

Accelerating Column Generation for a Survivable Network Design Problem

Markus Leitner^{*,*} Günther R. Raidl^{*} Ulrich Pferschy[◇]

^{*}*Department of Telematics and Network Engineering, Carinthia University of Applied Sciences
Primoschgasse 8, 9020 Klagenfurt, Austria, markus.leitner@fh-kaernten.at*

^{*}*Institute of Computer Graphics and Algorithms, Vienna University of Technology
Favoritenstraße 9-11, 1040 Vienna, Austria, raidl@ads.tuwien.ac.at*

[◇]*Institute of Statistics and Operations Research, University of Graz
Universitätsstraße 15, 8010 Graz, Austria, pferschy@uni-graz.at*

Abstract

We consider a network design problem occurring in the extension of fiber optic networks on the last mile which generalizes the (Price Collecting) Steiner Tree Problem by introducing redundancy requirements on some customer nodes. In this work we present a formulation for this problem based on exponentially many variables and solve its linear relaxation by column generation. Using alternative dual-optimal solutions in the pricing problem we are able to significantly reduce the effects of typical efficiency issues of simplex based column generation. Computational results clearly show the advantages of our proposed strategy with respect to the number of pricing iterations needed as well as by means of required running times.

Keywords: *Survivable Network Design, Steiner Tree Problem, Column Generation*

1 Introduction

We consider a real world network design problem occurring in the extension of fiber optic networks on the last mile. This problem, to which we refer as b_{\max} -Survivable Network Design Problem (SNDP) generalizes the (Price Collecting) Steiner Tree Problem (STP) on a graph by introducing redundancy requirements on some customer nodes. Formally, we are given an undirected graph $G = (V, E)$ modelling a spatial area in which an existing fiber optic infrastructure should be augmented to supply new customers. Edges $e \in E$ represent (potential) fiber optic routes with corresponding cable lengths $l_e \geq 0$ and installation costs $c_e \geq 0$. The node set V is the disjoint union of customer nodes C with associated prizes $p_k \geq 0, \forall k \in C$ (i.e. the expected return on investment) and spatial nodes S (switches, possible Steiner nodes). The set of customer nodes C is further partitioned into type-1 customer nodes C_1 and type-2 customer nodes C_2 . Already existing infrastructure is represented by a subgraph $I = (V_I, E_I)$ of G which we shrink into a single root node $0 \in V$ in a preprocessing step.

A solution $G' = (V', E')$, $V' \subseteq V, E' \subseteq E$ is a connected subgraph of G feasibly connecting a set of customers $C' \subseteq C$. While type-1 customer nodes $k \in C_1$ are feasibly connected if a path from 0 to k exists in G' , type-2 customer nodes $k \in C_2$ need two node disjoint paths from 0 to k . Furthermore, this biconnectivity condition for type-2 customers is relaxed if such a node $k \in C_2$ has an associated *maximum branch line length* $b_{\max}(k) > 0$. Such a type-2 customer node is feasibly connected, if there exists some *branch node* $j \in V'$ with two node disjoint paths to the root node 0 and a single path (*branch line*) from j




Figure 1: Solution with $b_{\max} = 0$.

Figure 2: Solution with $b_{\max} > 0$.

to k of maximum total length $b_{\max}(k)$. Figure 1 depicts a solution without considering b_{\max} -redundancy, i.e. $b_{\max} = 0$, while Figure 2 visualizes a solution to b_{\max} -SNDP with $b_{\max} > 0$.

We distinguish two problem variants with respect to the objective function. In the so called *Operative Planning Task* (OPT) we need to identify a cheapest subgraph G' feasibly connecting all customer nodes ($C' = C$), i.e.

$$c_{\text{OPT}}(G') = \min \sum_{e \in E'} c_e,$$

while we want to find the most profitable solution in the *Strategic Simulation Task* (SST), i.e.

$$c_{\text{SST}}(G') = \min \sum_{e \in E'} c_e + \sum_{k \in C \setminus C'} p_k.$$

As b_{\max} -SNDP contains the NP-hard STP as a special case ($C_2 = \emptyset$), we conclude that b_{\max} -SNDP is NP-hard, too.

