



A Partition-and-Merge Algorithm for Solving the Steiner Tree Problem in Large Graphs

Ming Sun¹, Xinyu Wu¹, Yi Zhou^{1(✉)}, Jin-Kao Hao², and Zhang-Hua Fu^{3,4}

¹ University of Electronic Science and Technology of China, Chengdu, China
{summ,zhou.yi}@uestc.edu.cn, wu.xinyu@outlook.com

² LERIA, Université d'Angers, 2 Boulevard Lavoisier, 49045 Angers, France
jin-kao.hao@univ-angers.fr

³ Shenzhen Institute of Artificial Intelligence and Robotics for Society,
Shenzhen, China

⁴ The Chinese University of Hong Kong, Shenzhen, China
fuzhanghua@cuhk.edu.cn

Abstract. The Steiner tree problem aims to determine a minimum edge-weighted tree that spans a given set of terminal vertices from a given graph. In the past decade, a considerable number of algorithms have been developed to solve this computationally challenging problem. However, existing algorithms typically encounter difficulties for solving large instances, i.e., graphs with a high number of vertices and terminals. In this paper, we present a novel partition-and-merge algorithm to effectively solve this problem in large graphs. The algorithm breaks the input network into small subgraphs and then merges the subgraphs in a bottom-up manner. In the merging procedure, partial Steiner trees in the subgraphs are also created and optimized by efficient local optimization. When the merging procedure ends, the algorithm terminates and reports the final solution for the input graph. We evaluated the algorithm on a wide range of benchmark instances, showing that the algorithm outperforms the best-known algorithms on large instances and competes favorably with them on small or medium-sized instances.

Keywords: Steiner tree problem · Network Design · Local search · Partition-and-merge · Large graphs

1 Introduction

The Steiner tree problem (STP) is a fundamental network design problem to determine the topology of optical networks [14], telephone networks [7], or even multimedia networks [17]. The input of the classic STP consists of an undirected edge-weighted graph (or network) G with a vertex set $V(G)$, an edge set $E(G)$, a non-negative weight function on its edges $c_G : E(G) \rightarrow \mathbb{R}^+$, and a set of *terminals* $A \subseteq V(G)$. The problem is to determine a minimum Steiner tree, that is, a

tree T spanning all terminals in A (and possibly other vertices) and minimizes $\sum_{e \in E(T)} c_G(e)$. In fact, this optimization problem has also been studied in many communities, including operations research and theoretical computer science.

The decision version of STP is one of Karp's 21 NP-Complete problems. In recent decades, a rich number of theoretical studies have been devoted to this problem. For instance, it is known that STP can be solved in time $O(2^t n^2 + nm)$ [3] where n , m and t denote the number of vertices, edges, and terminals, respectively. The current best approximation ratio is 1.39 [5], and there is no 1.01 approximation algorithm (unless $P = NP$) [9]. In practice, except in the area of network design, STP is also modeled in many other fields like computational biology [8], VLSI design [13], and social network analysis [18]. Notably, two international competitions, the 11th DIMACS Implementation Challenge in 2014 [15] and the 3rd Parameterized Algorithms and Computational Experiments (PACE) Challenge in 2018 [4] have been dedicated to solving STP practically, which demonstrates the interest of the research community for this challenging problem and its variants such as the directed Steiner tree problem, the prize-collecting Steiner tree problem, etc.

Due to the great importance of STP, there exists a considerable number of both exact and heuristic solution algorithms for the problem. Exact algorithms typically rely on two search frameworks, dynamic programming and integer linear programming. To our knowledge, exact algorithms can solve problem instances with up to thousands of vertices or hundreds of terminals in a reasonable time, mostly thanks to some new ILP techniques like vertex formulation and local branching [6, 10]. However, to deal with larger instances whose optimality cannot be determined by exact algorithms, heuristic algorithms become indispensable [10, 11, 20–22]. For example, for some VLSI graphs with up to ten thousand vertices, no exact algorithms can produce a feasible solution. According to the final report of the 11th DIMACS Challenge, PUW [20] and Staynerd [10] are considered the best global heuristics compared to other competing algorithms (*AB* [1], *mozartballs* [10], *polito* [2], and *scipjack* [12]) in the heuristic track. Interestingly, the exact algorithm *mozartballs*, which also uses some search techniques of Staynerd, is the winner in the exact track during the 11th DIMACS Challenge. As for the more recent PACE Competition, an evolutionary algorithm developed by the CIMAT Team ranked first in the heuristic track [4]. Considering that PUW and Staynerd were absent from PACE, we consider that PUW, Staynerd and CIMAT are among the most competitive heuristics for STP so far. Notably, all these heuristic algorithms use local search as their key optimization component.

