

Open Journal of Mathematical Optimization

Giovanni Righini

Efficient optimization of the Held–Karp lower bound

Volume 2 (2021), article no. 9 (17 pages)

<https://doi.org/10.5802/ojmo.11>

Article submitted on April 27, 2020, revised on August 12, 2021,
accepted on October 22, 2021.



This article is licensed under the
CREATIVE COMMONS ATTRIBUTION 4.0 INTERNATIONAL LICENSE.
<http://creativecommons.org/licenses/by/4.0/>



Efficient optimization of the Held–Karp lower bound

Giovanni Righini

University of Milan, Department of Computer Science
via Celoria 18, Milano
Italy
giovanni.righini@unimi.it

Abstract

Given a weighted undirected graph $G = (V, E)$, the Held–Karp lower bound for the Traveling Salesman Problem (TSP) is obtained by selecting an arbitrary vertex $\bar{p} \in V$, by computing a minimum cost tree spanning $V \setminus \{\bar{p}\}$ and adding two minimum cost edges adjacent to \bar{p} . In general, different selections of vertex \bar{p} provide different lower bounds. In this paper it is shown that the selection of vertex \bar{p} can be optimized, to obtain the largest possible Held–Karp lower bound, with the same worst-case computational time complexity required to compute a single minimum spanning tree. Although motivated by the optimization of the Held–Karp lower bound for the TSP, the algorithm solves a more general problem, allowing for the efficient pre-computation of alternative minimum spanning trees in weighted graphs where any vertex can be deleted.

Digital Object Identifier 10.5802/ojmo.11

Keywords Traveling salesman problem, Minimum spanning tree, Held–Karp lower bound, Union-Find data-structure.

1 Introduction and motivation

The Traveling Salesman Problem (TSP) is the well-known problem of computing a minimum cost Hamiltonian cycle on a given weighted graph [1, 5]. Since the problem is *NP*-hard, many techniques have been devised not only to solve it to optimality but also to compute upper and lower bounds to its optimal value. One of the most celebrated lower bounding techniques was proposed in 1970 by Held and Karp [6]; it is known to provide very tight lower bounds with limited computational effort and it is used as a sub-routine within state-of-the-art TSP solvers, like Concorde [1], and state-of-the-art TSP heuristics, like the Helsgaun implementation of the Lin–Kernighan algorithm [7]. The Held–Karp lower bound is still subject to investigation: besides being included in branch-and-bound algorithms for the symmetric TSP [9], its properties have been studied for the asymmetric TSP [10] and for the metric TSP [8]. Recent contributions include, for instance, hybridization with constraint programming [2] and randomized approximation algorithms [3].

The lower bounding procedure proposed by Held and Karp is based on Lagrangean relaxation: at each iteration the edge weights of the underlying graph $G = (V, E)$ are modified so that the costs of Hamiltonian cycles remain unaffected. Then, valid lower bounds for the TSP can be efficiently obtained from minimum spanning 1-trees, i.e. spanning trees with an additional edge, relying on the observation that a Hamiltonian cycle is a special case of a spanning 1-tree. This can be done in different ways: here below we list three of them (see Figure 1).

Lower bound LB_0 :

- a minimum spanning tree T^* of G ;
- a minimum cost edge not in T^* .

Lower bound $LB_H(\bar{\ell})$ (Helsgaun [7]):


- a minimum spanning tree T^* of G ;
- a minimum cost edge $[\bar{\ell}, j] \notin T^*$, incident to a leaf $\bar{\ell}$ of T^* .

Lower bound $LB_{HK}(\bar{p})$ (Held and Karp [6]):

- a minimum spanning tree $T_{\bar{p}}^*$ of the subgraph induced by $V \setminus \bar{p}$, where \bar{p} is an arbitrary selected vertex;
- two minimum cost edges of G with an endpoint in \bar{p} .



© Giovanni Righini;
licensed under Creative Commons License Attribution 4.0 International



■ **Figure 1** Left (LB_0): a minimum spanning 1-tree, made by a minimum spanning tree and a minimum cost edge (dashed). Center (LB_H): a Heldsgaun 1-tree, made by a minimum spanning tree and a minimum cost edge (dashed) incident to a leaf (filled). Right (LB_{HK}): a Held–Karp 1-tree, made by a minimum spanning tree disregarding a vertex (filled) and two minimum cost edges incident to it (dashed). The comparison illustrates the difference between the three corresponding lower bounds.

It is immediate to observe that $LB_H(\bar{\ell}) \geq LB_0$ for any choice of a leaf $\bar{\ell} \in T^*$ and $LB_{HK}(\bar{p}) \geq LB_0$ for any choice of $\bar{p} \in V$, because $LB_H(\bar{\ell})$ and $LB_{HK}(\bar{p})$ are the costs of 1-trees of G , while LB_0 is the cost of a minimum 1-tree of G .

Both $LB_H(\bar{\ell})$ and $LB_{HK}(\bar{p})$ depend on the selection of a vertex: then, both can be maximized by selecting $\bar{\ell}$ and \bar{p} in an optimal way. Let $LB_H^* = \max_{\bar{\ell} \in \Lambda(T^*)} \{LB_H(\bar{\ell})\}$, where $\Lambda(T^*)$ indicates the set of leaves of T^* , and $LB_{HK}^* = \max_{\bar{p} \in V} \{LB_{HK}(\bar{p})\}$. Since $LB_{HK}(\bar{p}) = LB_H(\bar{\ell})$ when $\bar{p} = \bar{\ell}$ and since $V \supset \Lambda(T^*)$, then it is guaranteed that $LB_H^* \leq LB_{HK}^*$.

Although LB_{HK}^* is more appealing than LB_H^* for the tightness of the TSP lower bound that can be achieved, the computation of LB_H^* is much easier: it just requires to select the leaf of T^* for which the second smallest cost of the incident edges is maximum. For this reason the Heldsgaun lower bound is attractive for practical implementation purposes.

The main motivation of this paper is the efficient optimization of $LB_{HK}(\bar{p})$, which is less trivial: a naive procedure is to compute as many *alternative minimum spanning trees* as the number n of vertices of G , each one spanning G with the exception of a vertex $\bar{p} \in V$. This paper presents and analyzes an algorithm that allows to compute all the *alternative edges* that are needed to rebuild a minimum cost spanning tree when a vertex is deleted from G . This goal is achieved with the same worst-case computational complexity required by the computation of a single minimum spanning tree of G . As a by-product, the algorithm allows to efficiently compute LB_{HK}^* .

The algorithm takes in input a weighted graph, a minimum spanning tree and a sorted list of all edges and it returns an optimal set of alternative edges. In particular, it computes a vertex and a corresponding 1-tree that provide the largest Held–Karp lower bound LB_{HK}^* .

Paper outline

The next section gives an overview of the main idea behind the algorithm, with the aim of facilitating the reader. The full details are presented in Sections 3, 4 and 5. Section 3 illustrates the main data-structures, their properties, their initialization and some procedures for updating them. Section 4 describes the main algorithm (Algorithm 7). Section 5 describes how the maximum value LB_{HK}^* of the Held–Karp lower bound is obtained.

2 Overview of the algorithm

Consider an undirected graph $G = (V, E)$ with edge costs $c : E \mapsto \mathfrak{R}$; the number of vertices is n and the number of edges is m . A minimum spanning tree T^* of G is assumed to be known and the edges of E are assumed to having been sorted by non-decreasing cost in a previous execution of Kruskal algorithm.


When a vertex $p \in V$ is deleted, T^* is disconnected into a forest F_p^* made of as many connected components as the neighbors of p in T^* . In the remainder they are called *p-components*, to mean that they arise when p is deleted. When vertex $p \in V$ is deleted, *alternative edges* must be found to reconnect the *p-components* and they must be selected to produce a new minimum spanning tree T_p^* .

