by the heuristic. Since it is a *primal* feasible solution, the
82 % lower bound isn't calculated using the Lagrangian dual value but
83 % using the original problem formulation. Thus, it is a bit simpler
84 % than the calculation of h(pi), but still analogous to the math
85 % described in the project description.
86 LBDS(t) = sum(hxt1sl, 3); % Count the number of connections made by the primal feasible s
87 % Note that since LBDS(t) = 0 for all t before this assignment is
88 % made, simply commenting the above line (along with the code
89 % under 1.1) will essentially solve task 1 instead of task 2.
90
91 % 3. Calculate subgradient direction d
92 % Again, since the output of gsp has been transformed into a
93 % suitable matrix, the calculation of the subgradient direction is
94 % very similar to the mathematical formula given by the project
95 % description.
96 sxj1 = sum(sum(xij,1),3); % Pre-calculate sum (for efficiency)
97 d(t,:) = ones(size(sxj1))-sxj1;% Calculate subgradient direction
98
99 % 4. Calculate step length s
100 % The step length calculation is also pretty much a verbatim
101 % translation of the mathematical formula into MATLAB code.
102 % Note that in task 2 we actually use the lower bound given by
103 % the heuristic in the formula, while LBDS=0 for task 1 as
104 % suggested by the project description.
105 s(t) = lambda*(h-LBDS(t))/sum(d(t,:).^2);% Calculate step length
106

```



```

107 % 5. Take step s in direction -d
108 % If this isn't the last iteration, the Lagrangian multipliers
109 % have to be updated for the next iteration by taking a step in
110 % the subgradient direction. Again (there's a pattern here, no?),
111 % the transformation of the gsp output makes this a verbatim
112 % translation of math into MATLAB code.
113 if t ~= maxIter % If this isn't the last it
114     pi(t+1,:) = max(zeros(size(pi(t,:))),pi(t,:)-s(t).*d(t,:));% Take a step in the subgra
115 end
116
117 % 6. Decrease value of lambda
118 % In accordance with the project description, lambda is decreased
119 % by 5% every iteration. This ensures 0<lambda<2.
120 lambda = lambda*0.95;
121
122 % 7. Save for plotting
123 % In order to present graphs showing the convergence of the
124 % subgradient algorithm along with the best solution found by
125 % the heuristic (for task 2), the data is saved in the
126 % preallocated variables.
127 if (LBDS(t) >= max(LBDS)) % && false % Uncomment "&& false" for task 1
128     optCom = hcom(any(hcom,2),:); % Save the best feasible "com"
129     optN1 = hn1; % Save the best feasible "n1"
130     optPi = pi; % Save the best feasible "pi"
131 end
132 UBDS(t) = h; % Save the upper bound of the current iteration
133
134 % 8. Progress bar (sort of)
135 % This simply prints some text every 25th iteration, to roughly
136 % indicate how far the algorithm has come and how long we have
137 % to wait before it will terminate.
138 if mod(t,25) == 0
139     disp(['Iteration_', num2str(t), '_of_', num2str(maxIter)])
140 end
141 end
142
143 % Plotting results
144 % The subgradient algorithm is done, and it is time to show the
145 % results. First, a plot showing the convergence of the subgradient
146 % algorithm and the heuristic (in task 2) is shown. Since all data
147 % has been saved, this is fairly standard and boring MATLAB code.
148 plot(1:maxIter, UBDS, 'k-'); hold on; % Plot the upper bounds
149 plot(1:maxIter, k*ones(size(UBDS)), 'k--'); % Indicate the number of connections we want
150 plot(1:maxIter, LBDS, 'k:'); % Plot the lower bounds (only useful in task 2)
151 axis([1 maxIter 0 max(k, max(UBDS))+2]); % Set sensible axes
152 xlabel('Iterations'); % Label the x axis
153 ylabel('Connections'); % Label the y axis
154 legend('Dual_objective_value', 'Connections_required', 'Primal_feasible_solution'); % Explain p
155 % Next (only for task 2, comment out if running task 1), the best
156 % primal feasible solution found by the heuristic is shown using
157 % the given visagrid function. Again, nothing strange happening.
158 figure; % Get a new figure
159 visagrid(dimX,dimY,optN1,optCom,optPi,25); % Show the feasible solution

```

A.1 The okcom function