In this article, we present a formulation for b_{\max} -SNDP based on exponentially many variables and solve its linear relaxation by column generation. Hereby, we utilise alternative dual-optimal solutions to reduce the number of iterations needed and dramatically decrease the required running time. Our computational results clearly document the beneficial effects of this strategy.

2 Previous Work

A lot of research has been conducted on similar, more prominent network design problems, namely the (Price Collecting) Steiner Tree Problem (STP) and several variants of the Survivable Network Design Problem (SNDP), see e.g. [12] for a survey on the STP and [5] for the SNDP. Wagner et al. presented exact approaches for b_{\max} -SNDP based on multicommodity flows [11] and directed connection cuts [10]. Another, even stronger model – but which does not consider the special case of b_{\max} -redundancy – has been described by Chimani et al. [2]. It is also based on connection cuts but models redundant connections by reversely directed paths. Recently, we approached b_{\max} -SNDP by means of Lagrangian Decomposition, Greedy Randomized Adaptive Search, Variable Neighborhood Search as well as hybrid methods combining Lagrangian Decomposition with Variable Neighborhood Descent [6].

3 The Connection Formulation

To model b_{\max} -SNDP as an Integer Linear Program (ILP) we consider the set of all possible feasible connections P_k for each customer $k \in C$. For type-1 customers $k \in C_1$, P_k corresponds to the set of all paths from the root node 0 to k , i.e.

$$P_k = \{p \subseteq E \mid p \text{ forms a path from } 0 \text{ to } k\},$$

while for type-2 customers $k \in C_2$, P_k can be expressed as follows:

$$P_k = \{p \subseteq E \mid p \text{ forms two node disjoint paths from } 0 \text{ to some node } j \text{ and one path from } j \text{ to } k \text{ whose length does not exceed } b_{\max}(k)\}.$$

We formulate the SST variant of our problem by the following *integer master problem* (IMP) using binary variables $f_p^k, \forall k \in C, \forall p \in P_k$ to indicate whether a corresponding connection $p \in P_k$ is realized ($f_p^k = 1$) or not, decision variables $x_e \in \{0, 1\}, \forall e \in E$ to specify whether an edge e is part of the solution ($x_e = 1$) or not, and binary variables $y_k, \forall k \in C$ to denote whether a feasible route to customer k is installed ($y_k = 1$) or not. y_k variables are fixed to one in the OPT variant.

$$\begin{aligned}
(1) \quad & \text{(IMP)} \quad z = \min \sum_{e \in E} c_e x_e + \sum_{k \in C} p_k (1 - y_k) \\
(2) \quad & \text{s.t.} \quad \sum_{p \in P_k} f_p^k - y_k = 0 \quad \forall k \in C \\
(3) \quad & x_e - \sum_{p \in P_k | e \in p} f_p^k \geq 0 \quad \forall k \in C, \forall e \in E \\
(4) \quad & x_e \in \{0, 1\} \quad \forall e \in E \\
(5) \quad & y_k \in \{0, 1\} \quad \forall k \in C \\
(6) \quad & f_p^k \in \{0, 1\} \quad \forall k \in C, \forall p \in P_k
\end{aligned}$$

The linear relaxation of IMP – the *linear master problem* (MP) – is given by substituting the integrality constraints (4)–(6) by

$$\begin{aligned}
(7) \quad & x_e \geq 0 \quad \forall e \in E \\
(8) \quad & y_k \geq 0 \quad \forall k \in C \\
(9) \quad & y_k \leq 1 \quad \forall k \in C \\
(10) \quad & f_p^k \geq 0 \quad \forall k \in C, \forall p \in P_k
\end{aligned}$$