A trivial way to accomplish this task is to repeat, for each deleted vertex $p \in V$, the computation of a minimum spanning tree T_p^* with Kruskal or Boruvka algorithm after having restored the state of a suitable data-structure

to represent the p -components of F_p^* . The vertices in each p -component can be identified in $O(n)$ and this is the complexity to set up the state of a Union-Find data-structure (for readers unfamiliar with Union-Find, it is illustrated in detail in Appendix A). Then, the computation of an alternative minimum spanning tree with Kruskal algorithm requires $O(m + n \log d(p))$, where $d(p)$ indicates the degree of vertex $p \in V$. Hence, repeating this naive procedure for all vertices would require $O(mn + n^2 \log n)$.

To achieve a better worst-case time complexity one could work the other way around: all edges are scanned in non-decreasing order of their cost, as in Kruskal algorithm, and for each edge $[i, j] \notin T^*$ the algorithm searches the vertices p such that the edge $[i, j]$ reconnects two different p -components of F_p^* . For each vertex p and for each pair of p -components in F_p^* , the first edge that is found to reconnect them, certainly belongs to T_p^* owing to the edge ordering. In this way all forests F_p^* can be populated in parallel with alternative edges for all $p \in V$ until all alternative minimum spanning trees T_p^* are produced.

This search, if performed as described above, would be very inefficient. However, it can be restricted (and accelerated) as follows. Consider the minimum spanning tree T^* and an edge $[i, j] \notin T^*$. Adding edge $[i, j]$ to T^* defines a unique cycle $C(i, j)$. The vertices along $C(i, j)$, except i and j , are those for which $[i, j]$ is a candidate alternative edge: if a vertex $p \in V$ along $C(i, j)$ is deleted from T^* , edge $[i, j]$ reconnects two different p -components. On the contrary, for all vertices $p \in V$ not belonging to $C(i, j)$, edge $[i, j]$ cannot be a candidate alternative edge, because both i and j belong to the same p -component. To efficiently scan $C(i, j)$, the algorithm illustrated in the remainder skips in constant time all vertices p along $C(i, j)$ for which an alternative edge reconnecting the p -components of i and j has already been found. The effect is illustrated in Figure 2. When edge $[i', j']$ is considered, it can be an alternative edge for vertices a, b, c, d and e . Later on, when the more expensive edge $[i'', j'']$ is considered, it can be an alternative edge for vertices f, g, h and l , but there is no point in considering it for a, b, c, d and e again.



■ **Figure 2** Cycles $C(i', j')$ and $C(i'', j'')$ partially overlap: since $[i', j']$ is cheaper than $[i'', j'']$, the vertices belonging to $C(i', j')$ can be skipped when $[i'', j'']$ is considered.

This idea allows to achieve an overall computational complexity $O(m + n \log n)$, since all operations executed when a candidate alternative edge is discarded give a contribution $O(1)$ for each edge, i.e. $O(m)$ overall, while all operations executed when a candidate alternative edge is inserted in one of the minimum alternative spanning trees give a contribution $O(\log n)$ for each insertion, i.e. $O(n \log n)$ overall.

3 The data-structures

This section describes the main data-structures used in the algorithm, their properties and their initialization. Some necessary terminology and definitions are also introduced.

3.1 The oriented minimum spanning tree

The minimum spanning tree T^* is initially oriented from an arbitrarily selected root vertex $r \in V$.

► **Definition 1.** The *predecessor* of any vertex $p \in V \setminus \{r\}$, indicated by $\text{Pred}(p)$, is defined as the vertex adjacent to p along the path between p and r in T^* .

► **Property 2.** Since T^* is a spanning tree, there exists a unique path between any vertex $p \in V \setminus \{r\}$ and r ; therefore the predecessor exists and is unique for each $p \in V \setminus \{r\}$.

$\text{Pred}(r)$ is set to a null value to indicate that the root vertex r has no predecessor.

► **Definition 3.** On the oriented tree T^* the *depth* of any vertex $p \in V$, indicated by $\text{Depth}(p)$, is defined as the number of edges between r and p .

From Definition 3 and Property 2 the following property follows.

► **Property 4.** For any two vertices u and v such that $v = \text{Pred}(u)$, it holds $\text{Depth}(u) = \text{Depth}(v) + 1$.

On the oriented tree T^* two additional values $\text{Dn}(p)$ and $\text{Up}(p)$ are defined for each vertex $p \in V$. Consider a depth-first-search visit of T^* and let us define a *move* to happen every time an edge is traversed in any direction.


► **Definition 5.** We define $\text{Dn}(r) = 1$. For all $p \in V \setminus \{r\}$, $\text{Dn}(p)$ is the *progressive number of the move* that reaches p from $\text{Pred}(p)$. For all $p \in V$, $\text{Up}(p)$ is the *progressive number of the move* that reaches $\text{Pred}(p)$ from p .

The following properties hold.

► **Property 6.**

- i. The indices $\text{Dn}(p)$ and $\text{Up}(p)$ have unique values in T^* .
- ii. Their values span the interval $[1, \dots, 2n]$.
- iii. For any vertex $v \in V$ and for each vertex $u \in V$ in the oriented subtree rooted at v , with $u \neq v$, $\text{Dn}(v) < \text{Dn}(u) < \text{Up}(u) < \text{Up}(v)$.

An example is illustrated in Figure 3.



■ **Figure 3** An oriented spanning tree with the values (Dn, Up) for each vertex.

Consider an edge $[i, j] \notin T^*$. When it is added to T^* a unique cycle, indicated by $C(i, j)$, is formed. Owing to the orientation of T^* , it is possible to define the *apex* of the cycle.

► **Definition 7.** The *apex* of a cycle $C(i, j)$ is the minimum depth vertex along it.

Owing to the orientation of T^* and to Property 6, the following properties hold.

► **Property 8.**

- i. For any cycle $C(i, j)$, the apex exists and it is unique;
- ii. if vertex p is the apex of $C(i, j)$, then $\text{Dn}(p) \leq \min\{\text{Dn}(i), \text{Dn}(j)\}$ and $\text{Up}(p) \geq \max\{\text{Up}(i), \text{Up}(j)\}$.

Let $\text{SubTree}(p, i, j)$ be a boolean function that tests Property 8ii, to check whether an endpoint of edge $[i, j]$ is the apex of $C(i, j)$ or not.

A consequence of Definition 7 and Property 8 is the following.

► **Property 9.** Given any edge $[i, j] \notin T^*$, if a vertex p verifies $\text{SubTree}(p, i, j)$ and it also belongs to $C(i, j)$, then it is the apex of $C(i, j)$.

Property 9 is exploited in Algorithm 8 and 9 described in Section 4.

Algorithm 1 The procedure that orients the minimum spanning tree.

```

1: procedure Orient
2:   for  $p \in V$  do
3:      $\delta(p) \leftarrow \emptyset$ 
4:   for  $[i, j] \in T^*$  do
5:      $\delta(i) \leftarrow \delta(i) \cup \{j\}$ 
6:      $\delta(j) \leftarrow \delta(j) \cup \{i\}$ 
7:    $r \leftarrow \text{Select}$ 
8:    $\text{Pred}(r) \leftarrow 0$ 
9:    $\text{Depth}(r) \leftarrow 0$ 
10:   $\alpha \leftarrow 0$ 
11:   $\text{DFS}(r)$ 

```

Algorithm 2 The recursive procedure to visit T^* in depth-first-search order.

```

1: procedure DFS( $p$ )
2:    $\alpha \leftarrow \alpha + 1$ 
3:    $\text{Dn}(p) \leftarrow \alpha$ 
4:   for  $k \in \delta(p)$  do
5:     if  $k \neq \text{Pred}(p)$  then
6:        $\text{Pred}(k) \leftarrow p$ 
7:        $\text{Depth}(k) \leftarrow \text{Depth}(p) + 1$ 
8:        $\text{DFS}(k)$ 
9:    $\alpha \leftarrow \alpha + 1$ 
10:   $\text{Up}(p) \leftarrow \alpha$ 

```

3.1.1 Procedure Orient

The values of Pred, Depth, Dn and Up are initially computed by the procedure Orient shown in Algorithm 1.