```
1 % TMA521 - Large scale optimization
2 % Spring 2013
3 % Project 1, tasks 1 & 2
4 % Simon Sigurdhsson
5
6 function [ ok, newnl ] = okcom( pi, k, com, nl )
7 %OKCOM Eliminates paths from com/nl if their cost is large.
8 % The okcom function, which contains code from page 6 of
9 % the project description, calculates the cost of each path
10 % in com/nl, adding the path to ok/newnl if the cost is less
11 % than one.
12 last = 0;
13 ok = zeros(k,2);
14 newnl = [];
15 for i = 1:k
16     first = last+1;
17     slask = find(nl(last+1:length(nl)) == com(i,1));
18     last = slask(1)+first-1;
19     if sum(pi(nl(first:last))) < 1
20         ok(i,:) = com(i,:);
21         newnl = [newnl; nl(first:last)];
22     end
23 end
24 end
```

A.2 The getxij function

```
1 % TMA521 - Large scale optimization
2 % Spring 2013
3 % Project 1, tasks 1 & 2
4 % Simon Sigurdhsson
5
6 function [ xij, xt1sl ] = getxij( dimX, dimY, k, com, nl, kok )
7 %GETXIJ Calculate x_{ijl} and x_{t_{1}s_{1}l} matrices
8 % The getxij function takes the output of gsp and transforms it
9 % into the actual problem variables x_{ijl} and x_{t_{1}s_{1}l},
10 % making it much easier to calculate the dual value, subgradient
11 % direction and step length required by the subgradient algorithm.
12 % To begin with, a couple of local variables are defined.
13 maxij = dimX*dimY*2; % The number of nodes in the problem
14 xij = zeros(maxij, maxij, k); % Output matrix, preallocated
15 xt1sl = zeros(1, 1, k); % Output matrix, preallocated
16 tempnl = nl; % Copy of nl, will be modified in loop
17 % Now, for all the "ok" paths in com/nl, we set the appropriate
18 % elements of the output matrices to 1 (remember that they are
19 % initialized to 0).
20 for i=kok'
21     % First, the part of nl containing the ith path is found.
22     % It is assumed that the paths in nl have the same order as
23     % the corresponding path endpoints in com, but that the path
24     % is stored "backwards".
25     % First, we extract the endpoints.
```

```

26     sn = com(i,1); en = com(i,2);
27     % Then, we extract the path corresponding to those endpoints,
28     % assuming that the path is the next one in nl (i.e. the first
29     % one in tempnl, in which we discard each path after finding it).
30     thisnl = tempnl(1:find(tempnl == sn));
31     % Discard the path we just found from tempnl.
32     tempnl = tempnl((find(tempnl == sn)+1):end);
33     % Now, we set all the appropriate elements of x_{ijl} (with l=i)
34     % to one. Since the path is backwards and the matrix is used as
35     % x_{jil}, for each element j in thisnl except the last, we set
36     % x_{ij}(j, j+1, i) to one. This is done using sub2ind, since just
37     % inserting the vectors with ordinary subscript indexing sets
38     % entire blocks of the matrix (which of course is incorrect).
39     xij(sub2ind(size(xij),thisnl(1:end-1),thisnl(2:end),i*ones(length(thisnl)-1,1))) = 1;
40     % Finally, we set x_{t_{1}s_{1}l} to one, since this path
41     % obviously forms a connection.
42     xt1sl(1,1,i) = 1;
43 end
44 end

```

A.3 The primal feasibility heuristic function

```

1 % TMA521 - Large scale optimization
2 % Spring 2013
3 % Project 1, task 2
4 % Simon Sigurdhsson
5
6 % This is the primal feasibility heuristic written for task 2, and it
7 % is explained further in the report.
8
9 function [ ncom, nnl ] = heuristic( dimX, dimY, pi, ~, com, nl)
10 %HEURISTIC Finds a primal feasible solution given a Lagrangian dual solution
11 % This heuristic, explained further by the report, basically transforms
12 % a dual solution to a primal feasible solution by re-routing and/or
13 % discarding paths from the dual solution. It is polynomial in the number
14 % of nodes of the problem (in fact,  $O(n^2)$ ), and will always terminate with
15 % a feasible (but potentially very bad) solution.
16 % First, we save the input variables since we'll be changing them a bit.
17 nnl = nl; ncom = com; opi = pi;
18 % Now, we find the number of times each node has been used by a path.
19 % Call this number  $0 \leq n_i \leq k$ . It is found by simply iterating over each
20 % node and setting  $n_j$  to the number of times that node occurs in nl.
21 nodeusage = zeros(size(pi));
22 for i=1:length(nodeusage)
23     nodeusage(i) = length(find(nl == i));
24 end
25 % If solution has no overused nodes, i.e.  $n_i < 2$  for all i, the solution
26 % primal feasible since all paths are vertex disjoint. In that case, we
27 % return the input solution as our primal feasible solution.
28 if (max(nodeusage) < 2)
29     return
30 end
31 % If we have overused/infeasible nodes, we save all of them in a vector

```