Today, large STP instances with at least several thousand vertices appear ubiquitously in many applications. For example, the well-known GEO instances have more than 100,000 vertices or up to 5,000 terminals. We observed that existing heuristic algorithms, including the aforementioned DIMACS and PACE winners, attain their limit when they are applied to such large instances, leading to large gaps to the best lower bounds. In this work, our aim is to propose a

novel heuristic algorithm that is able to better solve large-scale STP instances. The contributions can be summarized as follows.

- We design a partition-and-merge (PM) algorithm to deal with the large-scale nature of the given instance. PM first breaks the input graph into a number of smaller subgraphs. Then it finds partial solutions for these subgraphs and pairwise merges the subgraphs and solutions at the same time. The algorithm provides a high-quality complete solution when all the subgraphs are merged into the input graph. Meanwhile, we investigate different merging heuristics.
- We conduct extensive experiments, with a focus on very large instances, to evaluate the PM algorithm and different techniques. We show that our algorithm is very competitive in solving large and hard instances. Remarkably, for the largest publicly available STP instances, the EFST family, PM outperforms the known heuristic solvers including the winners of the 11th DIMACS and 3rd PACE competitions.

2 Basic Notations

A STP instance is composed of an edge-weighted undirected graph G and a subset of terminal vertex A . We use $V(G)$ to denote the vertex set, $E(G)$ to denote the edge set, and $c_G : E(G) \rightarrow \mathbb{R}^+$ to denote the edge-weight function of G . So $A \subseteq V(G)$. When the context is clear, we use n , m and t to denote the number of vertices, edges, and terminals, respectively. For two arbitrary vertices $u, v \in V(G)$, the *distance* $d_G(u, v)$ between them is the smallest total weight among the paths between u and v in G . For a vertex subset $S \subseteq V(G)$, $G[S]$ denotes the subgraph induced by S .

A *partition* of the graph G , $\mathcal{P} = \{H_1, \dots, H_{|\mathcal{P}|}\}$, is a set of vertex-induced subgraphs of G such that $V(H_i) \cap V(H_j) = \emptyset$ for any $1 \leq i < j \leq |\mathcal{P}|$ and $\bigcup_{1 \leq i \leq |\mathcal{P}|} V(H_i) = V(G)$. $(u, v) \in E(G)$ is a *cutting edge* of \mathcal{P} if u and v are in different subgraphs of \mathcal{P} .

Give an input instance with graph G and terminal set $A \subseteq V(G)$, a *Voronoi diagram* $\mathcal{D} = \{G_1, \dots, G_t\}$ is a special partition of G such that each subgraph G_i contains exactly one terminal, denoted by a_i , and for all other vertex $u \in V(G_i) \setminus \{a_i\}$, a_i is the nearest terminal of u in G , that is, $d_G(u, a_i) = \min_{a \in A} d_G(u, a)$. Clearly, the number of subgraphs in a Voronoi diagram is exactly t . For an undirected graph $G = (V, E)$, the Voronoi diagram can be computed by Dijkstra's shortest path algorithm in $O(m + n \log n)$ time (with ties breaking randomly) [22].

3 The Partition-and-Merge Algorithm

Our Partition-and-Merge (PM) algorithm is described in Alg. 1. The input of PM consists of an input graph $G = (V, E)$, a terminal set $A \subseteq V$, two parameters $\rho \leq |A|$ and $\gamma > 1$, and two maximum iteration numbers I_{short} and I_{long} .