The procedure pre-computes the star $\delta(p)$ of each vertex $p \in V$, from the list of edges of T^* (lines 2-6). A root vertex r with no predecessor is arbitrarily selected and its predecessor and depth are set to 0 (lines 7-9). Then T^* is visited in depth-first-search order from vertex r with a call to the recursive procedure DFS (line 11).

A counter α counts every move in either direction along the edges of T^* . Every time a vertex p is reached from $\text{Pred}(p)$ by a call to $\text{DFS}(p)$, α is increased by 1 (line 2) and its value sets $\text{Dn}(p)$ (line 3). Then all the neighbors of vertex p are visited (lines 4-8) with the only exception of $\text{Pred}(p)$ (line 5). Before calling $\text{DFS}(k)$ for a generic neighbor $k \in \delta(p)$, the algorithm sets $\text{Pred}(k)$ at the current vertex p (line 6) and $\text{Depth}(k)$ at the value $\text{Depth}(p) + 1$ (line 7). When the neighborhood of p has been completely explored, α is increased again (line 9) and it sets $\text{Up}(p)$ (line 10). Finally the search backtracks.

Complexity of Orient.

The computation of the star of each vertex (lines 2 to 6) takes $O(n)$, because T^* contains $n - 1$ edges. All the operations on lines 7-10 of Orient can be done in $O(1)$. The time complexity taken by all instructions on lines 2, 3, 9 and 10 of DFS is $O(n)$, because α ranges from 1 to $2n$ by Property 6. The total number of iterations of the loop on lines 4-8 of DFS is twice the number of edges of T^* , i.e. $2(n - 1)$, and the body of the loop (lines 6 and 7) is executed in $O(1)$. Therefore the time complexity for visiting T^* with DFS is $O(n)$.

3.2 Local subgraphs


► **Definition 10.** For each vertex $p \in V$ a *local subgraph* $\mathcal{G}(p) = (\mathcal{V}(p), \mathcal{T}(p))$ is defined: it has $|\mathcal{V}(p)| = d(p)$ vertices, where $d(p)$ is the degree of p in T^* .

The vertices in $\mathcal{V}(p)$ are in one-to-one correspondence with the neighbors of p in T^* and therefore to the p -components they belong to. Hence, with a little abuse of notation in the remainder the same indices are used to indicate the neighbors of vertex $p \in V$ and the corresponding vertices in the local subgraph $\mathcal{V}(p)$.

Edges in local subgraphs are called *links* in the remainder, in order to make the difference with the edges of graph \mathcal{G} .

The link set $\mathcal{T}(p)$ is a forest for each $p \in V$; it is initially empty; at the end of the execution, it is a spanning tree of $\mathcal{G}(p)$. Each time the algorithm detects an alternative edge $[i, j]$ to be used when vertex p is deleted, this piece of information is recorded in $\mathcal{T}(p)$ (this happens in Algorithm 9 and Algorithm 10 described in Section 4). Each link in $\mathcal{T}(p)$ is represented as a record with two fields: one is the link itself between two vertices of $\mathcal{G}(p)$, and the other is the alternative edge $[i, j]$ reconnecting the two corresponding p -components.

An example is shown in Figure 4. When vertex p is deleted from the spanning tree, edges $[i, j]$ and $[k, l]$ are used to reconnect the three resulting p -components at minimum cost. Two corresponding links $[a, b]$ and $[b, c]$ are inserted in the local subgraph $\mathcal{G}(p)$, forming a spanning tree $\mathcal{T}(p)$. Each link in $\mathcal{T}(p)$ has an associated alternative edge in G .



■ **Figure 4** Left: the spanning tree T^* and two alternative edges (dashed). Right: two corresponding links $[a, b]$ and $[b, c]$ in the local subgraph $\mathcal{G}(p)$.

► **Remark 11.** The local subgraphs of the leaves of T^* have no links, since they contain a single vertex. Hence, they are not used in the algorithm.

► **Definition 12.** Links in $\mathcal{T}(p)$ incident to the local vertex corresponding to $\text{Pred}(p)$ are *vertical links*; links in $\mathcal{T}(p)$ not incident to the local vertex corresponding to $\text{Pred}(p)$ are *horizontal links*.

► **Remark 13.** The local forest $\mathcal{T}(r)$ cannot include vertical links, since r has no predecessor.

Figure 5 shows an example of vertical and horizontal links corresponding to a same alternative edge. When the candidate alternative edge $[i, j]$ is considered, a vertical link between the local vertices i and a is inserted in $\mathcal{G}(k)$, because k is a vertex between i and the apex; a horizontal link between the local vertices k and j is inserted in $\mathcal{G}(a)$, because a is the apex.


3.2.1 Local Union-Find data-structure

Since a spanning tree $\mathcal{T}(p)$ must be eventually computed for each local subgraph $\mathcal{G}(p)$, a Union-Find data-structure is kept for each vertex $p \in V$. It is made by an array of $d(p)$ linked lists. Each list $\mathcal{L}(p, k)$ is initialized with a single element $k \in \mathcal{V}(p)$, corresponding to a neighbor k of p in T^* . For each list two additional scalar values are also recorded: $\text{Card}(p, k)$ records the cardinality of $\mathcal{L}(p, k)$; $\text{Head}(p, k)$ records the head of the list to which the local vertex k belongs.

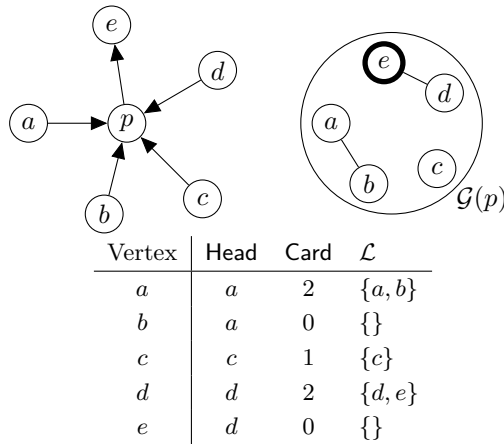
The operations executed on the Union-Find data-structure of vertex $p \in V$ when $\mathcal{T}(p)$ is enriched with a new link are the same as those reported in many textbooks on algorithms and data-structures (see for instance [4]). For readers unfamiliar with them, they are recalled in Appendix A, where the pseudo-code of the procedure `Merge` is also described. An example is shown in Figure 6.

3.2.2 Procedure Initialization

The procedure `Initialization`, shown in Algorithm 3, initializes the local data-structure of each vertex $p \in V$.



■ **Figure 5** An example of links added to local subgraphs when a candidate alternative edge is considered. In each local subgraph, the vertex corresponding to the predecessor is drawn in thick line.



■ **Figure 6** A vertex p with five neighbors (left), its local subgraph $\mathcal{G}(p)$ (center) and the corresponding state of the Union-Find data-structure (right). Oriented arcs indicate the predecessors in T^* . In $\mathcal{G}(p)$, the vertex corresponding to $e = \text{Pred}(p)$ is indicated by a thick line. Local vertices a and b have been connected with a horizontal link; consequently, the lists of a and b have been merged in a single list, so that $\mathcal{L}(a) = \{a, b\}$, $\mathcal{L}(b) = \emptyset$, $\text{Card}(a) = 2$, $\text{Card}(b) = 0$ and $\text{Head}(a) = \text{Head}(b) = a$. The local vertex d has been connected to $e = \text{Pred}(p)$ with a vertical link; consequently, the lists of d and e have been merged in a single list, so that $\mathcal{L}(d) = \{d, e\}$, $\mathcal{L}(e) = \emptyset$, $\text{Card}(d) = 2$, $\text{Card}(e) = 0$ and $\text{Head}(d) = \text{Head}(e) = d$.

The local vertex sets $\mathcal{V}(p)$ for all $p \in V$, initially empty (line 3), are populated in lines 5-6. The corresponding Union-Find data-structure is initialized in lines 7-12. The local forests $\mathcal{T}(p)$ are initially empty (line 14).

The value $\text{AltTreeCost}(p)$ indicates the total cost of the alternative edges that have been selected to reconnect the p -components for each vertex $p \in V$. It is initialized at 0 on line 15.