Let $\mu_k, \forall k \in C$ be the dual variables associated to the convexity constraints (2) and $\pi_{k,e}, \forall k \in C, \forall e \in E$ be the dual variables associated to the coupling constraints (3). Note that one would usually replace equalities (2) by inequalities (≥ 0) to restrict dual variables when solving such a model with column generation. However, as we do not need to consider customers $k \in C$ in the pricing problem if $\mu_k \leq 0$ (as will be explained in the following) no significant differences could be observed when using inequalities. Furthermore, let $F = \{f_p^k \mid k \in C, p \in P_k\}$ be the set of all f_p^k variables representing columns in MP. Since F consists of an exponential number of variables we define the *restricted master problem* (RMP) using only a small subset $\tilde{F} \subsetneq F$; otherwise RMP corresponds to MP. When solving RMP we obtain optimal dual variable values μ_k^* and $\pi_{k,e}^*$, defining reduced prices $\bar{c}_{k,p}$ for variables $f_p^k \in F \setminus \tilde{F}$:

$$\bar{c}_{k,p} = -\mu_k^* + \sum_{e \in p} \pi_{k,e}^*$$

The pricing problem is now to find $(k^*, p^*) = \operatorname{argmin}_{k \in C, p \in P_k} \{\bar{c}_{k,p}\}$. If $\bar{c}_{k^*, p^*} \geq 0$ we have obtained an optimal solution to MP. Otherwise, we add at least one column with negative reduced costs and resolve RMP.

Solving the Pricing Problem: More generally speaking, in the pricing subproblem we have to find a feasible connection for some $k \in C$ yielding negative reduced costs $\bar{c}_{k,p} = -\mu_k^* + \sum_{e \in p} \pi_{k,e}^*$ or prove that no such connection exists. For this purpose we need to determine the cheapest feasible connection on the graph in which edge $e \in E$ has costs $\pi_{k,e} \geq 0$, for each $k \in C$ with $\mu_k > 0$. When the costs of such a connection are less than μ_k , we have found an appropriate connection, i.e. the corresponding variable f_p^k can be added to RMP.

While for $k \in C_1$ finding the cheapest feasible connection is a simple shortest path calculation from 0 to k , we have to find a cheapest pair of node-disjoint paths from 0 to k for type-2 customers (without yet considering b_{\max} -redundancy). Suurballe and Tarjan [8] (see also [4]) presented an algorithm to efficiently

compute a shortest arc-disjoint pair of paths between two nodes in time $O(|E| + |V| \log |V|)$. By applying this algorithm on the split graph of the original graph we can compute a shortest node-disjoint pair of paths. The split graph is obtained by replacing each node $v \in V$ by a pair of nodes v' and v'' . For each such pair, we add an arc (v', v'') with zero costs. Each edge $e = (u, v)$ of G is replaced by two directed arcs (u'', v') , (v'', u') having costs c_e .

We use a simple extension of this algorithm to consider b_{\max} -redundancy by determining the cheapest combination of a node-disjoint pair of paths from 0 to some node j in the b_{\max} -neighborhood of a customer k and a single path from j to k whose length does not exceed $b_{\max}(k)$. To avoid unnecessary calculations we do only consider those possible branch nodes j for which the costs of the shortest path to the root node do not exceed half of the costs of the so far found cheapest connection.

4 Using Alternative Dual-Optimal Solutions

It is well known that (simplex based) column generation approaches often suffer from inefficiency resulting in a large number of needed pricing iterations as well as long computation times. Vanderbeck [9] describes five major efficiency issues of simplex based column generation. Several stabilization techniques to reduce their effects have been proposed, see e.g. [3] or [7] for reviews on those methods. From the issues described by Vanderbeck preliminary tests showed that *primal degeneracy* as well as the *heading-in effect* are mainly relevant in our case. Instead of using a problem-independent stabilization approach we analyse the dual of RMP to take advantage of problem specific characteristics. Let $\lambda_k \leq 0$ denote the dual variables associated to inequalities (9) and $\tilde{F}_k \subseteq \tilde{F}$ be the set of variables representing connections to customer $k \in C$ in RMP. Then the *restricted dual problem* (RDP) is given by (11)–(16).