PM maintains a partition of G , that is, $\mathcal{P} = \{H_1, \dots, H_{|\mathcal{P}|}\}$. As PM runs, the subgraphs in \mathcal{P} are merged in pairwise. The *merge operation* concatenates H_i and H_j into a larger graph H' by adding the cutting edges between H_i and H_j . In H' , the edge weights are the same as their weights in $E(H_i)$ and $E(H_j)$ from H_i and H_j , respectively, but the weights of the cutting edges between H_i and H_j are the same as in the original input graph G . It is important to notice that the weight of the edge e in $E(H_i)$ is possibly different from its weight in $E(G)$. The reason is that PM modifies the weights of the edges of the subgraphs in \mathcal{P} as the search runs.

Algorithm 1: The Partition-and-Merge Algorithm

Input: Graph G , terminal set $A \subseteq V(G)$, parameters $\rho \in [0, |A|]$, $\gamma \in [1, \infty)$, the maximum iteration numbers I_{short} and I_{long} for local search.

Output: Steiner tree T

```

1 Compute a Voronoi diagram  $\mathcal{P} = \{H_1, \dots, H_{|\mathcal{P}|}\}$  with respect to  $G$  and  $A$ .
2 Initially,  $level \leftarrow 1$ 
3 while  $|\mathcal{P}| > \rho$  do
4    $\mathcal{P} \leftarrow \text{OneRoundMerge}(G, \mathcal{P}, \lceil \frac{n\gamma^{level}}{t} \rceil)$ 
5    $level \leftarrow level + 1$ 
6 while  $|\mathcal{P}| > 1$  do
7   for each  $H_i \in \mathcal{P}$  do
8      $T_i \leftarrow \text{LocalOptimize}(G[V(H_i)], H_i, A \cap V(H_i), I_{short})$ 
9      $H'_i \leftarrow \text{SolutionStoring}(H_i, T_i)$ 
10     $\mathcal{P} \leftarrow \mathcal{P} \setminus \{H_i\} \cup \{H'_i\}$ 
11    $\mathcal{P} \leftarrow \text{OneRoundMerge}(G, \mathcal{P}, \lceil \frac{n\gamma^{level}}{t} \rceil)$ .
12    $level \leftarrow level + 1$ 
13 There is only one graph  $H_1$  in  $\mathcal{P}$ .
14  $T \leftarrow \text{LocalOptimize}(G, H_1, A, I_{long})$ 
15 return  $T$ 

```

3.1 Two-Phase Merging Procedure

Initially, as shown in line 1 of Alg. 1, PM builds \mathcal{P} by computing the Voronoi diagram, so each subgraph in \mathcal{P} contains one terminal vertex. The algorithm then performs a two-phase merging procedure, which gradually merges the subgraphs of \mathcal{P} into a whole graph. The parameter ρ is used to set the time to start the second phase and γ controls the size gap between different subgraphs in \mathcal{P} .

In the first phase (lines 3–5), *OneRoundMerge* is used to pairwise merge the graphs in \mathcal{P} . In general, *OneRoundMerge* selects two subgraphs by a specific heuristic rule and joins them to form a larger graph. We introduce different merging rules in Sect. 3.2. *OneRoundMerge* also makes sure that the subgraph after a round of merge has no more than $\lceil \frac{n\gamma^{level}}{t} \rceil$ vertices (so that the subgraphs

in \mathcal{P} have roughly the same scale). Clearly, $\lceil \frac{n\gamma^{level}}{t} \rceil$ increases geometrically with a factor of γ as *level* increases. For example, if $\gamma = 2$, this value doubles at every level.

The second phase, as in lines 6–12, begins when there are no more than ρ subgraphs in \mathcal{P} . In this phase, for each subgraph H_i , we additionally carry out a local search algorithm, *LocalOptimize*, to find a (high-quality) Steiner tree in $G[V(H_i)]$. It is worth noting that *LocalOptimize* finds a starting solution from H_i , but improves the solution in $G[V(H_i)]$. Assume T_i is a Steiner tree obtained by *LocalOptimize*($G[V(H_i)], H_i, A \cap V(H_i), I_{short}$). T_i spans terminals $A \cap V(H_i)$ in both H_i and $G[V(H_i)]$. Our intuition is to keep the information that T_i is a high-quality solution of graph H_i , so that the further search will benefit from this information. This is implemented by *SolutionStoring* method, which essentially changes the edge weights of H_i . We will introduce two different ways for *SolutionStoring* in Sect. 3.3. After local optimization and solution storing for each subgraph, we continue using *OneRoundMerge* with the same heuristic in the first phase to merge the subgraphs.