Each spanning tree $\mathcal{T}(p)$ in the local graph of vertex p requires $d(p) - 1$ edges. Since $\sum_{p \in V} d(p) = 2(n - 1)$, then the number of (not necessarily distinct) alternative edges that must be selected is $2(n - 1) - n$, i.e. $n - 2$. This is the initial value of the variable μ , that counts the number of missing alternative edges (line 16). The search for alternative edges terminates as soon as $\mu = 0$.

Algorithm 3 The procedure that initializes the local subgraphs.

```

1: procedure Initialization
2:   for  $p \in V$  do
3:      $\mathcal{V}(p) \leftarrow \emptyset$ 
4:   for  $[i, j] \in T^*$  do
5:      $\mathcal{V}(i) \leftarrow \mathcal{V}(i) \cup \{j\}$ 
6:      $\mathcal{V}(j) \leftarrow \mathcal{V}(j) \cup \{i\}$ 
7:      $\mathcal{L}(i, j) \leftarrow \{j\}$ 
8:      $\text{Card}(i, j) \leftarrow 1$ 
9:      $\text{Head}(i, j) \leftarrow j$ 
10:     $\mathcal{L}(j, i) \leftarrow \{i\}$ 
11:     $\text{Card}(j, i) \leftarrow 1$ 
12:     $\text{Head}(j, i) \leftarrow i$ 
13:  for  $p \in V$  do
14:     $\mathcal{T}(p) \leftarrow \emptyset$ 
15:     $\text{AltTreeCost}(p) \leftarrow 0$ 
16:   $\mu \leftarrow n - 2$ 

```

Complexity of Initialization.


The procedure `Initialization` contains three loops: the first and the third loop are repeated for each vertex in V , i.e. n times; the second loop is repeated for each edge of T^* , i.e. $n - 1$ times. All operations within the loops require $O(1)$. Therefore `Initialization` takes $O(n)$.

3.3 Oriented paths and trees

Consider a generic edge $[i, j] \notin T^*$ and the corresponding cycle $C(i, j)$, whose apex is indicated as $\text{apex}(i, j)$. Two oriented paths are defined as follows.

► **Definition 14.** The *oriented path* $P(i, j)$ goes from vertex i to $\text{apex}(i, j)$ in T^* ; the *oriented path* $P(j, i)$ goes from vertex j to $\text{apex}(i, j)$ in T^* .

It is possible that one of the two paths does not exist, when i or j is the apex of $C(i, j)$, as illustrated in Figure 7.



■ **Figure 7** Left: edge $[i, j]$ corresponds to two oriented paths, one including a and b , the other including c and d . Right: edge $[i, j]$ corresponds to a single path, because one of its endpoints is the apex of $C(i, j)$. In both cases, the apex is indicated by a thick line. Oriented arcs indicate predecessors in T^* .

► **Definition 15.** The *internal vertices* of a non-empty oriented path $P(i, j)$ (or $P(j, i)$) are the vertices along it excluding the endpoints, i.e. vertex i (or j) and $\text{apex}(i, j)$.

As observed in Section 2, $\text{apex}(i, j)$ and the internal vertices of $P(i, j)$ and $P(j, i)$ are those for which edge $[i, j]$ can be used as an alternative edge. In the local subgraph of $\text{apex}(i, j)$ edge $[i, j]$ corresponds to a horizontal link, while in the local subgraphs of the internal vertices along the two paths edge $[i, j]$ corresponds to a vertical link, as shown in Figure 5. For this reason, the apex is processed separately from the internal vertices along the two paths.

► **Definition 16.** The *root* of a non-empty oriented path $P(i, j)$ or $P(j, i)$ is its vertex t with $\text{Depth}(t) = \text{Depth}(\text{apex}(i, j)) + 1$.

For each internal vertex p of an oriented path $P(i, j)$ edge $[i, j]$ reconnects two p -components, one of them containing i and the other one containing $\text{Pred}(p)$ (and the same holds for $P(j, i)$ swapping i with j). For each candidate alternative edge $[i, j]$ the algorithm scans $P(i, j)$ from i and $P(j, i)$ from j up to their roots. Therefore all local forests of internal vertices of the two paths are considered to possibly insert a vertical link in each of them.

► **Definition 17.** Two oriented paths P' and P'' *overlap* if and only if they have at least an internal vertex in common.

When two (or more) oriented paths overlap, the algorithm merges them to form an *oriented tree*. The oriented trees made by the oriented paths have the following properties.

► **Property 18.**

- i. Each vertex belongs to at most one oriented tree;
- ii. each oriented tree has a unique *root*, that is the minimum depth root of its paths;
- iii. there is at least one vertical link in the local forest of each vertex in an oriented tree.

Owing to this property, every time the algorithm scans a path and it detects that part of it overlaps with an existing oriented tree, the oriented tree is skipped in $O(1)$ and the scan resumes directly from its root.

For this purpose, for each vertex $p \in V$ a variable $\text{Path}(p)$ records the index of the first path that introduces a vertical link in $\mathcal{G}(p)$.

► **Definition 19.** A path is *relevant* if and only if it inserts a vertical link in at least one subgraph.

► **Property 20.** Since the number of vertical links is bounded above by $n - 2$ (the limit case occurs when T^* is a path), this is also the maximum number of relevant paths.

Every time a new path is scanned, a path counter π is increased by 1 and the path root is initialized at 0. Every time a vertical link is inserted in some local forest $\mathcal{T}(p)$, the path root is updated to p . At the end of the scan, if the path root is still equal to 0, then the path is discarded as non-relevant and the path counter π is decreased by 1. No update to any data-structure is done, while the currently scanned path has not yet been recognized as relevant.

3.4 The Tree-Union-Find data-structure

A suitable data-structure is used to allow for efficiently merging relevant oriented paths into oriented trees. This data-structure is indicated as Tree-Union-Find, to distinguish it from the Union-Find data-structures associated with local subgraphs.

The Tree-Union-Find data-structure includes an array TList of linked lists, an array TCard, an array THead and an array TRoot. Each component of such arrays corresponds to an oriented tree, i.e. a set of overlapping relevant oriented paths: TList(τ) is the set of paths merged into the tree τ ; TCard(τ) is their number; THead(τ) is the index of a representative path among those in TList(τ); TRoot(τ) is the root of tree τ , i.e. the minimum depth vertex among the roots of the paths in TList(τ).

All arrays in the Tree-Union-Find data-structure are made by at most $n - 2$ elements, because no more than $n - 2$ relevant paths can exist, owing to Property 20.

Algorithm 4 shows the initialization of a new path. Algorithm 5 shows how non-relevant paths are detected and purged. Algorithm 6 illustrates the procedure that merges the subtrees containing two paths π' and π'' : TreeMerge assigns index τ' to the tree with maximum cardinality and τ'' to the other one, according to the values of TCard (lines 2-7). The root of the resulting tree is selected as the one with minimum depth among the root of τ' and the root of τ'' (lines 8-9). Then tree τ'' is appended to tree τ' , so that THead(π) $\leftarrow \tau'$ for all paths π in TList(τ'') (lines 10-11). Finally TCard and TList are updated, appending the shortest list to the longest one (lines 12-14).

4 The algorithm

After the introduction of the necessary definitions and properties and the description of the main data-structures, it is now possible to introduce the algorithm.