$$\begin{aligned}
(11) \quad & \text{(RDP)} \quad z = \sum_{k \in C} \lambda_k + p_k \\
(12) \quad & \sum_{k \in C} \pi_{k,e} \leq c_e \quad \forall e \in E \\
(13) \quad & \mu_k - \sum_{e \in p} \pi_{k,e} \leq 0 \quad \forall k \in C, \forall p \in \tilde{F}_k \\
(14) \quad & -\mu_k + \lambda_k \leq -p_k \quad \forall k \in C \\
(15) \quad & \pi_{k,e} \geq 0 \quad \forall k \in C, \forall e \in E \\
(16) \quad & \lambda_k \leq 0 \quad \forall k \in C
\end{aligned}$$

Let $D^* = (\lambda^*, \mu^*, \pi^*)$ be an optimal solution to RDP. Since only few connection variables $f \in \tilde{F}$ will be non-zero in an optimal solution to RMP, RMP is usually degenerate, i.e. alternative optimal solutions to RDP exist. In the following, we exploit different dual-optimal solutions to improve the convergence properties of our column generation algorithm. This approach can be interpreted as a generic stabilization technique that “centers” an actual LP solution.

Let $E' \subseteq E$ denote the subset of edges which are not part of any so far included connection, i.e. $E' = \{e \in E \mid \nexists f_p^k \in \tilde{F} : e \in p\}$. For edges $e \in E'$ all values $\pi_{k,e} \geq 0, \forall k \in C$ are dual optimal as long as $\sum_{k \in C} \pi_{k,e} \leq c_e$. Furthermore, for edges $e \in E \setminus E'$, we can increase the sum of dual variable values $\sum_{k \in C} \pi_{k,e}$ by $\delta_e = c_e - \sum_{k \in C} \pi_{k,e}$.

Since CPLEX¹ generates minimal values for dual variables (i.e. $\pi_{k,e} = 0, \forall k \in C, \forall e \in E'$, usually $\delta_e > 0$ for some edges $e \in E \setminus E'$) and $|E'|$ is typically quite large in the beginning, corresponding edge costs in the pricing subproblem are not meaningful. More precisely, a lot of irrelevant columns will be generated since many connections have equal costs. To reduce this harmful behavior one could simply use the alternative dual optimal solution $D' = (\lambda^*, \mu^*, \pi')$ with $\pi'_{k,e} = \frac{c_e}{|C|}, \forall k \in C, \forall e \in E'$ and $\pi'_{k,e} = \pi_{k,e}^* + \frac{\delta_e}{|C|}, \forall k \in C, \forall e \in E \setminus E'$. However, as will be illustrated by our computational

¹<http://www.ilog.com>

Algorithm 1: Column Generation

```

d = 1
create and add set of initial columns  $\tilde{F}$ 
 $E' = \{e \in E \mid \nexists f_p^k \in \tilde{F} : e \in p\}$ 
m = true
while m do
  m = false
  solve RMP
   $\delta_e = (c_e - \sum_{k \in C} \pi_{k,e}) / d, \forall e \in E'$ 
  forall k ∈ C do
    if  $\mu_k > 0$  then
       $c'_e = \begin{cases} \pi_{k,e} + \delta_e & \text{if } e \in E', \\ c_e/d & \text{else.} \end{cases} \quad \forall e \in E$ 
      p = shortest connection to k using edge costs c'
       $E_p = \{e \in E \mid e \in p\}$ 
      if  $\sum_{e \in E_p} c'_e < \mu_k$  then
        add corresponding variable  $f_p^k$  to RMP
        if  $E_p \not\subseteq E'$  then
          d = 1
           $E' = E' \cup E_p$ 
          m = true
  if m == false ∧ d < |C| then
    m = true
    d ++
  