Example. We consider the example in Fig. 1, which describes the partition-and-merge procedure. We set $\gamma = 2$ and $\rho = 6$. The size of each subgraph approximately doubles after each level of merging. The Steiner tree solution is computed when there are no more than 6 subgraphs.

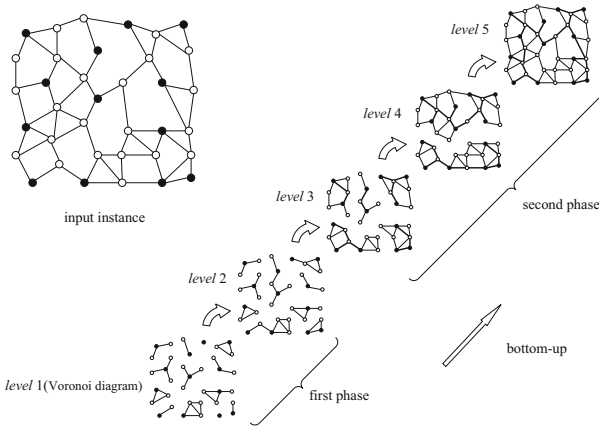


Fig. 1. The bottom-up partition-and-merge procedure. Terminal vertices are marked by black dots. The sub-solutions computed for subgraphs are marked by bold lines.

3.2 Merging Heuristics

As mentioned, the *OneRoundMerge* procedure repeatedly selects a pair of subgraphs (H_i, H_j) from \mathcal{P} and merges them into one subgraph. In Alg. 2, we show the general procedure of *OneRoundMerge* using the so-called *random rule*. We will also discuss the *random-edge* and *min-distance* rules afterward.

Algorithm 2: One round of merging by random heuristic

Input: Graph G , a partition \mathcal{P} , subgraph size limit ω

Output: A new partition of G

```

1 while true do
    Random rule:
2   From all pairs  $(H_k, H_l) \in \mathcal{P} \times \mathcal{P}$  such that  $|V(H_k)| + |V(H_l)| \leq \omega$  and
       $cut_{\mathcal{P}}(H_k, H_l) \neq \emptyset$ , randomly pick one  $(H_i, H_j)$ .
3   Merge  $H_i$  and  $H_j$  to  $H'$ .
4    $\mathcal{P} \leftarrow \mathcal{P} \setminus \{H_i, H_j\} \cup \{H'\}$ 
5   if none of the pairs of graphs in  $\mathcal{P}$  meet the random heuristic then
6     break
7 return  $\mathcal{P}$ 

```

Random Rule. By this rule, we randomly select two subgraphs that are *adjacent*, that is, there is at least one cutting edge between them, and the total number of vertices in the two subgraphs does not exceed ω . Then we merge the two subgraphs and repeat the procedure until no such pair of graphs exists. The intuition behind this heuristic is to keep each subgraph in \mathcal{P} connected and balanced in terms of vertex number after merging.

Random-Edge Rule. The random-edge rule, as shown in below, resembles the random rule except that it picks a cutting edge randomly rather than a pair of adjacent subgraphs randomly. Clearly, this rule concerns the number of cutting edges between every pair of adjacent subgraphs. Using this rule to replace line 3 in Alg. 2, we can obtain another version of *OneRoundMerge* with the random-edge heuristic.

Random-edge rule:
 From all cutting edges (x, y) such that $x \in V(H_k)$, $y \in V(H_l)$ and $|V(H_k)| + |V(H_l)| \leq \omega$, randomly pick one edge (u, v) .
 Suppose $u \in V(H_i)$ and $v \in V(H_j)$ without loss of generality.