Algorithm 4 The algorithm that initializes a new oriented path.

```

1: procedure InitPath
2:    $\pi \leftarrow \pi + 1$ 
3:    $\text{TList}(\pi) \leftarrow \{\pi\}$ 
4:    $\text{TCard}(\pi) \leftarrow 1$ 
5:    $\text{THead}(\pi) \leftarrow \pi$ 
6:    $\text{TRoot}(\pi) \leftarrow 0$ 

```

Algorithm 5 The algorithm that deletes non-relevant paths.

```

1: procedure PurgePath( $\pi$ )
2:   if  $\text{TRoot}(\text{THead}(\pi)) = 0$  then
3:      $\pi \leftarrow \pi - 1$ 
4:      $\text{TList}(\pi) \leftarrow \emptyset$ 

```

Algorithm 6 The algorithm that merges two subtrees containing paths π' and π'' .

```

1: procedure TreeMerge( $\pi', \pi''$ )
2:   if  $\text{TCard}(\text{THead}(\pi')) > \text{TCard}(\text{THead}(\pi''))$  then
3:      $\tau' \leftarrow \text{THead}(\pi')$ 
4:      $\tau'' \leftarrow \text{THead}(\pi'')$ 
5:   else
6:      $\tau' \leftarrow \text{THead}(\pi'')$ 
7:      $\tau'' \leftarrow \text{THead}(\pi')$ 
8:   if  $(\text{Depth}(\text{TRoot}(\tau'')) < \text{Depth}(\text{TRoot}(\tau')))$  then
9:      $\text{TRoot}(\tau') \leftarrow \text{TRoot}(\tau'')$ 
10:  for  $\pi \in \text{TList}(\tau'')$  do
11:     $\text{THead}(\pi) \leftarrow \tau'$ 
12:   $\text{TCard}(\tau') \leftarrow \text{TCard}(\tau') + \text{TCard}(\tau'')$ 
13:   $\text{TCard}(\tau'') \leftarrow 0$ 
14:   $\text{TList}(\tau') \leftarrow \text{TList}(\tau') \cup \text{TList}(\tau'')$ 

```

The main body of the algorithm is shown in Algorithm 7. The complexity of each step is indicated on the corresponding line and it is proven in the remainder. The algorithm takes in input a graph $G = (V, E)$, an edge weighting function c , a minimum spanning tree T^* and a sorted list \mathcal{S} of all edges $e \in E$. It returns the optimal vertex p^* and the corresponding 1-tree T' providing the largest Held–Karp lower bound LB_{HK}^* .

Algorithm 7 The main algorithm.

```

1: procedure Main( $G, c, T^*, \mathcal{S}$ )
2:   // Find all alternative edges //
3:   Orient  $\triangleright O(n)$ 
4:   Initialization  $\triangleright O(n)$ 
5:   Search  $\triangleright O(m + n \log n)$ 
6:   // Compute the maximum Held–Karp lower bound  $LB_{HK}^*$  //
7:   SelectVertex  $\triangleright O(m)$ 
8:   Return( $p^*, T', LB_{HK}^*$ )

```

Procedures **Orient** and **Initialization** have been already described in Section 3. Procedure **SelectVertex** is described in the next Section; it is used to optimize the Held–Karp lower bound.

This Section describes in detail the procedure **Search**, that computes n alternative minimum spanning trees.

The **Search** procedure, shown in Algorithm 8, is made by a main loop (lines 3-18) whose body is divided into three parts. At each iteration of the main loop a candidate alternative edge is considered and all the corresponding vertical and horizontal links are inserted. The loop is repeated until all necessary links have been inserted; this is detected from the test on the number of missing edges in the local subgraphs: $\mu = 0$ (line 18).

Algorithm 8 The procedure that detects all alternative edges.

```

1: procedure Search
2:    $\pi \leftarrow 0$ 
3:   repeat
4:     /* Edge selection */
5:     repeat
6:        $[i, j] \leftarrow \text{Extract}(\mathcal{S})$ 
7:     until  $[i, j] \notin T^*$ 
8:     /* Vertical links */
9:      $u \leftarrow i$ 
10:    if  $\neg \text{SubTree}(i, i, j)$  then
11:       $\text{PathScan}(i, j, u)$ 
12:     $v \leftarrow j$ 
13:    if  $\neg \text{SubTree}(j, i, j)$  then
14:       $\text{PathScan}(j, i, v)$ 
15:    /* Horizontal links */
16:    if  $\neg \text{SubTree}(i, i, j) \wedge \neg \text{SubTree}(j, i, j) \wedge \text{then}(\text{Stop}(u) \vee \text{Stop}(v))$ 
17:       $\text{ProcessApex}(u, v, i, j)$ 
18:  until  $\mu = 0$ 

```

The first part of the body of the main loop (lines 4-7) is a secondary loop in which the next candidate alternative edge (i.e. an edge not in T^*) is extracted from the sorted edge list \mathcal{S} . The candidate alternative edge is indicated by $[i, j]$.

In the second part (lines 8-14) `Search` scans the paths $P(i, j)$ and $P(j, i)$ to possibly insert vertical links. Two indices u and v are used to identify the current vertex on the paths: u along $P(i, j)$ and v along $P(j, i)$. The operations on the two paths are symmetrical and they are grouped in the procedure `PathScan` that is called twice (lines 11 and 14). `PathScan` is called only if the corresponding path exists; this is the reason for tests on lines 10 and 13. `SubTree`(p, i, j) is a boolean function that tests Property 8 (ii), to check whether an endpoint of edge $[i, j]$ is the apex of $C(i, j)$ or not. In this way empty paths are disregarded.

The effect of this second part is twofold: first, to insert all possible vertical links in local forests; second, to indicate how the search along each path terminates. This is represented by the value of the variables `Stop`(u) and `Stop`(v). Consider the side of vertex i . If `Stop`(u) is true, the current vertex u is within the cycle $C(i, j)$; hence, when the loop is over, `Pred`(u) is `apex`(i, j). If `Stop`(u) is false, the path on the side of i has been merged with a pre-existing oriented tree, whose root is at `apex`(i, j) or above (i.e. closer to the root) and therefore `Pred`(u) is out of $C(i, j)$. The same holds symmetrically for `Stop`(v) on the side of j .

In the third part (lines 15-17) a horizontal link is possibly inserted in the local forest of `apex`(i, j) by a call to procedure `ProcessApex` (line 17). However, this is done only if the apex is different from i and j (i.e. both oriented paths exist) and if at least one of the two current vertices u and v is within the cycle $C(i, j)$. If both u and v have reached the apex, then both the `apex`(i, j)-components of i and j are already connected with that of `Pred`(`apex`(i, j))) and therefore no horizontal link must be inserted in \mathcal{T} (`apex`(i, j))).

4.1 Procedure `PathScan` and vertical links

For each side of the cycle $C(i, j)$ the current vertex is initialized at an endpoint of the candidate alternative edge $[i, j]$ (lines 9 and 12). The variable `Stop` is initialized at true (line 2) and a new path is initialized by a call to `InitPath` (line 3) so that its index is π .

The current vertex, indicated by w , corresponds either to u or to v , depending on the call on line 11 or 14 of `Search`. The test on line 4 detects whether the predecessor of the current vertex w is at a larger depth than `apex`(i, j). As soon as `Pred`(w) reaches the apex or above, the loop ends, because vertical links must be inserted only in internal vertices of the oriented path.

Within the loop (lines 4-22) at each iteration an attempt is made to insert a vertical link between w and `Pred`(p) in $\mathcal{G}(p)$, where $p = \text{Pred}(w)$ (line 5). For this purpose the local Union-Find data-structure of vertex p is tested to check whether w and `Pred`(p) are already connected or not (line 6).