```

results we can do even better by initially using different dual-optimal solutions $D^{(k,d)} = (\lambda^*, \mu^*, \pi^{(k,d)})$, for all $k \in C$ – controlled by parameter d ($1 \leq d \leq |C|$) – which finally converge to D' for $d = |C|$. When considering client $k \in C$ in the pricing problem, we use dual values $\pi_{k,e}^{(k,d)} = \frac{c_e}{d}, \forall e \in E'$ and $\pi_{k,e}^{(k,d)} = \pi_{k,e}^* + \frac{\delta_e}{d}, \forall e \in E \setminus E'$. Note that assuming $\pi_{k',e}^{(k,d)} = 0, \forall k' \neq k \in C, \forall e \in E'$ and $\pi_{k',e}^{(k,d)} = \pi_{k',e}^*, \forall k' \neq k \in C, \forall e \in E \setminus E'$ we again only use dual optimal solutions when solving the pricing problem. As shown in Algorithm 1 parameter d is initially set to one and gradually incremented up to $|C|$ in case no column with negative reduced cost could be priced in and reset to one in case columns including new edges have been added to RMP. Since we essentially use D' if $d = |C|$ we can terminate the column generation process if no column with negative reduced costs could be found for $d = |C|$.

5 Computational Results

We tested our algorithms on real world instance sets from a German city [1] – see Table 1 – with an absolute time limit of 7200 seconds. ILOG CPLEX 11.1 has been used to solve RMP after each pricing iteration. We use the dual simplex approach to solve RMP since it turned out to perform better than the primal simplex method. Note that, instead of generating feasible or optimal solutions to our problem, we only focus on solving its linear relaxation in this work.

The set of columns \tilde{F} is initialized by (i) cheapest connections to each client $k \in C$, (ii) connections of a solution constructed with the Minimum Spanning Tree Augmentation Heuristic [6], and (iii) connections that emphasize pairs of customer nodes. For the latter, we consider each pair of customer nodes $k, k' \in C, k \neq k'$ and determine the cheapest connection to k' while treating all edges part of the cheapest feasible connection of k as pseudo-infrastructure, i.e. set their edge costs to zero. In each pricing iteration we add the cheapest connection to each customer $k \in C$ to RMP if it has negative reduced costs. In the following, CG^* refers to the standard column generation approach without adapting dual values, while CG' refers

Table 1: Instance set characteristics.

Set	#	$ V $	$ E $	$ C $	$ \overline{C} $	$ C_1 $	$ \overline{C}_1 $	$ C_2 $	$ \overline{C}_2 $	b_{\max}	$ V(b_{\max}) $
ClgSE-I1	25	190	377	5–8	5.9	3–5	3.8	2–3	2.1	30	3.79
ClgSE-I2	15	190	377	11–17	13.8	7–12	8.9	4–7	4.9	30	8.97
ClgSE-I3	15	190	377	8–12	9.6	5–8	6.0	3–6	3.6	30	6.04
ClgN1B-I1	20	2804	3082	11–14	11.8	8–11	8.5	3–4	3.3	100	8.49
ClgN1B-I2	19	2804	3082	7–11	9.0	3–6	4.1	4–6	5.0	100	3.99
ClgME-I1	25	1757	3877	6–10	7.2	4–7	5.0	2–3	2.3	100	4.96
ClgME-I2	15	1523	3290	11–14	12.2	8–11	8.7	3–4	3.5	100	8.71

Table 2: Number of instances where linear relaxation could be solved.

Set	#	OPT+RED			SST+RED			OPT+BMAX			SST+BMAX		
		CG^*	CG'	CG^k	CG^*	CG'	CG^k	CG^*	CG'	CG^k	CG^*	CG'	CG^k
ClgSE-I1	25	25	25	25	25	25	25	25	25	25	25	25	
ClgSE-I2	15	15	15	15	15	15	15	15	15	15	15	15	
ClgSE-I3	15	14	15	15	14	15	15	15	15	15	15	15	
ClgN1B-I1	20	7	20	20	9	20	20	12	20	20	12	20	
ClgN1B-I2	19	5	19	19	7	19	19	7	19	19	4	19	
ClgME-I1	25	0	1	19	0	2	19	0	1	17	0	1	
ClgME-I2	15	0	2	10	0	1	10	0	2	9	0	1	