We note that this rule is inspired by Karger’s minimum-cut randomized algorithm [16]. As we know, Karger’s algorithm finds the minimum cut of a graph with high probability by repeatedly merging a randomly selected edge. Thus, it is believed that this heuristic gives preference to the pair of adjacent subgraphs that are *closely connected*, that is, there are more cutting edges between them.

Min-Distance Rule. Let $A_i = A \cap V(H_i)$ and $A_j = A \cap V(H_j)$ where A is the input terminal set, H_i and H_j are two subgraphs in \mathcal{P} . For each pair of H_i and H_j , we compute $d_G(A_i, A_j)$, the minimum value over all the shortest lengths of path between any two terminals in $A_i \times A_j$, i.e., $d_G(A_i, A_j) = \min_{(u,v) \in A_i \times A_j} d_G(u, v)$. Then, the min-distance is shown as follows. Using this rule to replace line 3 in Alg. 2, we obtain *OneRoundMerge* with the min-distance heuristic.

Min-distance rule:

Find all pairs $(H_k, H_l) \in \mathcal{P} \times \mathcal{P}$ such that $|V(H_k)| + |V(H_l)| \leq \omega$ and $cut_{\mathcal{P}}(H_k, H_l) \neq \emptyset$, pick a pair (H_i, H_j) with smallest $d_G(A_i, A_j)$, ties breaking randomly

The running time of the min-distance rule bounded by $O(|E|)$ due to the early computation of the Voronoi diagram. For each non-terminal vertex $u \in H_i \subseteq A_i$, we already know its nearest terminal $a_u = argmin_{a \in A_i} d_G(a, u)$ after computing the Voronoi diagram. Thus, for each cutting edge $(u, v) \in E_{\mathcal{P}}$, we first obtain $d(u, v) = d_G(u, a_u) + c_G(u, v) + d_G(v, a_v)$ where a_u and a_v are the nearest terminals of u and v , respectively. Then, we randomly pick up an (u, v) of the smallest values $d(u, v)$ and merge its corresponding subgraphs H_i and H_j . Because we only iterate over the edges in $E_{\mathcal{P}}$, the above running time follows.




Fig. 2. Example for three merging rules. $G_1 - G_4$ are four subgraphs. Lines connecting the subgraphs represent the cutting edges between subgraphs. The minimum distance between the terminals of each two subgraphs is shown in the right-side of the figure.

Example. In Fig. 2, there are currently four subgraphs G_1, \dots, G_4 and any two subgraphs have fewer than ω vertices. Now, we need to select a pair of adjacent subgraphs to merge. If the random rule is used, the probability that each pair of adjacent subgraphs is selected is the same. If the random-edge rule is used, the probability of selecting subgraphs (G_1, G_3) is the highest, followed by (G_2, G_4) , (G_1, G_2) and (G_3, G_4) . If the min-distance rule is used, (G_3, G_4) is definitely selected for merging.

3.3 Solution Storing

The motivation of solution storing is to guide future search by the fact that T_i is a high-quality solution in H_i . To achieve this, we simply change the edge weight of H_i so that we can apply the same local search algorithm. Here, we introduce a random way to reweight the edges in H_i .

For any edge $e \in E(H_i)$, we score the edge weight $c_{H_i}(e)$ as follows, where α is a random real number from $[0, 0.75]$.

$$c_{H_i}(e) = \begin{cases} (1 - \alpha)c_G(e) & e \in E(T_i) \\ (1 + \alpha)c_G(e) & \text{otherwise} \end{cases}$$

Clearly, the weights of edges in $E(T_i)$ are reduced, while the weights of edges not in $E(T_i)$ are increased, implying that the edges of $E(T_i)$ are preferred.

3.4 Local Optimization

The local optimization procedure, $\text{LocalOptimize}(G, H, A, I_{max})$ is a population-based search algorithm modified from the state-of-the-art PUW solver [20]. The original graph G , the edge reweighted graph H , terminal set A , maximum number of iterations I_{max} .

In summary, LocalOptimize follows the scheme of genetic programming. It maintains a pool of elite solutions \mathcal{Q} . In each generation, it produces new solutions to update the elite pool – It first builds a new Steiner tree T from H (Note that H is a graph with biased edge weights by solution storing). Then, T is improved by local search, but the local search runs on the initial graph G without biased edge weights. Afterward, the improved solution T exchanges information with another solution from the pool by combination, resulting in a new solution T' . Finally, T and T' are used to update the pool.