Algorithm 9 The procedure that inserts vertical links.

```

1: procedure PathScan( $i, j, w$ )
2:   Stop( $w$ )  $\leftarrow$  true
3:   InitPath
4:   while  $\neg$ SubTree(Pred( $w$ ),  $i, j$ ) do
5:      $p \leftarrow$  Pred( $w$ )
6:     if (Head( $p, w$ ))  $\neq$  Head( $p, \text{Pred}(p)$ ) then
7:       /* Insert a vertical link */
8:        $\mathcal{T}(p) \leftarrow \mathcal{T}(p) \cup \{([w, \text{Pred}(p)], [i, j])\}$ 
9:       AltTreeCost( $p$ )  $\leftarrow$  AltTreeCost( $p$ ) +  $c(i, j)$ 
10:       $\mu \leftarrow \mu - 1$ 
11:      Merge( $p, w, \text{Pred}(p)$ )
12:      TRoot(THead( $\pi$ ))  $\leftarrow p$ 
13:       $w \leftarrow p$ 
14:      if Path( $p$ ) = 0 then
15:        Path( $p$ )  $\leftarrow \pi$ 
16:      else
17:        /* Skip Path( $p$ ) up to its root */
18:        if TRoot(THead( $\pi$ ))  $\neq$  0 then
19:          TreeMerge( $\pi, \text{Path}(p)$ )
20:         $w \leftarrow$  TRoot(THead(Path( $p$ )))
21:        if SubTree( $w, i, j$ ) then
22:          Stop( $w$ )  $\leftarrow$  false
23:  PurgePath( $\pi$ )

```

If the test succeeds (i.e. w and $\text{Pred}(p)$ are not connected), a vertical link is inserted and several update operations are done (lines 7-13). In particular, a new element is inserted in $\mathcal{T}(p)$ as a record with two fields: the link is $[w, \text{Pred}(p)]$, while the alternative edge is $[i, j]$ (line 8); the total cost of the alternative edges $\text{AltTreeCost}(p)$ is increased by $c(i, j)$ (line 9); the number of missing alternative edges is decreased by 1 (line 10). Then the local Union-Find data-structure of vertex p is updated, merging the connected components of w and $\text{Pred}(p)$ (line 11) with a call to `Merge` (see Appendix A for the pseudo-code). When a vertical link is inserted, the path is relevant and its root is updated to vertex p where the vertical link has been inserted (line 12). Finally the current vertex w is moved one step forward along the path to $p = \text{Pred}(w)$ (line 13).

That a vertical link has been inserted or not, two situations can occur: either no vertical link has been inserted in $\mathcal{T}(p)$ in previous iterations (and in this case $\text{Path}(p) = 0$) or a vertical link has already been inserted while scanning a path (and in this case the index of that path has been recorded in $\text{Path}(p)$). In the former case, detected by the test on line 14, $\text{Path}(p)$ is set to the index π of the current path (line 15) and no other operation is required. In the latter case, the current vertex w is moved directly to the root of the oriented subtree containing the path $\text{Path}(p)$ and the subtree containing the current path π is merged with the subtree containing path $\text{Path}(p)$. The `TreeMerge` procedure is called (line 19) only if π has already been detected as relevant (line 18). The current vertex is updated on line 20. Consequently, a new test is done to check whether w is still below the apex of $C(i, j)$ (line 21). If not, `Stop` is set to false (line 22).

When the loop is over, π is checked again: if it is non-relevant then it is purged (line 23).

4.2 Procedure ProcessApex and horizontal links

The procedure `ProcessApex`(u, v, i, j), shown in Algorithm 10, inserts a horizontal link in the local forest of `apex`(i, j) if and only if the `apex`(i, j)-components of i and j are not already connected in $\mathcal{T}(\text{apex}(i, j))$. If they are already connected, the procedure has no effect.

To check whether the horizontal link can be inserted, it is necessary to know the indices of the two vertices adjacent to the apex along $P(i, j)$ and $P(j, i)$. These correspond to u and v when `Stop`(u) and `Stop`(v) are true: if `Stop`(u) \wedge `Stop`(v), then $\text{Pred}(u) = \text{Pred}(v) = \text{apex}(i, j)$. However, this is not the case when `Stop` is false. Hence, the apex, indicated by p in Algorithm 10, is found as the predecessor of the current vertex for which `Stop` is true (lines 2-5). At least one of `Stop`(u) and `Stop`(v) is guaranteed to be true, owing to the test on line 16 of `Search`.

If one of the two current vertices, say u , has been moved to the apex or above, i.e. $\text{Stop}(u)$ is false, then the p -component of i is already connected with the p -component of $\text{Pred}(p)$ in $\mathcal{T}(p)$. Hence a test on the local Union-Find data-structure of vertex p must be done to check whether v and $\text{Pred}(p)$ are connected or not. This is the reason for setting a variable u' to u if $\text{Stop}(u)$ is true and to $\text{Pred}(p)$ if $\text{Stop}(u)$ is false (lines 6-9). The same is done for v' (lines 10-13).

If the test on the Union-Find data-structure succeeds (line 14), then a horizontal link is inserted; otherwise the procedure terminates with no effect.

Before inserting the horizontal link, it is necessary to find the index of both vertices adjacent to the apex p along $P(i, j)$ and $P(j, i)$. They are not both available if Stop is false for one of the two current vertices. Therefore, u or v is reset to the position just below the apex, in case it is not (lines 15-16 and 17-18). This is done by a procedure Find , that exploits the values of Up and Dn of all vertices adjacent to p . W.l.o.g. assume $\text{Stop}(u)$ be false. The execution of $\text{Find}(p, i)$ implies a search among the vertices of $\mathcal{V}(p)$ to find the vertex which lies on the path between i and p . This requires to find the (unique) vertex $u \in \mathcal{V}(p)$ such that $(\text{Dn}(u) \leq \text{Dn}(i)) \wedge (\text{Up}(u) \geq \text{Up}(i))$.

Once u and v have been set to the neighbors of the apex p along $P(i, j)$ and $P(j, i)$, the horizontal link between them is finally inserted in $\mathcal{T}(p)$. The two fields of the new element in $\mathcal{T}(p)$ record the link $[u, v]$ and the alternative edge $[i, j]$ (line 19); the cost of $\mathcal{T}(p)$ is increased by $c(i, j)$ (line 20); the number μ of missing edges is decreased by 1 (line 21). Finally the Union-Find data-structure of the local subgraph $\mathcal{G}(p)$ is updated by a call to Merge (line 22) to connect the components of u and v (see Appendix A for the pseudo-code).

Algorithm 10 The procedure that inserts horizontal links.

```

1: procedure ProcessApex( $u, v, i, j$ )
2:   if ( $\text{Stop}(u)$ ) then
3:      $p \leftarrow \text{Pred}(u)$ 
4:   else
5:      $p \leftarrow \text{Pred}(v)$ 
6:   if ( $\text{Stop}(u)$ ) then
7:      $u' \leftarrow u$ 
8:   else
9:      $u' \leftarrow \text{Pred}(p)$ 
10:  if ( $\text{Stop}(v)$ ) then
11:     $v' \leftarrow v$ 
12:  else
13:     $v' \leftarrow \text{Pred}(p)$ 
14:  if ( $\text{Head}(p, u') \neq \text{Head}(p, v')$ ) then
15:    if  $\neg \text{Stop}(u)$  then
16:       $u \leftarrow \text{Find}(p, i)$ 
17:    if  $\neg \text{Stop}(v)$  then
18:       $v \leftarrow \text{Find}(p, j)$ 
19:     $\mathcal{T}(p) \leftarrow \mathcal{T}(p) \cup \{[u, v], [i, j]\}$ 
20:     $\text{AltTreeCost}(p) \leftarrow \text{AltTreeCost}(p) + c(i, j)$ 
21:     $\mu \leftarrow \mu - 1$ 
22:     $\text{Merge}(p, u, v)$ 

```

Complexity of Search.

The number of iterations of the outer loop of Search (lines 3-18) is $O(m)$; therefore the time taken by all constant time operations in Search (including SubTree that takes constant time), PathScan and ProcessApex (including InitPath and PurgePath), yield an overall $O(m)$ contribution to the time complexity.

The loop on lines 5-7 can be implemented so that it takes $O(1)$ time per iteration, i.e. $O(m)$ overall. This requires to sort the edges of T^* with the same criterion used to sort the edges in \mathcal{S} (non-decreasing cost plus some additional lexicographic criterion to break ties). This pre-sorting step requires $O(n \log n)$ because T^* contains $n - 1$ edges. In this way the test on line 7 is a comparison between the edge $[i, j]$ extracted from \mathcal{S} and the first

edge of T^* not yet extracted from \mathcal{S} and it takes constant time. Hence, the pre-sorting step (not indicated in the pseudo-code) and the loop on lines 5-7 require $O(m + n \log n)$ time.

The block of operations on lines 7-13 of `PathScan` is executed $O(n)$ times, because a vertical link is inserted each time and the number of possible vertical links is $O(n)$. Since all operations excepted `Merge` have $O(1)$ complexity, their overall contribution to the time complexity of the algorithm is $O(n)$.