to the simpler adaptation strategy (i.e. equally increasing dual variable values), and CG^k denotes the adaptive strategy described by Algorithm 1. Furthermore, *RED* refers to the problem variant without considering b_{\max} -redundancy ($b_{\max} = 0$) while *BMAX* takes b_{\max} -redundancy into account. Table 2 compares the number of instances solved by CG^* , CG' , and CG^k , while Table 3 compares median run times of those three variants and of the Lagrangian Decomposition (LD) approach from [6] which is equally strong from a theoretic point of view and also practically generates lower bounds identical to those obtained by column generation.

Both strategies to adapt dual variable values perform significantly better than CG^* with respect to the number of solved instances as well as median runtimes. Furthermore, while CG' failed to solve most of the ClgM instances, CG^k was successful on 60% to 76% of those instances. While LD outperforms column generation for most of the ClgM instances, CG^k is the fastest (except for ClgN1B-I1 with $b_{\max} > 0$) method for the other instance sets. We further observed, that CG^k usually finds the optimal LP bound relatively quickly and spends around 60% to 80% of its total runtime to prove optimality (*tailing-off effect*) which facilitates an early termination criterion in a possible extension to a branch and price algorithm. CG^* and CG' need almost all of their (even longer) runtimes to find the optimal LP bound. Finally, Table 4 compares the number of pricing iterations needed for instances that could be solved, while Table 5 depicts the relative amount of time spent for repeatedly solving RMP (i.e. solving the LPs with CPLEX). We conclude, that while CG' often needs fewer iterations than CG^k , both are able to considerably reduce the number of pricing iterations compared to CG^* . Furthermore, CG^k clearly performs best with respect to reducing the effects of primal degeneracy.

Table 3: Median run times.

Set	OPT+RED				SST+RED				OPT+BMAX				SST+BMAX			
	CG^*	CG'	CG^k	LD	CG^*	CG'	CG^k	LD	CG^*	CG'	CG^k	LD	CG^*	CG'	CG^k	LD
ClgSE-I1	3.4	0.3	0.2	1.0	2.0	0.3	0.2	1.1	2.1	0.3	0.4	3.8	2.1	0.5	0.4	3.9
ClgSE-I2	46.2	4.8	1.7	5.2	37.4	5.5	1.7	3.9	45.7	3.0	2.2	16.8	45.3	3.9	1.9	17.2
ClgSE-I3	7.5	0.6	0.6	2.1	4.4	0.6	0.7	2.0	9.1	0.9	1.0	10.6	9.9	0.9	0.9	15.1
ClgN1B-I1	7200.0	33.5	16.5	93.7	7200.0	26.4	22.6	89.5	2605.3	70.3	198.2	1015.3	3795.9	79.7	169.6	753.1
ClgN1B-I2	7200.0	206.3	15.7	62.1	7200.0	159.8	20.9	54.3	7200.0	251.5	149.9	463.6	7200.0	294.5	172.0	427.3
ClgME-I1	7200.0	7200.0	232.5	77.4	7200.0	7200.0	948.4	91.5	7200.0	7200.0	2249.8	3386.9	7200.0	7200.0	4441.5	3129.2
ClgME-I2	7200.0	7200.0	1417.5	75.5	7200.0	7200.0	964.9	80.0	7200.0	7200.0	3887.7	2113.4	7200.0	7200.0	3005.4	1621.6

Table 4: Average relative number of pricing iterations and corresponding standard deviations.