For detailed description of the local optimization algorithm, please refer to the full version.

4 Experiments

4.1 Settings

Our algorithm is programmed in C++ and compiled in g++ with the optimization option '-O3'.¹ All experiments were conducted on a Linux server with Intel Xeon Gold 6130 (2.10 GHz) processor.

Parameters Setting. In the experiments, ρ is set to 64 and γ is set to 2. For all test instances, we set the total time limit $t_{total} = 3600$ seconds (1h). The two parameters for controlling the number of iterations in LocalOptimize , I_{short} and I_{long} , are set to 256 and 2048, respectively. However, with these configurations, it is possible that we may not obtain a complete solution after 3600s for some graphs. This happens when there are still more than one subgraph in \mathcal{P} after this time frame. Thus, we additionally require that each LocalOptimize runs at most $\frac{0.7(\gamma-1)}{\gamma(\rho-1)} \times t_{total}$ seconds for each subgraph. This implies that at least $0.3 \times t_{total}$ seconds are allocated for running LocalOptimize with the final graph, that is, line 14 in Alg. 1.

4.2 Performance Assessment

Results on the Set of 113 Large Instances We first test the algorithm in 113 Large Instances². They are grouped in 8 datasets shown in Table 1, which indicates

¹ Source codes are publicly available at <https://github.com/xyu03/PM.git>.

² All instances are downloaded from site <http://dimacs11.zib.de/downloads.html>.

the number of instances, average vertex number, average edge number, average terminal number, and source of the dataset. GEO, I and EFST contain larger instances (with more vertices, edges, and terminals) compared to ES, VLSI and TSPFST.

Table 1. Information about the 8 large datasets

Dataset	Num	ave. <i>n</i>	ave. <i>m</i>	ave. <i>t</i>	Source
ES	16	4245	6235	1563	GeoSteiner generator
TSPFST	10	4687	7043	1871	TSPLIB
VLSI	10	17926	30973	452	VLSI applications
GEO	22	105906	150767	1845	Telecommunication networks
I	10	65120	106450	3211	Telecommunication networks
EFST(R25K)	15	39483	47893	25000	GeoSteiner generator
EFST(R50K)	15	79128	96293	50000	GeoSteiner generator
EFST(R100K)	15	158147	192190	100000	GeoSteiner generator

For this set of instances, we adopt the following best STP heuristic algorithms as our reference algorithms.

- PUW [20], a multi-start local search algorithm which ranks first at Formula 1 in the single-thread heuristic track of the 11th DIMACS Competition.
- Staynerd [10], a winner of most category of instances in the single-thread heuristic track of the 11th DIMACS Competition.
- CIMAT [4], the champion of the heuristic track of the PACE Challenge (2018).
- MPCH [19], a partition-based algorithm focusing on solving large graphs.

For each instance, we perform each of the above algorithms (except MPCH³) 10 times independently. For each algorithm and each instance, we take the best solution obtained in the 10 trials. We summarize these results for each dataset in Table 2. Detailed results of these 113 instances are provided in the full version of the paper. For each dataset and each algorithm, we indicate the average solution gaps of the algorithm⁴ and the number of instances for which the algorithm finds the best solution. In terms of solution gap, our PM algorithm performs better than other algorithms on the largest instances of the GEO, I, and EFST family datasets. It is also remarkable that PM is the best on all I and EFST instances and 16/22 on the GEO instances. Meanwhile, for the relatively smaller instances in this group, the algorithms compete with each other.

³ Because the code of MPCH is unavailable, we extract its results from [19].

⁴ For an instance, the gap of a solution T is $\frac{\sum_{e \in E(T)} c_G(e) - lb}{lb} \times 100\%$, lb being the best-known lower bound for the instance.

Table 2. The best solutions found by each algorithms on the set of 113 large instances. The smallest average gaps are marked in bold.