The test on line 14 of `PathScan` can succeed at most n times. Therefore the total contribution to the time complexity of the $O(1)$ operation on line 15 is $O(n)$.

The block of operations on lines 17-22 of `PathScan` takes $O(1)$ with the exception of the contribution of `TreeMerge`. The complexity analysis requires to distinguish the case in which π is non-relevant ($\text{TRoot}(\text{THead}(\pi)) = 0$) from the case in which π is relevant ($\text{TRoot}(\text{THead}(\pi)) > 0$).

An empty path can exist only twice for each candidate alternative edge, once for each side of $C(i, j)$. Therefore the execution of this block with non-relevant π can occur $O(m)$ times. In these cases `TreeMerge` is not executed and thus the total contribution is $O(m)$.

The number of relevant paths in the `Tree-Union-Find` data-structure is $O(n)$, because all roots are different. Therefore, a call to `TreeMerge` can occur at most $O(n)$ times. Hence, the number of times the block is executed with relevant π is $O(n)$ and therefore the total contribution of these iterations is $O(n)$ plus the contribution of `TreeMerge`.

The number of times the loop on lines 4-22 of `PathScan` is executed is the sum of both types of iterations, i.e. with relevant π and non-relevant π . Therefore the tests on lines 6 and 14 are executed $O(m + n)$ times.

The procedure `ProcessApex` is called at most once for each edge, i.e. $O(m)$ times. Therefore all $O(1)$ operations on lines 2-13 contribute $O(m)$.

The test on line 14 can succeed $O(n)$ times, because the number of horizontal links that can be inserted is $O(n)$. Excluding the contribution of `Merge` and `Find`, the total contribution of the operations on lines 15-22 to the time complexity is $O(n)$.

The execution of `Find(p, i)` implies a search among the vertices of $\mathcal{V}(p)$ (i.e. the neighbors of p in T^*) to find the vertex which lies on the path between p and i on T^* , i.e. the (unique) vertex u satisfying $(\text{Dn}(u) \leq \text{Dn}(i)) \wedge (\text{Up}(u) \geq \text{Up}(i))$. This task can be accomplished in $O(\log d(p))$ (which is not worse than $O(\log n)$) by binary search, because the vertices of $\mathcal{V}(p)$ are sorted by their values of `Dn` and `Up` after the execution of `Orient`. Since the number of calls to `Find` is $O(n)$, as observed above, the overall contribution of `Find` to the total time complexity is $O(n \log n)$.

At each call of `Merge` two lists in the local subgraph $\mathcal{G}(p)$ of some vertex $p \in V$ are merged. This can happen $O(n)$ times overall, because $O(n)$ alternative edges must be found. For the well-known property of the `Union-Find` data-structure (recalled in Appendix A) for each local graph with $d(p)$ vertices, the time taken by the update operations is $O(d(p) \log d(p))$. Summing up these contributions over all vertices results in an $O(n \log n)$ contribution to the complexity of the whole algorithm, because $\sum_{p \in V} d(p) \log d(p) \leq \sum_{p \in V} d(p) \log n = \log n \sum_{p \in V} d(p) = 2|T^*| \log n = 2(n - 1) \log n$.

The total time taken by `TreeMerge` to merge trees is $O(n \log n)$, because of the properties of `Union-Find`: every time two or more trees are merged, their lists are merged so that the shortest one is appended to the longest one. This guarantees that no representative is updated more than $\log n$ times since the number of oriented trees is $O(n)$. Note that to achieve this result, it is necessary to accept that the root of an oriented tree does not necessarily belong to its representative path.

The conclusion of this analysis is that the worst-case time complexity of `Search` is $O(m + n \log n)$ which is the same of a single run of `Kruskal` algorithm or `Prim` algorithm implemented with `Fibonacci` heaps.

5 Computation of the optimal Held–Karp lower bound LB_{HK}^*

If the computation of the n alternative minimum spanning trees is aimed at the optimization of the Held–Karp lower bound $LB_{HK}(\bar{p})$, the procedure `SelectVertex` is executed, as shown in Algorithm 7, to compute the most profitable vertex \bar{p} , the resulting 1-tree and the corresponding value LB_{HK}^* .

The procedure `SelectVertex`, illustrated in Algorithm 11, finds the two minimum cost edges $e_1(p)$ and $e_2(p)$ incident in each vertex p (their costs are indicated by c_1 and c_2) and evaluates what the resulting 1-tree lower bound would be if p were removed from T^* , replacing its star with the alternative edges recorded in $\mathcal{T}(p)$ plus the two edges $e_1(p)$ and $e_2(p)$. The vertex p^* providing the largest value is returned, together with the corresponding

value z , indicating the best lower bound that can be achieved, LB_{HK}^* . The value $\text{StarCost}(p)$ indicates the total cost of the edges of T^* incident to each vertex $p \in V$.

In the last two lines of Algorithm 11 the optimal Held–Karp 1-tree T' is produced. Each element in $\mathcal{T}(p)$ is stored as a record with two fields, one indicating the link in the local forest and the other indicating the corresponding alternative edge. In the pseudo-code shown in Algorithm 11 the notation $\widehat{\mathcal{T}}(p)$ indicates the set of alternative edges corresponding to the elements of $\mathcal{T}(p)$.

Complexity of SelectVertex.

The computation of $\text{StarCost}(p)$ for each vertex requires to scan each edge of the graph twice. Hence its complexity is $O(m)$. The operation of deleting from T^* the edges incident to p^* (line 23) takes $O(n)$. The operation of merging the two lists of edges T^* and $\widehat{\mathcal{T}}(p)$ (line 23) takes $O(n)$. Therefore the overall complexity of **SelectVertex** is $O(m)$.

Algorithm 11 The procedure that computes the optimal Held–Karp lower bound.

```

1: procedure SelectVertex
2:    $z \leftarrow -\infty$ 
3:   for  $p \in V$  do
4:      $\text{StarCost}(p) \leftarrow 0$ 
5:      $c_1 \leftarrow \infty$ 
6:      $c_2 \leftarrow \infty$ 
7:     for  $k \in \delta(p)$  do
8:        $\text{StarCost}(p) \leftarrow \text{StarCost}(p) + c(p, k)$ 
9:       if  $c(p, k) < c_1$  then
10:         $c_2 \leftarrow c_1$ 
11:         $c_1 \leftarrow c(p, k)$ 
12:         $e_2 \leftarrow e_1$ 
13:         $e_1 \leftarrow [p, k]$ 
14:       else
15:         if  $c(p, k) < c_2$  then
16:           $c_2 \leftarrow c(p, k)$ 
17:           $e_2 \leftarrow [p, k]$ 
18:       if  $\text{AltTreeCost}(p) - \text{StarCost}(p) + c_1 + c_2 > z$  then
19:         $z \leftarrow \text{AltTreeCost}(p) - \text{StarCost}(p) + c_1 + c_2$ 
20:         $p^* \leftarrow p$ 
21:         $e_1^* \leftarrow e_1$ 
22:         $e_2^* \leftarrow e_2$ 
23:    $T' \leftarrow T^* \setminus \delta(p^*) \cup \widehat{\mathcal{T}}(p^*) \cup \{e_1^*, e_2^*\}$ 
24:   Return( $p^*, T', z$ )

```

6 Conclusions

The value of the 1-tree-based lower bound for the TSP can be optimized, not only following Helsgaun’s idea of optimally selecting a leaf of a minimum spanning tree and an additional minimum cost edge incident to it, but also by improving upon the original idea by Held and Karp of selecting a vertex \bar{p} excluded from a minimum spanning tree and two additional minimum cost edges incident to it. The algorithm presented here shows that the vertex \bar{p} that Held and Karp selected in an arbitrary way can be selected in an optimal way with the same worst-case time complexity required to compute a minimum spanning tree on the same weighted graph, i.e. $O(m + n \log n)$; this is also the time complexity required by each iteration of the Held–Karp method when the vertex \bar{p} is selected in an arbitrary way. The resulting lower bound is guaranteed to be larger than or equal to the lower bound provided by Helsgaun’s method.

