

Thinning out Steiner trees

A node-based model for uniform edge costs

Matteo Fischetti · Markus Leitner ·
Ivana Ljubić · Martin Luipersbeck ·
Michele Monaci · Max Resch ·
Domenico Salvagnin · Markus Sinnl

Received: date / Accepted: date

Abstract The Steiner Tree Problem is a challenging NP-hard problem. Many hard instances of this problem are publicly available, that are still unsolved by state-of-the-art branch-and-cut codes. A typical strategy to attack these instances is to enrich the polyhedral description of the problem, and/or to implement more and more sophisticated separation procedures and branching strategies. In this paper we investigate the opposite viewpoint, and try to make the solution method as simple as possible while working on the modeling side. Our working hypothesis is that the extreme hardness of some classes of instances mainly comes from over-modeling, and that some instances can become quite easy to solve when a simpler model is considered. In other words, we aim at “thinning out” the usual models for the sake of getting a more agile framework. In particular, we focus on a model that only involves node variables, which is rather appealing for the “uniform” cases where all edges have the same cost. We show that this model allows one to quickly produce very good (sometimes proven optimal) solutions for notoriously hard instances from the literature. In particular, we report improved solutions for several SteinLib instances, including the (in)famous hypercube ones. Even though we do not claim our approach can work well in all cases, we report surprisingly good results for a number of unsolved instances. In some cases, our approach takes just few seconds to prove optimality for instances never solved (even after days of computation) by the standard methods.

Keywords Mixed Integer Programming · Exact Computation

Mathematics Subject Classification (2000) 90C10 · 90C27

Matteo Fischetti · Michele Monaci · Domenico Salvagnin
DEI, University of Padua, Italy
E-mail: {matteo.fischetti,michele.monaci,domenico.salvagnin}@unipd.it

Markus Leitner · Ivana Ljubić · Martin Luipersbeck · Markus Sinnl
Department of Statistics and Operations Research, University of Vienna, Austria
E-mail: {markus.leitner,ivana.ljubic,martin.luipersbeck,markus.sinnl} @univie.ac.at

Max Resch
Vienna University of Technology, Austria
E-mail: max.resch@alumni.tuwien.ac.at

1 Introduction

The Steiner Tree Problem (STP), in any of its various versions, is a challenging NP-hard problem that involves two related decisions: choosing the nodes to cover, and then covering them at minimum cost. Once the first decision has been taken, the second one is just trivial as it amounts to solving a minimum-cost tree spanning the selected nodes.

In this paper we introduce a new Mixed-Integer Linear Programming (MIP) approach for solving hard instances of the Steiner tree problem (close) to optimality. Instead of modeling graph connectivity using edge or arc variables (where many of them can exist), we propose to model it by using node variables only. Our model is particularly suited for “uniform” cases where all edges have the same cost. Besides the fact that these node-based models contain significantly less variables, they also avoid equivalences induced by uniform edge weights. For very dense graphs, or those containing a lot of symmetries, this strategy significantly outperforms the standard models where connectivity is modeled by using edge variables.

Our approach works for different variants of the Steiner tree problem, including the (rooted) Prize-Collecting STP (PCSTP), the Node-Weighted STP (NWSTP), the Maximum-Weight Connected Subgraph Problem (MWCS), and also the degree-constrained STP (DCSTP). To have a unified framework, in the following we will therefore focus on a slightly more general variant of the PCSTP, with a (potentially empty) set of terminal nodes. This general problem definition covers both the classical Steiner tree problem in graphs and its Prize-Collecting counterpart, as special cases. Necessary adaptations for the remaining problems will be explained below. It is worth mentioning that our code (with three variants submitted under the names *mozartballs*, *staynerd* and *mozartduet*) was the winner of the DIMACS Challenge on Steiner trees [1] in most of the categories (see [11] for more details). This article contains a summary of the main ingredients of this implementation.

Definition 1 (The prize-collecting Steiner tree problem (PCSTP)) Given an undirected graph $G = (V, E)$ with a (possibly empty) set of *real terminals* $T_r \subset V$, edge costs $c : E \mapsto \mathbb{R}^+$ and node revenues $p : V \mapsto \mathbb{R}^+$, the PCSTP is to find a subtree \mathcal{T} that spans all real terminals and such that the cost

$$c(\mathcal{T}) = \sum_{e \in E[\mathcal{T}]} c_e + \sum_{i \notin V[\mathcal{T}]} p_i$$

is minimized.

In the classical PCSTP version studied in the previous literature, $T_r = \emptyset$, and the problem can be equivalently stated as searching for a subtree that maximizes the difference between the collected node revenues ($\sum_{i \in V[\mathcal{T}]} p_i$) and the costs for establishing the links of that tree ($\sum_{e \in E[\mathcal{T}]} c_e$). One objective value can be transformed into another by subtracting $c(\mathcal{T})$ from the sum of all node revenues ($P = \sum_{i \in V} p_i$).

In general, each node in $V \setminus T_r$ is considered as a *Steiner node*, i.e., it can be used as an intermediate node to connect real terminals, or those with positive revenues.

Observe that there always exists an optimal PCSTP solution in which nodes with zero revenue are not leaves. The same holds for each node $i \in V$ such that $p_i > 0$

and $\min_{\{i,j\} \in E} c_{ij} > p_i$. (Note that we impose strict inequality in the latter condition.) Hence, only a specific subset of nodes in the PCSTP can be leaves of an optimal solution. We will refer to those nodes as *potential terminals*.

Definition 2 (Potential terminals) Among the nodes $i \in V \setminus T_r$, only those with revenues $p_i > 0$ such that at least one adjacent edge is strictly cheaper than p_i are considered as potential leaves. These nodes are referred to as *potential terminals*, and the associated set is denoted by T_p :

$$T_p = \{v \in V \setminus T_r \mid \exists \{u, v\} \text{ s.t. } c_{uv} < p_v\}.$$

In the following, we will call the set $T = T_r \cup T_p$, the set of *terminal nodes*. Our general problem definition covers the Steiner tree problem in graphs, since in this case all node revenues are equal to zero ($p_v = 0$, for all $v \in V$) and the set of real terminals is non empty, i.e., $\emptyset \neq T = T_r \subset V$.

2 A node-based MIP model

Node-based models for solving the Maximum-Weight Connected Subgraph Problem have been compared, both theoretically and computationally, in a recent publication [2]. Since there are no edge-costs involved in the objective of the MWCS, a natural MIP modeling approach is to derive a formulation in the space of node variables only. The first node-based model for the MWCS has been proposed in [5], and has been shown to computationally outperform extended formulations (involving both edge and node variables) on a benchmark set of instances from bioinformatics applications. However, as demonstrated in [2], the cycle-elimination model of [5] provides arbitrarily bad lower bounds and can be computationally improved by considering the notion of node-separators whose definition is provided below.

For STP/PCSTP instances, uniform edge costs can be embedded into node revenues (as shown below), so that using node-based MIPs appears natural in this case.

Definition 3 (Node Separators) For two distinct nodes k and ℓ from V , a subset of nodes $N \subseteq V \setminus \{k, \ell\}$ is called (k, ℓ) node separator if and only if after eliminating N from V there is no (k, ℓ) path in G . A separator N is minimal if $N \setminus \{i\}$ is not a (k, ℓ) separator, for any $i \in N$. Let $\mathcal{N}(k, \ell)$ denote the family of all (k, ℓ) separators.

Note that in order to make sure that a subset of chosen nodes is connected, it is sufficient to impose connectivity between the pairs of terminals (due to the minimization of the objective function). Therefore, we are mainly interested in node separators between pairs of terminals.

Let $N_\ell = \cup_{k \in T, k \neq \ell} \mathcal{N}(k, \ell)$ be the family of all such node separators with respect to a node $\ell \in T$. We will refer to elements from N_ℓ as ℓ -separators. Finally, let $\mathcal{N} = \cup_{\ell \in T} N_\ell$ be the set of all node subsets that separate two arbitrary terminals.

Let us assume that we are dealing with an undirected graph G with node revenues p_v 's and uniform edge costs $c_e = c$ for all $e \in E$.