Dataset	ave.gap(%)					#best				
	MPCH	PUW	Staynerd	CIMAT	PM	MPCH	PUW	Staynerd	CIMAT	PM
ES	0.84	0.0082	0.0069	0.0987	0.0078	–	0	3	0	1
TSPFST	0.96	0.0100	0.0203	0.1068	0.0380	–	2	2	0	0
VLSI	1.64	0.0292	0.3963	0.0665	0.0730	–	2	0	1	1
GEO	–	0.8136	1.0044	1.0753	0.7744	–	4	0	0	16
I	0.02	0.0130	0.0154	0.0394	0.0107	–	0	0	0	10
EFST(R25K)	–	0.0798	0.0347	0.2761	0.0101	–	0	0	0	15
EFST(R50K)	–	0.1131	–	0.3387	0.0188	–	0	0	0	15
EFST(R100K)	–	0.1344	–	0.3644	0.0353	–	0	0	0	15

Results on the Set of 30 DIMACS Instances. We first test 30 instances that are used in the final ranking phase for the single thread heuristic challenge of the 11th DIMACS Implementation Challenge⁵. They are representative instances from the well-known SteinLib, Vienna, Copenhagen, PUCN, and GAPS libraries.

For this set of instances, we compare our results with those of six best performing reference algorithms: AB [1], mozartballs [10], politico [2], PUW [20], scipjack [12] and Staynerd [10], which participated in the final heuristic track of the 11th DIMACS Challenge. Following the time constraints in the competition results of the DIMACS challenge heuristic track (1 thread), all algorithms are run under a cutoff time of 2 h.

According to Table 3, in terms of finding the best solutions, politico and PM each have 6 solutions that outperform the remaining algorithms. Mozartballs performs best on 'Average' and 'ave. gap(%)'. In particular, a detailed data analysis reveals that PM has the best solution for fnl4461fst-p, alue7080-p, es10000fst01-p, cc3-12p-p, G106ac-p, I064ac-p, which are the largest instances in the group.

Table 3. The best solutions found by each algorithms on DIMACS competition instances

Instance	AB	mozartballs	politico	PUW	scipjack	staynerd	PM
ave.gap(%)	776.0807	0.3640	3.8920	0.4174	1.5015	0.3737	0.6219
#best	0	0	6	4	0	1	6

The Effectiveness of Different Merging Rules. We show the average solution gaps obtained by using each rule in Table 4. We use the Fredman test to check the statistical significance of the differences in terms of solution quality, which

⁵ <https://dimacs11.zib.de/contest/results/SPGp.1.7200.heur.html>.

reveals a p -value of 0.00. The p -value (< 0.05) clearly indicates the differences between the variants with the three different merging rules. The random-edge and min-distance strategies perform better than the random strategy, while min-distance shows a slight advantage over random-edge.

Table 4. The average gaps (%) of PM variants using different merging rules.

Dataset	random	random-edge	min-distance
ES	0.02069	0.00922	0.00783
TSPFST	0.01832	0.01831	0.03799
VLSI	0.09954	0.08684	0.07299
GEO	0.80160	0.78055	0.77435
I	0.01144	0.01090	0.01074
EFST(R25K)	0.01423	0.00960	0.01008
EFST(R50K)	0.02740	0.01868	0.01877
EFST(R100K)	0.04900	0.03646	0.03533

5 Conclusion

In this paper, we introduced a new partition-and-merge algorithm to effectively solve large-scale STP instances. To ensure the effectiveness of the algorithm, we investigated different merging heuristics, local search algorithms, and new solution storing techniques.

This is the first STP algorithm based on a multi-level merging and search strategy. The algorithm has some attractive features for parallel computing. For example, we can run local optimizing procedures on different threads in parallel as they are independent.

Acknowledgements. We thank the authors of Staynerd, PUW, CIMAT and other solvers in the paper for providing their codes. We also thank the organizers of the 11th DIMACS challenge and the 3rd PACE competition for hosting the data. This work is partially funded by the National Natural Science Foundation of China under grant No. 62372093 and the Natural Science Foundation of Sichuan Province of China under grant No. 2023NSFSC1415.