Set	OPT+RED		SST+RED		OPT+BMAX		SST+BMAX	
	$\frac{CG^k}{CG^*}$	$\frac{CG^k}{CG'}$	$\frac{CG^k}{CG^*}$	$\frac{CG^k}{CG'}$	$\frac{CG^k}{CG^*}$	$\frac{CG^k}{CG'}$	$\frac{CG^k}{CG^*}$	$\frac{CG^k}{CG'}$
ClgSE-I1	0.28 (0.19)	1.57 (0.89)	0.32 (0.22)	1.41 (0.68)	0.33 (0.18)	1.86 (0.90)	0.37 (0.19)	1.72 (0.87)
ClgSE-I2	0.33 (0.17)	1.56 (0.67)	0.37 (0.28)	1.49 (0.61)	0.40 (0.21)	1.80 (0.65)	0.41 (0.28)	1.70 (1.04)
ClgSE-I3	0.38 (0.28)	2.25 (0.93)	0.39 (0.30)	2.15 (0.69)	0.33 (0.21)	2.33 (1.00)	0.35 (0.25)	2.31 (1.02)
ClgN1B-I1	0.12 (0.03)	2.26 (1.11)	0.15 (0.05)	2.04 (1.00)	0.17 (0.09)	2.72 (1.60)	0.16 (0.05)	2.32 (1.46)
ClgN1B-I2	0.09 (0.04)	1.10 (0.70)	0.11 (0.02)	0.91 (0.33)	0.12 (0.03)	0.93 (0.27)	0.15 (0.04)	0.85 (0.24)
ClgME-I1	- (-)	0.37 (0.00)	- (-)	0.29 (0.16)	- (-)	0.95 (0.00)	- (-)	1.01 (0.00)
ClgME-I2	- (-)	0.28 (0.06)	- (-)	0.50 (0.00)	- (-)	0.85 (0.18)	- (-)	1.69 (0.00)

Table 5: Average relative time and corresponding standard deviations for solving RMP.

	Set	OPT			SST		
		CG^*	CG'	CG^k	CG^*	CG'	CG^k
RED	ClgSE-I1	0.85 (0.10)	0.63 (0.18)	0.49 (0.13)	0.84 (0.10)	0.67 (0.13)	0.49 (0.10)
	ClgSE-I2	0.94 (0.04)	0.83 (0.10)	0.51 (0.14)	0.94 (0.05)	0.84 (0.09)	0.52 (0.10)
	ClgSE-I3	0.85 (0.09)	0.70 (0.13)	0.46 (0.09)	0.87 (0.09)	0.69 (0.14)	0.48 (0.10)
	ClgN1B-I1	0.96 (0.05)	0.85 (0.11)	0.52 (0.08)	0.97 (0.05)	0.85 (0.10)	0.58 (0.11)
	ClgN1B-I2	0.99 (0.01)	0.92 (0.08)	0.53 (0.15)	0.99 (0.01)	0.94 (0.04)	0.62 (0.16)
	ClgME-I1	1.00 (0.00)	0.98 (0.03)	0.73 (0.19)	1.00 (0.00)	0.98 (0.03)	0.79 (0.16)
	ClgME-I2	1.00 (0.00)	0.99 (0.02)	0.85 (0.17)	1.00 (0.00)	0.99 (0.02)	0.86 (0.15)
BMAX	ClgSE-I1	0.63 (0.17)	0.46 (0.14)	0.28 (0.09)	0.65 (0.15)	0.47 (0.15)	0.32 (0.10)
	ClgSE-I2	0.85 (0.11)	0.72 (0.12)	0.37 (0.08)	0.86 (0.12)	0.72 (0.14)	0.37 (0.09)
	ClgSE-I3	0.72 (0.16)	0.53 (0.14)	0.29 (0.07)	0.75 (0.15)	0.56 (0.14)	0.29 (0.06)
	ClgN1B-I1	0.52 (0.29)	0.34 (0.19)	0.07 (0.04)	0.58 (0.29)	0.37 (0.20)	0.10 (0.06)
	ClgN1B-I2	0.81 (0.10)	0.63 (0.16)	0.08 (0.06)	0.82 (0.11)	0.66 (0.15)	0.15 (0.07)
	ClgME-I1	0.82 (0.16)	0.81 (0.23)	0.33 (0.36)	0.82 (0.14)	0.84 (0.18)	0.30 (0.33)
	ClgME-I2	0.94 (0.05)	0.92 (0.12)	0.62 (0.32)	0.94 (0.04)	0.93 (0.13)	0.62 (0.33)