```

32 % in order to iterate over them. We want to return a solution for which
33 % length(infeasiblenodes) = 0.
34 infeasiblenodes = find(nodeusage >= 2);
35 % While infeasiblenodes isn't empty (i.e. its length isn't 0), we eliminate
36 % paths passing through nodes found in the list.
37 while(~isempty(infeasiblenodes))
38     % For every infeasible node, we try to either replace a path with a
39     % new one, or simply discard it.
40     for i=1:length(infeasiblenodes)
41         % First, we reset the nl, com and pi variables by copying the
42         % com and nl variables from the current "state" (i.e. the current
43         % output variables), and pi from the input variable.
44         nl = nnl; com = ncom; pi = opi;
45         % We find the first path that passes through the first
46         % infeasible node in our list. This code is based on the code
47         % in okcom.m, and functions like it. The result is a pair of
48         % variables first,last containing the positions of the path in
49         % nl, and a pair of variables sl,tl containing the start and
50         % end node of the path, respectively.
51         node = infeasiblenodes(i);
52         first = 0; last = 0; sl = 0; tl = 0;
53         for i = 1:size(com,1)
54             first = last+1;
55             slask = find(nl(last+1:length(nl)) == com(i,1));
56             last = slask(1)+first-1;
57             if sum(nl(first:last) == node) > 0
58                 tl = nl(first);
59                 sl = nl(last);
60                 break
61             end
62         end
63         % Since we found out path (we always will), we remove it
64         % from nl as well as removing its corresponding row in com.
65         % If we can replace it, we will.
66         nl(first:last) = [];
67         com(find(ismember(com, [sl tl], 'rows') == 1), :) = [];
68         % Having removed our path from nl, we modify pi by setting an
69         % infinite cost for all nodes in nl. This means that when finding
70         % an alternative to the path we've removed, the cost of creating
71         % a path that isn't vertex-disjoint to the others will be infinite
72         % and we can discard such solutions easily.
73         pi(nl) = Inf;
74         % Using this new pi, we find a new cheapest path between the nodes
75         % we removed earlier.
76         newnl = gsp(dimX, dimY, pi(:), 1, [tl sl]);
77         % If we didn't find a new path, we try again with the next node
78         % in the list of infeasible nodes (the modified variables will be
79         % reset at the beginning of the next iteration). If we did find
80         % a new path, we break out of the loop.
81         if length(newnl) < 2 || sum(pi(newnl)) == Inf % If a path wasn't found
82             continue % Continue to next iteration
83         else % Else
84             break % Break out of the loop
85         end

```

```

86     end
87     % Since we're outside the infeasible-node loop, we must have either found a
88     % new path, in which case we append it to com/nl, or we didn't, in which case
89     % it has been removed from com/nl and we simply update the output variables
90     % to match the new set of paths.
91     if length(newnl) >= 2 && sum(pi(newnl)) ~= Inf % If a new path was found
92         % Append the found path to com/nl
93         ncom = [com; s1 t1];
94         nnl = [nl; flipud(newnl)];
95     else
96         % Copy the paths without appending anything
97         ncom = com;
98         nnl = nl;
99     end
100    % Before the end of the loop, we recalculate the node usage in the same way
101    % as before the loop, and update the infeasiblenodes vector to contain the
102    % new set of infeasible nodes (which will always be smaller).
103    for i=1:length(nodeusage)
104        nodeusage(i) = length(find(nl == i));
105    end
106    infeasiblenodes = find(nodeusage >= 2);
107 end
108 end

```