References

1. Althaus, E., Blumenstock, M.: Algorithms for the maximum weight connected subgraph and prize-collecting Steiner tree problems (2014). <http://dimacs11.cs.princeton.edu/workshop.html>
2. Biazzo, I., Braunstein, A., Zecchina, R.: Performance of a cavity-method-based algorithm for the prize-collecting Steiner tree problem on graphs. Phys. Rev. E **86**(2), 026706 (2012)

3. Björklund, A., Husfeldt, T., Kaski, P., Koivisto, M.: Fourier meets möbius: fast subset convolution. In: Proceedings of the 39th ACM Symposium on Theory of Computing, pp. 67–74 (2007)
4. Bonnet, É., Sikora, F.: The pace 2018 parameterized algorithms and computational experiments challenge: the third iteration. In: IPEC 2018 (2018)
5. Byrka, J., Grandoni, F., Rothvoß, T., Sanità, L.: Steiner tree approximation via iterative randomized rounding. *J. ACM* **60**(1), 1–33 (2013)
6. Carrabs, F., Cerulli, R., Pentangelo, R., Raiconi, A.: Minimum spanning tree with conflicting edge pairs: a branch-and-cut approach. *Ann. Oper. Res.* **298**(1), 65–78 (2021)
7. Cheng, X., Du, D.Z.: Steiner Trees in Industry, vol. 11. Springer (2013)
8. Cheng, X., Li, Y., Du, D.Z., Ngo, H.Q.: Steiner trees in industry. In: Handbook of Combinatorial Optimization, pp. 193–216. Springer (2004)
9. Chlebík, M., Chlebíková, J.: The Steiner tree problem on graphs: Inapproximability results. *Theoret. Comput. Sci.* **406**(3), 207–214 (2008)
10. Fischetti, M., et al.: Thinning out Steiner trees: a node-based model for uniform edge costs. *Math. Program. Comput.* **9**(2), 203–229 (2017)
11. Fu, Z.H., Hao, J.K.: Swap-vertex based neighborhood for Steiner tree problems. *Math. Program. Comput.* **9**(2), 297–320 (2017)
12. Gamrath, G., Koch, T., Maher, S.J., Rehfeldt, D., Shinano, Y.: SCIP-jack-a solver for STP and variants with parallelization extensions. *Math. Program. Comput.* **9**(2), 231–296 (2017)
13. Held, S., Korte, B., Rautenbach, D., Vygen, J.: Combinatorial optimization in VLSI design. *Combin. Optim.-Meth. Appl.* **31**, 33–96 (2011)
14. Hsieh, C.Y., Liao, W.: All-optical multicast routing in sparse splitting WDM networks. *IEEE J. Sel. Areas Commun.* **25**(6), 51–62 (2007)
15. Johnson, D.S., Koch, T., Werneck, R.F., Zachariasen, M.: 11th Dimacs implementation challenge in collaboration with Icerm: steiner tree problems (2014)
16. Karger, D.R.: Global min-cuts in RNC, and other ramifications of a simple min-cut algorithm. In: SODA, vol. 93, pp. 21–30. Citeseer (1993)
17. Kun, Z., Yong, Q., Hong, Z.: Dynamic multicast routing algorithm for delay and delay variation-bounded Steiner tree problem. *Knowl.-Based Syst.* **19**(7), 554–564 (2006)
18. Lappas, T., Liu, K., Terzi, E.: Finding a team of experts in social networks. In: Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 467–476 (2009)
19. Luiipersbeck, M.: A new partition-based heuristic for the Steiner tree problem in large graphs. Ph.D. thesis, Technische Universität Wien (2013)
20. Pajor, T., Uchoa, E., Werneck, R.F.: A robust and scalable algorithm for the Steiner problem in graphs. *Math. Program. Comput.* **10**(1), 69–118 (2018)
21. Ribeiro, C.C., Uchoa, E., Werneck, R.F.: A hybrid grasp with perturbations for the Steiner problem in graphs. *INFORMS J. Comput.* **14**(3), 228–246 (2002)
22. Uchoa, E., Werneck, R.F.: Fast local search for the Steiner problem in graphs. *J. Exp. Algorithmics (JEA)* **17**, 1–2 (2012)