In order to derive a node-based model, we will first shift edge costs into node costs as follows:

$$\tilde{c}_v = c - p_v, \quad \forall v \in V.$$

Let $P = \sum_{v \in V} p_v$ be the sum of all node revenues in G . Binary node-variables y_v , $\forall v \in V$, will be set to one if node v is part of the solution, and the node-based MIP model can be derived as follows:

$$\text{(PCSTP}_y) \quad \min \sum_{v \in V} \tilde{c}_v y_v + (P - c) \quad (1)$$

$$y(N) \geq y_i + y_j - 1 \quad \forall i, j \in T, i \neq j, \quad \forall N \in \mathcal{N}(i, j) \quad (2)$$

$$y_v = 1 \quad \forall v \in T_r \quad (3)$$

$$y_v \in \{0, 1\} \quad \forall v \in V \setminus T_r \quad (4)$$

where $y(N) = \sum_{v \in N} y_v$.

Connectivity constraints (2) are used to ensure connectivity of the underlying solution. Basically, whenever two distinct terminals $i, j \in T$ are part of a solution, at least one node from any node-separator $N \in \mathcal{N}(i, j)$ has to be chosen as well, in order to ensure that there exists a path between i and j .

There is a difference between our model and the one considered in [2], where a more general MWCS variant on digraphs has been studied. In this latter variant, a root node needs to be established, i.e., a node r with in-degree zero and such that there is a directed path from r to any node j being part of the solution. Consequently, an additional set of node variables was needed to locate the root, node separators were defined on digraphs, and connectivity constraints have been lifted with respect to root variables. Since our input graphs are undirected, we rely on the undirected model to keep the number of variables as small as possible. Only very recently, connectivity constraints (2) were given more attention in the literature concerning polyhedral studies. In particular, [26] study the connected subgraph polytope, involving node variables only, and show that (2) define facets if and only if N is a minimal node-separator separating i and j .

Node-degree inequalities. The following node-degree inequalities are also valid for our model:

$$y(A_i) \geq \begin{cases} y_i, & \text{if } i \in T \\ 2y_i, & \text{otherwise} \end{cases} \quad (5)$$

where $A_i = \{v \in V \mid \exists \{v, i\} \in E\}$ is the set of all neighboring nodes of i . For terminals, these constraints ensure that at least one of their neighbors is part of the solution (assuming $\sum_{v \in V} y_v \geq 2$, which can be safely assumed after preprocessing single-node solutions). Clearly, they are just a special case of the inequalities (2), but can be used to initialize the model for a branch-and-cut approach. For the remaining nodes, constraints (5) make sure that each such node that belongs to a solution will be used as an intermediate node, i.e., at least two of its neighbors have to be included in the solution as well. These constraints are not implied by (2), in fact they can improve the quality of lower bounds of the original model.

It is sufficient to consider only minimal node-separators in inequalities (2) (since they dominate the remaining ones). In order to derive minimal node separators associated to node-degree constraints, we observe that nodes from $V \setminus T$ that are only adjacent to i and other nodes from A_i do not play a role in connecting i to the remaining terminals. Let J be the set of all such neighbors of a given $i \in T$, i.e., $J = \{j \in A_i \setminus T \mid A_j \subseteq A_i \cup \{i\}\}$. If $V \setminus A_i$ contains other terminals, then $A_i \setminus J$ is a minimal i -separator, i.e., $A_i \setminus J \in \mathcal{N}_i$ and we can correspondingly strengthen the node-degree inequalities (5).

Finally, observe that for potential terminals $i \in T_p$, the node-degree inequality can be lifted as follows:

$$2 \sum_{v \in A_i: c_{vi} < p_i} y_v + \sum_{v \in A_i: c_{vi} \geq p_i} y_v \geq 2y_i \quad \forall i \in T_p.$$

2-Cycle inequalities. Observe the following: if node $i \in V$ is adjacent to a node $j \in T_p$, so that $c_{ij} < p_j$, then if i is part of the optimal solution, j has to be included as well, i.e.,

$$y_i \leq y_j \quad i \in V, j \in T_p, c_{ij} < p_j \quad (6)$$

2.1 Separation of connectivity constraints

Whenever we want to cut off a fractional solution \tilde{y} , we can separate the connectivity cuts (2) by applying a maximum flow algorithm. For each pair of distinct terminals (i, j) such that $\tilde{y}_i + \tilde{y}_j - 1 > 0$, one would need to find a minimum (i, j) -cut in a support digraph D which is constructed as follows. First, each edge $e \in E$ is replaced by two directed arcs. Then, each node $v \in V \setminus \{i, j\}$ is replaced by an arc (v', v'') whose capacity is defined as \tilde{y}_v , all arcs entering v are now directed into v' , and all arcs leaving v are now directed out of v'' . Capacities of these arcs are set to ∞ . Since all arcs except the node-arcs are of infinite capacity, the maximum (i, j) -flow returns the desired violated connectivity constraint.

According to our computational experience, however, the above procedure is rather time consuming (all terminal pairs need to be considered, and for each pair, the maximum flow is calculated). As there is always a certain trade-off between the quality of lower bounds obtained by separating fractional points and the associated computational effort, we refrain from the separation of fractional points in our default implementation.

Consequently, to ensure the validity of our branch-and-cut approach, we need to cut off infeasible integer points enumerated during the branching procedure (or detected by the heuristics of the MIP solver, given that the solver was not provided a complete information about the structure of the problem). Infeasible points are cut off by means of a `LazyCutCallback` in our setting based on the commercial MIP solver IBM ILOG CPLEX. For a given pair of distinct terminal nodes $i, j \in T$ such that $\tilde{y}_i = \tilde{y}_j = 1$, our separation procedure runs in linear time (with respect to $|E|$) and works as outlined below.

To derive our algorithm, we use the following well-known property of node-separators (see, e.g., [17]):

Lemma 1 *Let $N \in \mathcal{N}(i, j)$ be an (i, j) separator for $i, j \in T$, $i \neq j$, and let C_i and C_j be connected components of $G - N$ such that $i \in C_i$, $j \in C_j$. Then N is a minimal (i, j) separator iff every node in N is adjacent to at least one node in C_i and to at least one node in C_j .*

Let \tilde{y} be an integer solution, and let $G_{\tilde{y}} = (V, E_{\tilde{y}})$ denote the “support graph” induced by \tilde{y} , where

$$E_{\tilde{y}} = \{\{i, j\} \in E \mid \tilde{y}_i = \tilde{y}_j = 1\}$$

If \tilde{y} is infeasible, then $G_{\tilde{y}}$ contains at least two connected components, say C_i and C_j , with $i \in C_i$, $j \in C_j$, and $\tilde{y}_i = \tilde{y}_j = 1$. Let $A(C_i)$ be the set of neighboring nodes of C_i in G , i.e., $A(C_i) = \{v \in V \setminus C_i \mid \exists \{u, v\} \in E, u \in C_i\}$. Clearly, $\{i, j\} \notin E$ and hence $A(C_i) \in \mathcal{N}_i$. However, $A(C_i)$ is not necessarily a minimal (i, j) node separator, and Algorithm 1 below describes how to derive a minimal separator starting from $A(C_i)$.

Algorithm 1: A linear time algorithm for detecting a minimal node separator between two components C_i and C_j in $G_{\tilde{y}}$.

Data: Infeasible solution defined by a vector $\tilde{y} \in \{0, 1\}^n$ with $\tilde{y}_i = \tilde{y}_j = 1$, C_i being the connected component of $G_{\tilde{y}}$ containing i , and $j \notin C_i$.

Result: A minimal node separator N that violates inequality (2) with respect to i, j .

Delete all edges in $E[C_i \cup A(C_i)]$ from G

Find the set R_j of nodes that can be reached from j

Return $N = A(C_i) \cap R_j$

In practice, set R_j can be found by just running a standard Breadth-First Search (BFS) on the original graph G , starting from node j , with the additional rule that nodes in $A(C_i)$ are never put in the BFS queue.

Proposition 1 *Algorithm 1 returns a minimal node separator $N \in \mathcal{N}(i, j)$ in time $O(|E|)$.*

Proof By definition of N , i and j are not connected in $G - N$. To see that N is a minimal (i, j) separator, consider $G - N$ and let C'_i and C'_j be two connected components, containing i and j , respectively. Clearly, $C_i \subset C'_i$ and $C'_j = R_j \setminus N$. Hence, by Lemma 1, it follows that N is minimal. \square

In case $p > 2$ connected components exist in $G_{\tilde{y}}$ (each of them containing at least one terminal), one can repeat the procedure described in Algorithm 1 for each pair of distinct components C_i and C_j .