6 Conclusions and Future Work

In this paper we presented a column generation approach for b_{\max} -SNDP, a problem which occurs in the design of the last mile in fiber optic networks, based on exponentially many variables corresponding to feasible client connections. By using alternative dual-optimal solutions when solving the pricing subproblem, we could achieve a dramatic speedup with respect to computation time. However, since primal degeneracy still has harmful effects on several instances, we want to combine this approach with a problem independent stabilization method to further decrease the time needed to solve the LP relaxation of our model in future work. Furthermore, we want to extend our approach to a branch-and-price algorithm to obtain provably optimal solutions as well as combine column generation with metaheuristic methods for generating high quality solutions with small gaps for large instances. Finally, from a more general point of view we believe that the usage of alternative dual-optimal solutions might increase performance of column generation approaches of other problems of similar structure, too.

Acknowledgements

This work is supported by the Austrian Research Promotion Agency (FFG) under grant 811378 and by the Austrian Science Fund (FWF) under grant P20342-N13.

References

- [1] P. Bachhiesl. The OPT- and the SST-problems for real world access network design – basic definitions and test instances. Working Report 01/2005, Carinthia Tech Institute, Department of Telematics and Network Engineering, Klagenfurt, Austria, 2005.
- [2] M. Chimani, M. Kandyba, and P. Mutzel. A new ILP formulation for 2-root-connected prize-collecting Steiner networks. In *15th Annual European Symposium on Algorithms (ESA'07)*, volume 4698 of *LNCS*, pages 681–692. Springer, 2007.
- [3] O. du Merle, D. Villeneuve, J. Desrosiers, and P. Hansen. Stabilized column generation. *Discrete Mathematics*, 194(1-3):229–237, 1999.
- [4] K. Kar, M. Kodialam, and T. V. Lakshman. Routing restorable bandwidth guaranteed connections using maximum 2-route flows. *IEEE/ACM Transactions on Networking*, 11(5):772–781, 2003.
- [5] H. Kerivin and A. R. Mahjoub. Design of survivable networks: A survey. *Networks*, 46(1):1–21, 2005.
- [6] M. Leitner and G. R. Raidl. Lagrangian decomposition, metaheuristics, and hybrid approaches for the design of the last mile in fiber optic networks. In M. J. Blesa et al., editors, *Hybrid Metaheuristics 2008*, volume 5296 of *LNCS*, pages 158–174. Springer, 2008.
- [7] M. E. Lübbecke and J. Desrosiers. Selected topics in column generation. *Operations Research*, 53(6):1007–1023, 2005.
- [8] J. W. Suurballe and R. E. Tarjan. A quick method for finding shortest pairs of disjoint paths. *Networks*, 14:325–335, 1984.
- [9] F. Vanderbeck. Implementing mixed integer column generation. In G. Desaulniers, J. Desrosiers, and M. M. Solomon, editors, *Column Generation*, pages 331–358. Springer US, 2005.
- [10] D. Wagner, U. Pferschy, P. Mutzel, G. R. Raidl, and P. Bachhiesl. A directed cut model for the design of the last mile in real-world fiber optic networks. In B. Fortz, editor, *Proceedings of the International Network Optimization Conference 2007*, pages 103/1–6, Spa, Belgium, 2007.
- [11] D. Wagner, G. R. Raidl, U. Pferschy, P. Mutzel, and P. Bachhiesl. A multi-commodity flow approach for the design of the last mile in real-world fiber optic networks. In K.-H. Waldmann and U. M. Stocker, editors, *Operations Research Proceedings 2006*, pages 197–202. Springer, 2007.
- [12] P. Winter. Steiner problem in networks: a survey. *Networks*, 17(2):129–167, 1987.