The advantage of LB_{HK}^* with respect to LB_H^* obviously depends on the specific instance. Just to have a rough estimate of the expected improvement that LB_{HK}^* can provide, some limited computational tests were

also carried out. Two data-sets were generated at random. Data-set A is made of 100 random graphs with 100 vertices and density 15%; edge weights were generated with uniform probability distribution in the range $[0, 1000]$. Data-set B is made of 100 Euclidean instances; each of them is defined by 100 points placed at random in a square of size 100, with uniform probability distribution for both coordinates; graphs are complete and edge weights correspond to Euclidean distances.

Results are reported in Table 1. Column \overline{LB}_{HK} indicates the average value of LB_{HK} , i.e. the expected value of the lower bound when the vertex \bar{p} is chosen at random as suggested by Held and Karp. Column LB_H^* reports the optimized value of the 1-tree lower bound with Helsgaun method. Column LB_{HK}^* reports the optimized value of the Held–Karp lower bound. The rightmost column reports the number of times LB_{HK}^* was found to be tighter than LB_H^* . All values have been rounded to the closest integer. The experimental comparison shows that not only LB_{HK}^* dominates LB_H^* but it is strictly better for more than half of the instances. The improvement of LB_{HK}^* over LB_H^* is instance-dependent: in our tests it is zero or negligible for some instances, significant for others. The improvement of LB_{HK}^* over \overline{LB}_{HK} is always relevant (about 10 – 20%). Computing time was negligible and therefore it is not reported in Table 1.

Data-set	\overline{LB}_{HK}	LB_H^*	LB_{HK}^*	n. impr.
A	8990	11686	11759	55/100
B	530	595	597	63/100

■ **Table 1** Comparison between optimized and non-optimized 1-tree-based lower bounds.

It should be remarked that the algorithm illustrated in this paper computes a lower bound to the TSP, not a solution: it is conceived as a sub-routine within a more complex algorithm, such as a branch-and-bound algorithm, devised to solve some combinatorial problem implying Hamiltonian cycles. Therefore, a meaningful evaluation of the practical usefulness of the algorithm can be obtained only by evaluating the effect it produces on the computing time of the algorithm in which it is embedded. Hence the efficient code implementation of this algorithm can be a research topic deserving investigation.

Independently of the application to the Held–Karp lower bound, this paper provides an $O(m + n \log n)$ algorithm to pre-compute the necessary alternative edges, so that a minimum spanning tree can be immediately restored in a graph if a vertex is deleted (e.g. in case of node failure in a telecommunication network).

Acknowledgements

While developing his master thesis, Luca Diedolo made some preliminary experiments comparing LB_{HK}^* with LB_H^* , triggering the study presented in this paper. Two anonymous reviewers and the editor provided very useful and detailed comments and suggestions.

A Union-Find

The Union-Find data-structure is used to represent a partition of a ground set N into subsets that are iteratively merged until a single set coinciding with the ground set is obtained. The Union-Find data-structure allows to efficiently perform two operations:

- retrieving which is the subset in which a given element $k \in N$ is contained;
- merging two subsets in a single one.

A list $\mathcal{L}(k)$ is associated with each element $k \in N$. Each list has a distinguished element acting as the head of the list. The Union-Find data-structure also includes two vectors, indicated by **Head** and **Card**, with as many components as the number of elements in the ground set. For each $k \in N$, **Head**(k) indicates the element that is the head of the list containing k , while **Card**(k) indicates the number of elements in $\mathcal{L}(k)$. Initially $\mathcal{L}(k) = \{k\}$, **Card**(k) = 1 and **Head**(k) = k for all $k \in N$.

The list to which a given element $k \in N$ belongs can be retrieved in $O(1)$, because it is indicated by **Head**(k). Hence, the test whether two elements $u \in N$ and $v \in N$ belong to the same subset or not translates into testing whether **Head**(u) = **Head**(v).

When two distinct lists with heads in $u \in N$ and $v \in N$ must be merged, the list with smallest cardinality is appended to the list with largest cardinality. This is done by updating pointers in $O(1)$, but it requires to also update **Head** and **Card**. Assuming $\mathcal{L}(v)$ is appended to $\mathcal{L}(u)$, the values of **Card** are updated as follows:

$\text{Card}(u) \leftarrow \text{Card}(u) + \text{Card}(v)$ and $\text{Card}(v) \leftarrow 0$. Then, $\text{Head}(k)$ must be set to u for all elements $k \in \mathcal{L}(v)$. This operation takes linear time, for each iteration. However, appending the shortest list to the longest one guarantees that the value of Card for all elements in the shortest list doubles at least. Hence, no element is subject to the update operation more than $\log n$ times, where $n = |N|$. Therefore the update operations require at most $O(\log n)$ for each element, i.e. $O(n \log n)$ overall.

Procedure Merge is called by PathScan and ProcessApex, when a link is inserted in a local forest. It is illustrated in Algorithm 12.

Algorithm 12 The procedure that merges two lists in the Union-Find data-structure of a local subgraph $\mathcal{G}(p)$.

```

1: procedure Merge( $p, u, v$ )
2:   if  $\text{Card}(p, \text{Head}(p, u)) < \text{Card}(p, \text{Head}(p, v))$  then
3:      $k' \leftarrow v$ 
4:      $k'' \leftarrow u$ 
5:   else
6:      $k' \leftarrow u$ 
7:      $k'' \leftarrow v$ 
8:   for  $k \in \mathcal{L}(p, \text{Head}(p, k''))$  do
9:      $\text{Head}(p, k) \leftarrow \text{Head}(p, k')$ 
10:   $\text{Card}(p, \text{Head}(p, k')) \leftarrow \text{Card}(p, \text{Head}(p, k')) + \text{Card}(p, \text{Head}(p, k''))$ 
11:   $\text{Card}(p, \text{Head}(p, k'')) \leftarrow 0$ 
12:   $\mathcal{L}(p, \text{Head}(p, k')) \leftarrow \mathcal{L}(p, \text{Head}(p, k')) \cup \mathcal{L}(p, \text{Head}(p, k''))$ 

```

References

- 1 David L. Applegate, Robert E. Bixby, Vašek Chvátal, and William J. Cook. *The Traveling Salesman Problem: A Computational Study*. Princeton Series in Applied Mathematics. Princeton University Press, 2006.
- 2 Pascal Benchimol, Jean-Charles Régin, Louis-Martin Rousseau, Michel Rueher, and Willem-Jan van Hoeve. Improving the Held and Karp Approach with Constraint Programming. In *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems. CPAIOR 2010*, volume 6140 of *Lecture Notes in Computer Science*, pages 40–44. Springer, 2010.
- 3 Chandra Chekuri and Kent Quanrud. Approximating the Held-Karp Bound for Metric TSP in Nearly-Linear Time. In *Proceedings of the 58th Annual IEEE Symposium on Foundations of Computer Science*, pages 789–800. IEEE Computer Society, 2017.
- 4 Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms*. MIT Press, 2009.
- 5 Gregory Gutin and Abraham P. Punnen, editors. *The Traveling Salesman Problem and Its Variations*. Springer, 2007.
- 6 Michael Held and Richard M. Karp. The Traveling-Salesman Problem and Minimum Spanning Trees. *Oper. Res.*, 18:1138–1162, 1970.
- 7 Keld Helsgaun. An effective implementation of the Lin-Kernighan traveling salesman heuristic. *Eur. J. Oper. Res.*, 126:106–130, 2000.
- 8 Christine L. Valenzuela and Antonia J. Jones. Estimating the Held-Karp lower bound for the geometric TSP. *Eur. J. Oper. Res.*, 102:157–175, 1997.
- 9 Ton Volgenant and Roy Jonker. A branch and bound algorithm for the symmetric traveling salesman problem based on the 1-tree relaxation. *Eur. J. Oper. Res.*, 9:83–89, 1982.
- 10 David P. Williamson. Analysis of the Held-Karp lower bound for the asymmetric TSP. *Oper. Res. Lett.*, 12:83–88, 1992.