We conclude this section by observing that, for the pure STP case with real terminals only, connectivity constraints translate into

$$y(N) \geq 1, \quad \forall N \in \mathcal{N}$$

where \mathcal{N} is the family of all node separators between arbitrary real terminal pairs. Our model can therefore be interpreted as set covering problem with an exponential number of elements to be covered. As demonstrated by our computational experiments, this property can be exploited to derive specialized set covering heuristics for pure STP instances with uniform costs, with a significant performance boost.

3 Algorithmic framework

The proposed node-based model can be solved by means of a Branch-and-Cut (B&C) algorithm, to be initialized with a high-quality feasible solution and with a suitable set of relevant connectivity constraints.

Our initial MIP model, called the *basic model* in what follows, is made by (1), (3), (4), plus the node-degree inequalities (5) and the 2-cycle inequalities (6).

The overall algorithmic framework is shown in Algorithm 2. In an *initialization phase*, an initial solution pool $\mathcal{S}_{\text{init}}$ is generated by means of some problem-dependent heuristics. In a subsequent *local branching phase*, multiple calls of the B&C algorithm are used to explore the neighborhood of starting solutions chosen from the solution pool $\mathcal{S}_{\text{init}}$. All connectivity constraints separated during this phase are globally valid, hence they are stored in a global cut pool *CutPool*, and added to the initial MIP model before each B&C re-execution. The incumbent solution (denoted by *Sol*) is updated correspondingly.

The local branching phase implements a multiple restart policy (with different seed values), intended to gather relevant information about the problem at hand, namely good primal solutions and a relevant set of connectivity constraints; see e.g. [15] for a recent application of multiple restarts to MIPs with uncertainty. The availability of such information at the root node of each B&C re-execution turns out to be very important, as it triggers a more powerful preprocessing as well as the generation of a bunch of useful general-purpose cuts (in particular, $\{0, 1/2\}$ -cuts [3, 6]) based on the problem formulation explicitly available on input.

The algorithm terminates after proving the optimality, or after reaching the given time limit.

Algorithm 2: Proposed algorithmic framework.

Data: Input graph G , instance of the STP/PCSTP/NWSTP/MWCS, iteration and time limits

Result: A (sub)-optimal solution Sol .

$\mathcal{S}_{\text{init}} = \text{InitializationHeuristics}()$

$k = 1, \text{CutPool} = \emptyset, \mathcal{S} = \emptyset$

Choose Sol from the solution pool $\mathcal{S}_{\text{init}}$.

while ($k \leq \text{maxLBiter}$) and (time limit not exceeded) **do**

$(\text{CutPool}, Sol, \mathcal{S}) = \text{LocalBranching}(\text{CutPool}, Sol, \text{seed}, \mathcal{S})$

$k = k + 1$

 Choose Sol from the solution pool $\mathcal{S}_{\text{init}}$. Change *seed*.

end

$Sol = \text{BranchAndCut}(\text{CutPool}, Sol, \text{TimeLim})$

return Sol

3.1 Local branching

Local branching (LB) has been proposed in [14] as a solution approach that uses the power of a general-purpose MIP solver as a black box to strategically explore promis-

ing solution subspaces. LB is in the spirit of large-neighborhood search metaheuristics, with the main difference that the underlying local search black-box algorithm is a MIP with a specific local branching constraint that restricts the search for an optimal solution within a certain neighborhood of a given reference solution.

The LB framework is built on top of our B&C solver. As already mentioned, our solver deals with two sets of inequalities: those in the *basic model*, that are always part of the model, and connectivity constraints (2) that are dynamically separated and stored in a global *CutPool* to be used in every subsequent B&C call.

Given a reference solution Sol , let $W_1 = \{v \in V \mid v \in Sol\}$ and $W_0 = V \setminus W_1$. The *symmetric* local branching constraint makes sure that the new solution is within radius r (say) from the solution Sol with respect to the Hamming distance between the two solutions, i.e.

$$\sum_{v \in W_0} y_v + \sum_{v \in W_1} (1 - y_v) \leq r$$

Alternatively, one may consider *asymmetric* local branching constraint, requiring that the new solution contains at least $|Sol| - r$ nodes from Sol :

$$\sum_{v \in W_1} (1 - y_v) \leq r. \tag{7}$$

Notice that, for a fixed radius, the neighborhood of the asymmetric version is larger and leads to potentially better solutions—though it is more time consuming to explore. For example, for $r = 0$ the asymmetric version is equivalent to fixing to 1 all the variables from Sol , so that many feasible solutions are still available even in the 0-neighborhood around Sol . After some preliminary tests, we decided to use the asymmetric LB constraint (7) in our implementation, with a small radius r ranging from 10 to 30. Working with a small radius is indeed crucial for the success of proximity methods such as LB, as recently pointed out in [16].

Since the goal of the B&C in this context is to quickly find high-quality solutions, we do not necessarily search for an optimal solution within the given neighborhood, but we rather impose limits on the number of incumbent solutions found (*SolLim*), and a time limit per iteration (*TimeLim*).

The neighborhood is systematically explored by starting with the initial radius r_{min} , and increasing it by r_{delta} each time the subproblem could not provide an improved solution. Each time an improved solution is found, the neighborhood radius is reset to r_{min} . The whole framework is executed until a given number of iterations (*MaxIter*) is reached, or the radius exceeds r_{max} . Note that the radius limit r_{max} , in

combination with a consistent increase of the current radius, implicitly imposes a limit on the overall number of iterations without any improvement.

Algorithm 3: Basic Local Branching.

Data: Starting solution Sol , $CutPool$, \mathcal{S} , $seed$, lower and upper bound for radius (r_{min}, r_{max}), radius step (r_{delta}), maximum n. of iterations $MaxIter$.

Result: Improved solution Sol , enlarged cut pool $CutPool$.

$r = r_{min}, k = 1$

while ($k \leq MaxIter$) and ($r \leq r_{max}$) **do**

 Add local branching constraint (7) centered on Sol with radius r

$(Sol', CutPool') = \text{BranchAndCut}(CutPool, Sol, TimeLim, SolLim)$

 Remove the local branching constraint from the current model

if $cost(Sol') < cost(Sol)$ **then**

$Sol = Sol', r = r_{min}$

$\mathcal{S} = \mathcal{S} \cup Sol$

else

$r = r + r_{delta}$

end

$CutPool = CutPool \cup CutPool'$

$k = k + 1$

end

return ($CutPool, Sol, \mathcal{S}$)

3.2 Benders-like (Set Covering) heuristic

Local branching has a primal nature, in the sense that it produces a sequence of feasible solutions of improved cost. In addition, it needs a starting feasible solution, that in some cases can be time consuming to construct. As a matter of fact, for some very large/hard classes of instances we found that a dual approach is preferable, that produces a sequence of infeasible (typically, disconnected) solutions and tries to repair them to enforce feasibility. Algorithm 4 illustrates a general dual scheme that can be viewed as a heuristic version of the well-known Benders' exact solution approach to general MIPs.

In our experiments, we found that the above approach works very well for uniform STP instances of very large size, for which the standard MIP approach seems not very appropriate as even the LP relaxation of the model takes an exceedingly large computing time to be solved. Our *set-covering based heuristic* is an implementation of Algorithm 4 for these hard instances, and is based on their set covering interpretation. Indeed, as already observed, the basic model turns out to be a compact set covering problem where columns correspond to Steiner nodes, rows to real terminals not adjacent to any other real terminal, and column j covers row i iff $\{i, j\} \in E$. The approach can be outlined as follows.

At each iteration of the while loop, the relaxation to be heuristically solved is constructed through a procedure that automatically extracts a set covering relaxation

Algorithm 4: A conceptual Benders-like heuristic.

```

Data: Time/iteration limits.
Result: Feasible solution  $Sol$ , cut pool  $CutPool$ .
 $Sol =$  dummy solution of very large cost
 $CutPool = \emptyset$ 
while (time/iteration limit not exceeded) do
    Heuristically solve a relaxation of the current model (including all cuts in  $CutPool$ ) and let
     $Sol^R$  be the (possibly disconnected) solution found
    Add the local branching constraint (7) centered on  $Sol^R$  to the unrelaxed model
     $(Sol', CutPool') = \text{BranchAndCut}(CutPool, Sol, TimeLim, SolLim)$ 
    Remove the local branching constraint from the unrelaxed model
    if  $cost(Sol') < cost(Sol)$  then
        |  $Sol = Sol'$ 
    end
     $CutPool = CutPool \cup CutPool'$ 
end
return  $Sol$ 

```

from the current model. This is done by simply (1) projecting all fixed variables (including y variables for hard terminals) out of the model, and (2) skipping all constraints, if any, that are not of type $y(S) \geq 1$ for some node set S .

We then proceed by heuristically solving the set covering relaxation through an implementation of the Caprara-Fischetti-Toth (CFT) heuristic [7, 8]. This is a very efficient set covering heuristic based on Lagrangian relaxation, that is specifically designed for hard/large instances.

Given a hopefully good set covering solution Sol^R , we *repair* it in a very aggressive way by introducing a local branching constraint in asymmetric form with radius $r = 0$, and then by applying our B&C solver (with its ad-hoc connectivity cut separation) with a short time/node limit. As already observed, this local branching constraint in fact corresponds to fixing $y_j = 1$ for all j such that $Sol_j^R = 1$. As a result, the size/difficulty of the MIP model after fixing is greatly reduced, hence the node throughput of B&C solver becomes acceptable even for large instances—while setting a larger radius would not result in a comparable speedup.

All violated connectivity cuts generated by the B&C are added to $CutPool$ and hence to the current model. This makes solution Sol^R (if disconnected) infeasible even for the next set covering relaxation and thus the procedure can be iterated until an overall time/iteration limit is reached.

To improve diversification, our implementation uses the following two mechanisms:

- the procedure that extracts the set covering model makes a random shuffle of the rows/columns, so as to affect in a random way the performance of the CFT heuristic;
- before the repairing phase, we randomly skip (with a uniform probability of 20%) some variable fixings, meaning that approximately 80% of the variables y_j 's with $Sol_j^R = 1$ are actually fixed.

As such, the performance of our final heuristic (though deterministic) is affected by the initial random seed, a property that can be very useful to produce different solutions in a multi-start scheme.

Finally, we observe that our current CFT implementation is sequential and cannot exploit multiple processors. We therefore decided to also run the refining B&C in single-thread mode, thus obtaining an overall sequential code that can be run in parallel and with different random seeds on each single core (in the multi-thread mode). The best found solution is finally returned.

This Benders-like heuristic is embedded in the overall algorithmic framework shown in Algorithm 2 as `InitializationHeuristics` for uniform STP instances on bipartite graphs with a large percentage of terminals. Indeed, these kinds of graph are very regular and the basic model likely gives a reasonable approximation of the STP problem—in the sense that most connectivity constraints are automatically satisfied. In addition, the while-loop in Algorithm 2 that would apply standard local branching after the set-covering based heuristic is unlikely to be effective for these graphs, so we set $maxLBiter = 0$ and skip it in this case.

We conclude this section by observing that the “relaxation” to be heuristically solved in the Benders-like scheme is not intended to produce valid dual bounds, as its purpose is to feed the refining procedure with good (possibly disconnected) solutions. As a matter of fact, one can think of the relaxation as a “blurred” version of the original problem, which retains some of its main features but is not bothered by too many details (namely, connectivity conditions) that would overload the model. Following this general idea, we implemented the following variant of our set-covering based heuristic, which is intended for non-uniform instances on bipartite graphs with a large percentage of terminals. The idea is that we heuristically move edge-cost information into node revenues, so as to get a “blurred” uniform instance on the same underlying graph. To be specific, for each node v we just subtract the average cost of its adjacent edges from the revenue p_v , and in the end we reset all edge costs to $c = 0$. Needless to say, these modified costs/revenues are only used within the CFT heuristic, while the original ones are used in the refining phase.

4 Computational results

Our algorithms are applied to the following problems from the DIMACS challenge: STP, (rooted) PCSTP, MWCS, and also degree-constrained STP (DCSTP). We next summarize the results we obtained on the set of hard (unsolved) cases of the SteinLib [18] instance library. Additionally, we consider non-trivial MWCS instances posted at the website of the DIMACS challenge.

Detailed computational results, covering all nontrivial instances from the challenge are provided in the Appendix. Among these nontrivial instances, we distinguish between *easy* and *difficult* ones. The following criteria are applied: we first run all the instances using the exact B&C approach (without sophisticated heuristics and local branching) with a time-limit of one hour (with a fixed seed value) and record the obtained results. If an instance remains unsolved, it is considered as *difficult*. Our heuristic framework (consisting of the initialization and local branching phase, see

Algorithm 2) is then applied to all *difficult* instances with a time limit of one hour (10 independent runs with different seeds).

The experiments were performed on a cluster of computers, each consisting of 20 cores (2.3 GHz) and with 64GB RAM available for 20 cores. Reported computing times are in wall-clock seconds. To limit the overall time needed to complete our experiments, we decided to allow up to five simultaneous 4-core runs on the same computer, which however implies a significant slowdown due to shared memory.

All algorithms have been implemented in C++ and compiled with GCC 4.9.2. For data structures we used OGDF [24] and the dtree library [13]. CPLEX 12.6 was used as MIP-solver with an imposed memory limit of 12GB RAM.

For solving uniform instances, we used the proposed node-based model (which will be referred to as y -model in the following). For solving non-uniform instances, we used the Steiner arborescence model studied in [22] and referred to as (x, y) -model in the following. In all cases, the models are incorporated into the algorithmic framework shown in Algorithm 2.

4.1 Implementation details of the (x, y) -model

In addition to node variables y defined above, this model uses arc variables x that model the rooted arborescence structure of the solution. For the STP, a random terminal is chosen as a root, and for the unrooted PCSTP an artificial root is introduced and connected to all terminals (with an additional constraint that out-degree of the root is one). This model has been studied in [22], and in this section we explain the main differences between our implementation and the one given in [22].

Initialization Heuristic: A pool of initial feasible solutions is constructed as follows. Several terminals are chosen as root nodes, for each of which a solution is calculated by applying the shortest path Steiner tree heuristic (see, e.g., [4]). In the PCSTP case, for a small number of iterations the set $T_p \cup T_r$ is perturbed and subsets of terminals of different size are considered as fixed terminals T' . Then for each chosen set T' , the same construction heuristic as for the STP is applied. Each solution is also improved through a local search (see below).

For sparse, (almost) planar non-uniform instances of the STP, our framework computes an additional, enhanced initial solution by applying a parallel variant of the partitioning heuristic described in [21]. Based on a randomly chosen solution from the pool, the input graph is partitioned into a set of smaller subgraphs (containing terminals and their closest Steiner nodes). The STP is then solved to optimality (or with a small time limit) on each of these subgraphs independently. The obtained disconnected (and thus, infeasible) solution is then repaired by a shortest-path like heuristic, followed by a local search. When run in a multi-thread mode, the STP solution on each subgraphs is assigned to a single thread.

Primal Heuristic: We apply a few rounds of rounding on the set of y variables, with different threshold values. Rounded up variables again define T' on which we run the shortest path Steiner tree heuristic (with a modified cost function that reflects the LP-values at the current branch-and-bound node) and prune Steiner leaves.

Local Search Heuristics: Solutions obtained by the initialization or the primal heuristic are further improved by applying several local search procedures: Key-Path-Exchange, (Key-)Node-Removal and Node-Insertion. For further implementation details about these heuristics, see [25]. We only accept strictly improving moves, except for the uniform (or almost uniform) instances, for which (due to existence of many symmetric solutions) all moves with non-decreasing objective values are performed.

Additional Valid Inequalities: In [22] the authors observe that the (x,y) -model for PCSTP contains a lot of symmetries, and propose to get rid of them by fixing the terminal node with the smallest index (among those taken in the solution) to be the root node. This is imposed by adding additional asymmetry constraints. The latter constraints added by [22] were given in a disaggregated form, whereas in our implementation we add their stronger, aggregated variant: $\sum_{i>j} x_{ri} \leq 1 - y_j$, for all $j \in T$. We also add 2-cycle inequalities (6) to our starting model.

Separation Algorithms: For fractional solutions, connectivity constraints are separated by the calculation of the maximum flow (implemented through `UserCutCallback` of CPLEX). In order to find more diversified cuts as early as possible, we first add a single violated cut between the root and a single terminal, then we update the capacities on the arcs involved in this cut-set, and then iterate until all terminals are separated. Among the two cut-sets returned by the maximum flow algorithm of [9], we choose the one closer to the terminal (cut-sets closer to the root typically involve similar edges, and hence may imply many redundant cuts). For sparse or randomly generated non-uniform instances, the (x,y) -model provides very strong lower bounds, so that in most cases branching is not needed. On the other hand, for more difficult instances with large LP-gaps, *tailing off* plays a crucial role in the performance of the B&C algorithm. Rather than staying too long at the root node of the B&C tree (and separating too many cuts), we resort to branching if the lower bound does not improve much. During branching, we refrain from separation of fractional solutions.

To further speed up separation we skip the maximum flow calculation for terminals that are already reachable from the root in the support graph. We do not apply nested cuts, but instead separate at most one connectivity constraint for each terminal per iteration. Avoiding the separation of fractional solutions while branching and adding only one cut per terminal seemed to perform best on average in preliminary experiments.

For the PCSTP we only separate cuts between the root and those terminals i such that $y_i \geq 0.5$ in the given fractional solution to separate. Furthermore, instead of adding a connectivity constraint as in [22], we may replace the right-hand side with one if the revenue on the sink side of the cut is larger than the objective of the incumbent solution. This implies that in an optimal solution, at least one more node on the sink side has to be connected to the root.

Infeasible integer solutions are separated by searching for connected components in the support graph. For each subset S inducing a connected component of an infeasible solution, a connectivity constraint $x(\delta^+(S)) \geq y_i$ is added to the model, for $i \in S$ with the highest revenue.

4.2 Initial filter

Our implementation also contains a *filter* that analyzes the structure and costs of the input graph. According to these properties, the algorithm decides the actual MIP model, as well as the kind of initialization heuristics and preprocessing to apply. We note that in the proposed algorithm these rules are only used for the problem types STP, PCSTP and MWCS (which are transformed to their PCSTP representation). DC-STP and RPCSTP instances are solved solely by B&C on the (x,y) -model without any sophisticated settings.

Table 1 lists general filter rules for model and parameter selection. By default, the y -model is applied to instances with uniform edge costs, while the (x,y) -model is applied to all others. For sparse, uniform graphs with a relatively low number of terminals, we switch to the (x,y) -model, as preliminary experiments have shown that the y -model is less efficient for this specific class. For the y -model, the LP at the root node is restarted two times in order to enforce the generation of general purpose cutting planes.

For the (x,y) -model a tailing-off bound is specified, which defines the allowed ratio between the lower bounds of two consecutive B&C iterations in the root node. In addition, a tailing-off tolerance value is defined, which specifies the number of times this bound is allowed to be violated in a row. If the number of violations exceeds the specified tolerance-parameter, the algorithm switches from separation of fractional solutions to branching. The two parameters are chosen based on graph density. Instances are roughly divided into two classes, sparse ($|E|/|V| \leq 3$) and dense ($|E|/|V| > 3$) graphs. The tailing-off bound is only activated for dense graphs, while for sparse graphs we avoid branching as long as possible. In addition, we identify very dense ($|E|/|V| > 5$) instances, in which we set a lower tailing-off tolerance than in the regular case.

To increase the effectiveness of our local search procedure, we allow moves with non-decreasing objective for instances whose edge weights are located within a small range, which produces better solutions, but is also more time consuming.

Table 1 General filter rules. An instance is labeled as `uniform` if all edges have the same weight. The attribute `weightRange` is defined as the difference between the minimum and maximum edge weight. `ratioT` denotes the terminal ratio $|T|/|V|$. An instance is labeled with `sparse` if $|E|/|V| \leq 3$, otherwise with `dense`. An instance is assigned the additional attribute `verydense` if $|E|/|V| > 5$.

Rule	Applied settings
<code>uniform</code>	→ y -model
<code>uniform</code> \wedge <code>sparse</code> \wedge <code>ratioT</code> < 0.1	→ (x,y) -model
\neg <code>uniform</code>	→ (x,y) -model
<code>dense</code>	→ use tailing-off bound, high tolerance
<code>verydense</code>	→ use tailing-off bound, low tolerance
<code>weightRange</code> < 10	→ allow non-improving moves during local search (cf. Section 4.1)

Table 2 lists filter rules that select problem-specific settings and algorithms for STP instances. Given that the (x, y) -model has been selected by the filter, the model may be initialized by a set of connectivity cuts generated by dual ascent algorithm [27], and with a starting solution generated by a parallel implementation of a partitioning heuristic (see [21]). Based on the terminal ratio, the dual ascent connectivity cuts are either added at the beginning of the B&C or separated. Similarly, flow-balance constraints and GSECs of size two may be chosen to be separated if the graph is large, sparse and has a low terminal ratio. Dense STP instances are instead preprocessed by using the Special Distance Test [12].

Table 2 STP-specific rules. An instance is labeled as `big` if $|V| > 10000$ and as `dense` if $|E|/|V| > 3$. `ratioT` denotes $|T|/|V|$. `xy-model` is true if the (x, y) -model has been selected.

Rule	Applied settings
<code>xy-model</code> \wedge <code>ratioT</code> < 0.01	\rightarrow separate dual ascent cuts
<code>xy-model</code> \wedge <code>ratioT</code> ≥ 0.01	\rightarrow add all dual ascent cuts
<code>xy-model</code> \wedge <code>ratioT</code> < 0.1 \wedge <code>sparse</code> \wedge <code>big</code>	\rightarrow separate flow-balance constr., GSECs of size 2
<code>xy-model</code> \wedge <code>big</code> \wedge <code>sparse</code>	\rightarrow partition-based construction heuristic
<code>verydense</code>	\rightarrow preprocessing (special distance test)

Table 3 lists additional rules for the selection of initialization heuristics. In the case of bipartite graphs, the set-covering heuristic is applied for both uniform and non-uniform instances to generate high-quality starting solutions. For non-uniform instance, each node v is assigned the weight $\frac{1}{|\delta(v)|} \sum_{e \in \delta(v)} c_e$ (“blurred” set-covering heuristic). We note that in the case of PCSTP instances we choose to apply the blurred set-covering heuristic also for the large non-uniform instances, as there exists a significant slowdown of the (x, y) -model for larger instances, at which point the blurred set-covering heuristic performs better. By default, the following settings are also applied when executing local branching: The primal heuristic is executed in every branch-and-bound node for the (x, y) -model and in each 1000 nodes for the y -model. The GSECs of size two in the (x, y) -model are separated and stored in the cut pool. The local branching time limit per iteration is set to 120 seconds. The neighborhood radius is initialized with 10 and increased by 10 per iteration up to a maximum of 30.

Table 3 Initialization heuristic rules. An instance is labeled with `bipartite` if the instance graph is bipartite with respect to the node sets $V \setminus T$ and T . An instance is labeled as `small` if $|V| < 1000$

Rule	Applied settings
<code>bipartite</code> \wedge <code>uniform</code>	\rightarrow set-covering heuristic
<code>bipartite</code> \wedge \neg <code>uniform</code> \wedge <code>stp</code>	\rightarrow blurred set-covering heuristic
<code>hypercube</code> \wedge \neg <code>uniform</code> \wedge \neg <code>small</code> \wedge <code>pcstp</code>	\rightarrow blurred set-covering heuristic
\neg <code>bipartite</code>	\rightarrow local branching

4.3 Results for uniform STP instances

4.3.1 PUC dataset:

A preliminary study of hypercube (hc) instances [23] confirmed their very high symmetry even in the space of node variables. Indeed, after removing the y variables of terminal nodes (that are all fixed to one), the symmetry group produces a single orbit that contains all the remaining variables. This implies that one can w.l.o.g. fix $y_v = 1$ for an arbitrary Steiner node $v \in V \setminus T$, thus significantly reducing the symmetry group—that however remains nontrivial.

It is well known that symmetry plays an prominent (negative) role when solving an instance to optimality through enumerative methods. However, symmetry can even be of help when heuristics come into play, as the impact of bad heuristic choices is mitigated somehow by symmetry. As our main goal was to provide improved heuristic solutions for the unsolved cases, we did not investigate symmetry properties any further, and actually preferred not to apply the symmetry-breaking fixing $y_v = 1$ mentioned above.

Instead, we concentrated on the set-covering heuristic described in Section 3.2. Table 4 reports the very preliminary results we obtained by an over-simplified version of this approach, where the CFT set-covering heuristic was only available to us as a binary executable—namely, the TURNI commercial implementation [19,20]. As a result, this piece of software could not be fully embedded into our code and we had to use it off-line through input/output text files. In addition, we could not iterate the approach, i.e., we only refined a single set-covering solution. TURNI was therefore run only once and with a short time limit—one minute was enough for finding very good solutions for all cases, except for hc11u for which we allowed for 30 minutes. To our pleasant surprise, even in this very simplified setting the repairing phase turned out to be very fast and effective in producing connected solutions whose cost was very close to the cost of the set covering solution.

Table 4 Our very preliminary results for unsolved uniform STP instances of the PUC class

Instance	PreviousBest	NewBest	TURNI time (s.)	Repair time (s.)
bip52u	234	234	60.0	2.14
bip62u	220	219	60.0	0.03
bipa2u	341	338	60.0	0.05
hc9u	292	292	60.00	1.53
hc10u	581	575	60.00	4.52
hc11u	1154	1145	1800.00	0.07
hc12u	2275	2267	60.00	6.68

Significantly better results were obtained at a later stage of our experiments, when we implemented the CFT heuristic ourselves and were able to fully embed it into the Benders-like scheme of Algorithm 4; see Table 8.

4.3.2 GAPS and SP datasets:

For the subgroup “skutella” (s3-s5) of the artificially generated uniform instance set GAPS for the STP, LP-gaps of the standard connectivity-based (x,y) -model are large. Standard MIP approaches for these instances have difficulties in closing the integrality gap. Table 5 reports our results obtained on instances s1 to s5 from GAPS, and clearly demonstrates the power of our y -model.

Table 5 Uniform STP instances from the GAPS dataset. Proven optimal solutions in boldface. Column Time gives the computing time for proving the optimality (or, the time limit, otherwise). Columns UB and LB show upper and lower bounds obtained by the (x,y) -model, within the time limit of two hours, respectively. (*) Reached memory limit of 12GB after the specified time limit.

Instance	V	E	T	y-model		(x,y)-model			
				OPT	Time (s.)	UB	LB	Gap	Time (s.)
s1	64	192	32	10	0.03	10	10	0.0%	0.01
s2	106	399	50	73	0.04	73	73	0.0%	1.36
s3	743	2947	344	514	0.15	514	505	1.78%	1090.61*
s4	5202	20783	2402	3601	1.31	3601	3523	2.21%	3444.81*
s5	36415	145635	16808	25210	22.28	25210	24056	4.80%	7200.00

Additionally, for two previously unsolved instances from the set SP (with uniform edge costs as well), namely w13c29 and w23c23, we provide optimal values. For these two latter instances, both models were able to prove the optimality, with significant speed-ups achieved by the y -model. Table 5 reports results using ten different seeds. Here, the best running time for the (x,y) -model is competitive with the y -model, but in the average case it is much slower. This is due to the fact that solving the instance to optimality is highly dependent on finding an optimal solution.

Table 6 Uniform STP instances from the SP dataset. Proven optimal solutions in boldface. Previous best known solutions given in brackets. Runs have been performed using ten seeds. The Time columns give the best, average and the standard deviation for running time. For the y -model, no primal heuristics have been used besides the ones by CPLEX. For the (x,y) -model, we use local search (cf. Section 4.1).

Instance	V	E	T	OPT	y-model Time (s.)			(x,y)-model Time (s.)		
					BEST	AVG	STD	BEST	AVG	STD
w13c29	783	2262	406	507 (508)	0.31	0.87	0.46	14.46	38.28	30.04
w23c23	1081	3174	552	689 (694)	43.91	132.59	59.96	183.93	2600.15	1362.61

4.4 Results for MWCS instances

The MWCS can be transformed into the PCSTP with uniform edge costs (see [2]). We have tested both y - and (x,y) -model on the MWCS dataset, and the obtained computational results (for the most challenging instances) are reported in Table 7. Results on the JMPALMK dataset are available in the Appendix - those instances have all been solved within one second.

Table 7 ACTMOD PCSTP instances from the DIMACS website. For these runs no primal heuristics have been used besides the ones by CPLEX.

Instance	V	E	T	OPT	y-model		(x,y)-model	
					Time (s.)	Gap	Time (s.)	Gap
drosophila001	5226	93394	5226	8273.98263	7.98	0.00	86.12	0.00
drosophila005	5226	93394	5226	8121.313578	9.48	0.00	76.32	0.00
drosophila0075	5226	93394	5226	8039.859460	7.45	0.00	68.48	0.00
HCMV	3863	29293	3863	7371.536373	0.96	0.00	6.11	0.00
lymphoma	2034	7756	2034	3341.890237	0.28	0.00	1.24	0.00
metabol_expr_mice_1	3523	4345	3523	11346.927189	5965.76	0.00	1.08	0.00
metabol_expr_mice_2	3514	4332	3514	16250.235191	1.21	0.00	1.57	0.00
metabol_expr_mice_3	2853	3335	2853	16919.620407	4.00	0.00	0.89	0.00

4.5 Summary of improved results for STP and PCSTP instances

Tables 8 to 13 list the most important results after applying the set-covering heuristic or local branching. These datasets have been selected from the set of instances unsolved after one hour by B&C (detailed results are available in the Appendix).

Each run has been computed by starting the algorithm with ten different seeds, each with one hour time limit. The choice between the local branching and set-covering heuristic is performed by the filter as described in Section 4.2. The final B&C step of the framework is always skipped.

Each table is structured as follows: The first four columns list the instance name, the number of nodes, edges and terminals. The next pair of columns (BEST) shows objective value and time of computation for the best found solution. The following pairs (AVG and STD) list the average and standard deviation of these two values over all ten runs. For tests on previously known STP datasets (PUC, I640) the column ‘Impr.’ lists the improvement w.r.t. the best known published solution (by August 1st, 2014) according to the DIMACS challenge website (<http://dimacs11.cs.princeton.edu/instances/bounds20140801.txt>). In all other cases, the improvement is given w.r.t. the best primal solutions produced during the exact runs after one hour.

Table 8 PUC STP instances. Column ‘Impr.’ shows the improvement w.r.t. the previous best known values published on the DIMACS challenge website [10]. Results computed by the set-covering heuristic in the case of hypercube (hc) and bipartite (bip) instances, otherwise through (x,y) -model-based local branching. Improved solutions given in boldface.

Instance	V	E	T	BEST		AVG		STD		Impr.
				UB	Time	UB	Time	UB	Time	
bip42p	1200	3982	200	24657	38.36	24660.80	664.65	2.04	1125.16	0
bip52p	2200	7997	200	24549	805.57	24566.90	1403.00	13.91	1357.42	-14
bip52u	2200	7997	200	233	1390.10	233.80	287.94	0.42	597.96	1
bip62p	1200	10002	200	22906	3.60	22907.00	55.37	1.05	80.78	-36
bip62u	1200	10002	200	219	6.21	219.00	12.28	0.00	5.04	1
bipa2p	3300	18073	300	35355	547.18	35360.90	1342.88	4.38	879.59	24
bipa2u	3300	18073	300	337	185.06	337.00	310.89	0.00	215.22	4
cc10-2p	1024	5120	135	35257	875.45	35353.20	704.89	75.12	705.21	122
cc10-2u	1024	5120	135	342	206.27	342.60	818.02	0.52	1078.16	0
cc11-2p	2048	11263	244	63680	744.33	63895.70	976.37	103.40	726.59	146
cc11-2u	2048	11263	244	615	1388.68	616.90	1203.83	0.99	951.61	-1
cc12-2p	4096	24574	473	122166	1884.06	123096.00	1912.64	467.97	799.12	-1060
cc12-2u	4096	24574	473	1183	1559.46	1186.30	1937.01	1.77	804.16	-4
cc3-10p	1000	13500	50	12784	3471.19	12826.20	1801.62	43.46	1139.72	76
cc3-10u	1000	13500	50	125	61.88	125.00	615.79	0.00	683.51	0
cc3-11p	1331	19965	61	15599	458.95	15633.30	812.14	35.44	965.08	10
cc3-11u	1331	19965	61	153	29.69	153.00	269.26	0.00	580.93	0
cc3-12p	1728	28512	74	18879	1290.13	18936.60	1771.08	31.50	1139.80	-41
cc3-12u	1728	28512	74	185	59.70	185.00	900.54	0.00	985.39	1
cc3-5p	125	750	13	3661	0.75	3661.00	10.53	0.00	13.25	0
cc3-5u	125	750	13	36	0.01	36.00	0.02	0.00	0.01	0
cc5-3p	243	1215	27	7299	16.41	7299.00	238.22	0.00	208.56	0
cc5-3u	243	1215	27	71	0.65	71.00	2.61	0.00	1.24	0
cc6-3p	729	4368	76	20340	1266.76	20395.90	1543.97	46.02	983.95	116
cc7-3p	2187	15308	222	57080	1385.54	57328.70	1197.71	153.94	888.00	8
cc7-3u	2187	15308	222	551	383.80	554.10	1267.21	1.52	1078.48	1
cc9-2p	512	2304	64	17202	1603.44	17274.40	1579.81	28.51	984.36	94
cc9-2u	512	2304	64	167	14.95	167.30	753.10	0.48	1018.65	0
hc7p	128	448	64	7905	2479.97	7915.80	875.57	6.00	746.95	0
hc8p	256	1024	128	15337	2494.80	15349.50	1057.47	7.49	1128.89	-15
hc8u	256	1024	128	148	0.05	148.00	0.07	0.00	0.01	0
hc9p	512	2304	256	30319	1232.02	30342.30	1824.93	14.08	777.70	-61
hc9u	512	2304	256	292	0.27	292.00	0.39	0.00	0.08	0
hc10p	1024	5120	512	59981	267.51	60041.30	1013.51	33.38	816.95	513
hc10u	1024	5120	512	575	11.17	575.00	86.97	0.00	85.92	6
hc11p	2048	11264	1024	119500	3327.76	119533.00	1708.94	35.11	1129.07	279
hc11u	2048	11264	1024	1145	663.27	1145.40	1319.21	0.52	873.14	9
hc12p	4096	24576	2048	236267	2782.93	236347.10	2514.01	55.44	565.26	682
hc12u	4096	24576	2048	2261	2756.85	2262.50	2805.22	1.27	747.01	14

Table 9 I640 STP instances. Column ‘Impr.’ shows the improvement w.r.t. the previous best known values published on the DIMACS challenge website [10]. Results computed by local branching using the (x,y) -model. Improved solutions given in boldface.

Instance	V	E	T	BEST		AVG		STD		Impr.
				UB	Time	UB	Time	UB	Time	
i640-311	640	4135	160	35766	117.65	35779.00	1520.97	21.75	1219.66	0
i640-312	640	4135	160	35768	1410.35	35793.20	1478.45	25.38	1104.32	3
i640-313	640	4135	160	35535	292.62	35538.20	923.67	4.13	921.39	0
i640-314	640	4135	160	35533	1610.03	35547.00	1673.70	12.53	679.53	5
i640-315	640	4135	160	35720	156.24	35733.50	866.76	21.87	695.92	21

Table 10 PUCN STP instances (uniform version of the PUC code-coverage dataset). Column ‘Impr.’ shows the improvement with respect to the best solution computed by B&C. Results computed by local branching on the y -model. Time limits set to one hour.

Instance	V	E	T	BEST		AVG		STD		Impr.
				UB	Time	UB	Time	UB	Time	
cc10-2n	1024	5120	135	180	89.50	181.00	690.13	0.67	952.37	2
cc11-2n	2048	11263	244	327	39.84	328.00	658.63	0.67	904.00	4
cc12-2n	4096	24574	473	617	930.40	621.80	933.09	2.57	639.18	9
cc3-10n	1000	13500	50	75	0.31	75.00	1.19	0.00	1.03	0
cc3-11n	1331	19965	61	92	0.53	92.00	1.24	0.00	0.81	0
cc3-12n	1728	28512	74	111	2.07	111.00	6.61	0.00	3.83	0
cc7-3n	2187	15308	222	290	51.42	290.70	449.57	0.82	410.34	3
cc9-2n	512	2304	64	98	73.52	98.90	598.66	0.57	575.48	2

Table 11 PUCNU PCSTP instances (uniform PCSTP version of the PUC dataset). Column ‘Impr.’ shows the improvement with respect to the best solution computed by B&C. Results computed by local branching on the y -model. Time limits set to one hour.

Instance	V	E	T	BEST		AVG		STD		Impr.
				UB	Time	UB	Time	UB	Time	
bipa2nu	3300	18073	300	324	17.99	324.00	534.20	0.00	580.83	1
cc10-2nu	1024	5120	135	168	253.05	169.20	1059.52	0.63	1140.60	3
cc11-2nu	2048	11263	244	305	158.46	306.50	612.65	1.27	562.96	7
cc12-2nu	4096	24574	473	568	1427.16	571.00	921.48	1.49	893.82	13
cc3-10nu	1000	13500	50	61	0.52	61.00	4.21	0.00	5.57	0
cc3-11nu	1331	19965	61	79	11.62	79.30	365.37	0.48	571.75	1
cc3-12nu	1728	28512	74	95	12.23	95.00	412.33	0.00	538.46	1
cc7-3nu	2187	15308	222	271	565.00	274.10	639.51	1.29	751.37	9

Table 12 H and H2 PCSTP instances. Column ‘Impr.’ shows the improvement with respect to the best solution computed by B&C. Results computed by local branching on the y -model for uniform instances (hc*u(2)) and on the (x,y) -model for non-uniform instances (hc*p(2)). For large non-uniform hypercubes (starting from hc10p(2)), the blurred version of the set-covering heuristic is applied instead. Time limits set to one hour.

Instance	V	E	T	BEST		AVG		STD		Impr.
				UB	Time	UB	Time	UB	Time	
hc8p	256	1024	256	15206	537.28	15228.50	1514.80	15.69	1234.15	58
hc9p	512	2304	512	30043	3062.51	30084.00	1967.64	25.21	1073.46	209
hc10p	1024	5120	1024	59866	919.12	59965.20	1078.30	46.36	922.72	728
hc10u	1024	5120	1024	559	2349.42	559.90	773.58	0.32	1081.27	3
hc11p	2048	11264	2048	119191	3600.53	119377.60	1851.08	96.35	1114.04	1806
hc11u	2048	11264	2048	1116	2284.54	1117.20	1568.57	0.79	1193.87	4
hc12p	4096	24576	4096	235860	2542.60	236103.40	2417.05	147.20	685.50	4621
hc12u	4096	24576	4096	2221	310.53	2223.10	1078.51	1.45	1035.82	5
hc8p2	256	1024	256	15231	178.35	15255.00	1271.71	17.83	996.89	60
hc9u2	512	2304	512	190	10.01	190.00	15.40	0.00	3.03	0
hc10p2	1024	5120	1024	59930	2119.75	59966.40	1849.67	21.55	983.48	601
hc10u2	1024	5120	1024	380	47.66	380.30	1287.84	0.48	1119.22	2
hc11p2	2048	11264	2048	119236	330.52	119381.80	1106.48	89.67	703.96	1762
hc11u2	2048	11264	2048	750	2009.52	751.40	1596.31	0.70	1131.72	11
hc12p2	4096	24576	4096	235687	3172.89	235985.00	2099.70	217.33	1075.97	4793
hc12u2	4096	24576	4096	1494	101.86	1494.10	671.75	0.32	1031.32	13

Table 13 I640 PCSTP instances. Column ‘Impr.’ shows the improvement with respect to the best solution computed by B&C. Results computed by local branching on the (x,y) -model. Time limits set to one hour.

Instance	V	E	T	BEST		AVG		STD		Impr.
				UB	Time	UB	Time	UB	Time	
i640-311	640	4135	160	33503	1062.15	33518.80	2044.98	20.39	1093.29	44
i640-312	640	4135	160	32721	152.92	32721.00	1846.09	0.00	1147.93	69
i640-313	640	4135	160	32401	200.14	32403.00	1302.32	6.32	906.75	0
i640-314	640	4135	160	32871	1572.85	32893.10	1590.90	20.36	1147.87	88
i640-315	640	4135	160	32616	491.25	32631.30	1621.18	9.31	1059.35	50
i640-321	640	204480	160	28787	309.34	28788.30	1121.02	1.57	1234.34	16
i640-322	640	204480	160	28458	943.72	28461.20	989.01	3.88	828.89	9
i640-323	640	204480	160	28153	9.22	28153.50	667.85	0.53	687.52	4
i640-324	640	204480	160	28746	44.53	28747.50	864.61	0.97	1116.14	14
i640-325	640	204480	160	28385	31.32	28386.00	1588.58	0.94	1588.99	1
i640-341	640	40896	160	29702	1585.79	29720.50	1679.07	14.21	984.23	37
i640-342	640	40896	160	29806	1591.95	29828.80	1889.09	16.78	915.53	32
i640-343	640	40896	160	30056	307.35	30059.90	1093.91	8.23	691.12	0
i640-344	640	40896	160	29921	365.16	29943.80	1370.59	12.75	913.09	21
i640-345	640	40896	160	30004	2580.47	30029.60	2289.36	18.68	1141.93	78

Note that the tables report the results obtained within the time limit of one hour only. By extending the time limit, or by using more than four threads in parallel, the obtained values can further be improved. For example, for the most difficult ones we obtain:

Problem	Instance	Best UB	Time	#Threads
STP	hc11u	1144	474	8
STP	hc12u	2256	4817	8
STP	hc12p	236158	4411	4
PCSTP	hc12u2	1492	632	16

Furthermore, we were able to prove optimality for the following difficult instances, where according to the SteinLib website [18] the instance i640-313 has previously been unsolved (for the PCSTP version of i640-313, we report similar performance):

Problem	Instance	OPT	Time (s.)	#Threads	#Memory (GB)
STP	cc6-3u	197	390	4	12
STP	i640-313	35535	16012	20	32
PCSTP	i640-313	32401	15358	20	32

5 Conclusions

We have presented a simple model for the Steiner tree problem, involving only node variables. Besides drastically reducing the number of the required variables, the removal of edge variables avoids a number of issues related to equivalent (possibly symmetrical) trees spanning a same nodeset. In this view, we are “thinning out” the usual edge-based model with the aim of getting a more agile framework. Our model is mainly intended for instances with uniform edge costs, but it can be extended to the general case by a simple edge splitting mechanism (possibly to be applied on the fly). Computational results show that our approach can dramatically improve the performance of an exact solver, and in some cases converts very hard problems into trivial-to-solve ones. At the recent DIMACS challenge on Steiner trees (see [11]), the approach presented in this article won most of the categories of both, exact and heuristic challenge, for the STP, PCSTP, MWCS and DCSTP.

Acknowledgements

The research of the Padova’s group was supported by the University of Padova (Progetto di Ateneo “Exploiting randomness in Mixed Integer Linear Programming”), and by MiUR, Italy (PRIN project “Mixed-Integer Nonlinear Optimization: Approaches and Applications”). The research of the Vienna’s group was supported by the Austrian Research Fund (FWF, Projects P 26755-N19 and I 892-N23). The work of the last author was supported by a STSM Grant from COST Action TD1207; the authors would like to acknowledge networking support by the same COST action.

References

1. 11th DIMACS Implementation Challenge in Collaboration with ICERM: Steiner tree problems (2014). URL <http://dimacs11.cs.princeton.edu/home.html>
2. Álvarez-Miranda, E., Ljubić, I., Mutzel, P.: The maximum weight connected subgraph problem. In: *Facets of Combinatorial Optimization: Festschrift for Martin Grötschel*, pp. 245–270. Springer (2013)
3. Andreello, G., Caprara, A., Fischetti, M.: Embedding $\{0, 1/2\}$ -cuts in a branch-and-cut framework: A computational study. *INFORMS Journal on Computing* **19**(2), 229–238 (2007)
4. de Aragão, M.P., Werneck, R.F.F.: On the implementation of MST-based heuristics for the Steiner problem in graphs. In: D.M. Mount, C. Stein (eds.) *ALENEX, Lecture Notes in Computer Science*, vol. 2409, pp. 1–15. Springer (2002)
5. Backes, C., Rurainski, A., Klau, G.W., Müller, O., Stöckel, D., Gerasch, A., Küntzer, J., Maisel, D., Ludwig, N., Hein, M., Keller, A., Burtscher, H., Kaufmann, M., Meese, E., Lenhof, H.P.: An integer linear programming approach for finding deregulated subgraphs in regulatory networks. *Nucleic Acids Research* **40**, 1–13 (2012)
6. Caprara, A., Fischetti, M.: $\{0, 1/2\}$ -Chvátal-Gomory cuts. *Mathematical Programming* **74**(3), 221–235 (1996)
7. Caprara, A., Fischetti, M., Toth, P.: A heuristic algorithm for the set covering problem. In: *Integer Programming and Combinatorial Optimization*, pp. 72–84. Springer Berlin Heidelberg (1996)
8. Caprara, A., Fischetti, M., Toth, P.: A heuristic method for the set covering problem. *Operations Research* **47**(5), 730–743 (1999)
9. Cherkassky, B.V., Goldberg, A.V.: On implementing push-relabel method for the maximum flow problem **19**, 390–410 (1994.)
10. Best bounds as of August 1, 2014 for SteinLib instances (2014). <http://dimacs11.cs.princeton.edu/instances/bounds20140801.txt>
11. Results of the 11th DIMACS Competition on Steiner Tree Problems (2015). <http://dimacs11.cs.princeton.edu/contest/results/results.html>
12. Duin, C.: Preprocessing the Steiner Problem in Graphs. In: D.Z. Du, J. Smith, J. Rubinstein (eds.) *Advances in Steiner Trees, Combinatorial Optimization*, vol. 6, pp. 175–233. Springer US (2000)
13. Eisenstat, D.: dtree: dynamic trees à la carte (2014). <http://www.davideisenstat.com/dtree/>
14. Fischetti, M., Lodi, A.: Local branching. *Mathematical Programming* **98**(1-3), 23–47 (2003)
15. Fischetti, M., Monaci, M.: Cutting plane versus compact formulations for uncertain (integer) linear programs. *Mathematical Programming Computation* **4**(3), 239–273 (2012)
16. Fischetti, M., Monaci, M.: Proximity search for 0-1 mixed-integer convex programming. *Journal of Heuristics* **20**(6), 709–731 (2014)
17. Golombic, M.C.: *Algorithmic Graph Theory and Perfect Graphs (Annals of Discrete Mathematics, Vol 57)*. North-Holland Publishing Co., Amsterdam, The Netherlands (2004)
18. Koch, T., Martin, A., Voß, S.: SteinLib: An updated library on Steiner tree problems in graphs. Tech. Rep. ZIB-Report 00-37, Konrad-Zuse-Zentrum für Informationstechnik Berlin, Takustr. 7, Berlin (2000). URL <http://elib.zib.de/steinlib>
19. Kroon, L., Fischetti, M.: Crew scheduling for netherlands railways destination: customer. In: *Computer-aided scheduling of public transport*, pp. 181–201. Springer Berlin Heidelberg (2001)
20. Kroon, L., Huisman, D., Abbink, E., Fioole, P.J., Fischetti, M., Maróti, G., Schrijver, A., Steenbeek, A., Ybema, R.: The new Dutch timetable: The OR revolution. *Interfaces* **39**(1), 6–17 (2009)
21. Leitner, M., Ljubić, I., Luijpersbeck, M., Resch, M.: A Partition-Based Heuristic for the Steiner Tree Problem in Large Graphs. In: M.J. Blesa, C. Blum, S. Voß (eds.) *Hybrid Metaheuristics - Proceedings, Lecture Notes in Computer Science*, vol. 8457, pp. 56–70. Springer (2014)
22. Ljubić, I., Weiskircher, R., Pferschy, U., Klau, G.W., Mutzel, P., Fischetti, M.: An algorithmic framework for the exact solution of the prize-collecting Steiner tree problem. *Mathematical Programming* **105**(2-3), 427–449 (2006)
23. Rosseti, I., de Aragão, M.P., Ribeiro, C.C., Uchoa, E., Werneck, R.F.: New benchmark instances for the Steiner problem in graphs. In: *Extended Abstracts of the 4th Metaheuristics International Conference (MIC'2001)*, pp. 557–561. Porto (2001)
24. OGDf: The Open Graph Drawing Framework (2015). <http://www.ogdf.net/>
25. Uchoa, E., Werneck, R.F.: Fast local search for Steiner trees in graphs. In: G.E. Blelloch, D. Halperin (eds.) *ALENEX*, pp. 1–10. SIAM (2010)
26. Wang, Y., Buchanan, A., Butenko, S.: On imposing connectivity constraints in integer programs (2015). Submitted. Available at www.optimization-online.org/DB_HTML/2015/02/4768.html

-
27. Wong, R.T.: A dual ascent approach for Steiner tree problems on a directed graph. *Mathematical Programming* **28**(3), 271–287 (1984). DOI 10.1007/BF02612335