A multi-start local search algorithm for the Hamiltonian completion problem on undirected graphs

Jorik Jooker, Pieter Leyman, Patrick De Causmaecker

KU Leuven Kulak, Department of Computer Science, CODES, Etienne Sabbelaan 53, 8500 Kortrijk (Belgium)
jorik.jooker@kuleuven.be, pieter.leyman@kuleuven.be, patrick.decausmaecker@kuleuven.be

Abstract

This paper proposes a local search algorithm for a specific combinatorial optimisation problem in graph theory: the Hamiltonian Completion Problem (HCP) on undirected graphs. In this problem, the objective is to add as few edges as possible to a given undirected graph in order to obtain a Hamiltonian graph. This problem has mainly been studied in the context of various specific kinds of undirected graphs (e.g. trees, unicyclic graphs and series-parallel graphs). The proposed algorithm, however, concentrates on solving HCP for general undirected graphs. It can be considered to belong to the category of matheuristics, because it integrates an exact linear time solution for trees into a local search algorithm for general graphs. This integration makes use of the close relation between HCP and the minimum path partition problem, which makes the algorithm equally useful for solving the latter problem. Furthermore, a benchmark set of problem instances is constructed for demonstrating the quality of the proposed algorithm. A comparison with state-of-the-art solvers indicates that the proposed algorithm is able to achieve high-quality results.

Keywords: Metaheuristics; Matheuristics; Combinatorial optimisation; Hamiltonian completion problem; Minimum path partition problem

1 Introduction

The problem of determining whether a given undirected graph contains a Hamiltonian cycle is one of Karp's famous 21 NP-complete problems [22]. Being such a famous problem, it has been studied intensively for many years by the research community. This research has also led to many variations of this problem. One such variation is an interesting generalisation of the original problem in an optimisation context: the Hamiltonian Completion Problem (HCP), which is the main topic of this paper.

In this paper, we develop a matheuristic [8] for HCP on undirected graphs. We will identify a restricted version of this problem that can be solved exactly in polynomial time. The solution of this

restricted problem is integrated into a multi-start local search algorithm that attempts to solve HCP on general undirected graphs. This algorithm will repeatedly perturb the solution of the restricted problem to obtain a high-quality solution for the original problem.

The remainder of this paper is organised as follows: in Section 2 a formal description of HCP is given. We further motivate this problem by discussing applications in Section 3. Section 4 contains a brief overview of earlier work that has been done for this problem. We highlight those results from literature that are relevant to understand the remainder of this paper. In Section 5 the multi-start local search algorithm that we propose for solving HCP is explained. The explanation is built in a bottom-up fashion: we first describe the separate components of the algorithm, after which the complete algorithm is explained. In Section 6 we will describe how a benchmark set of problem instances is generated. These problem instances will be used to compare the algorithm from this paper with state-of-the-art algorithms. The results of these experiments are described in Section 7. Finally, a conclusion and possible avenues for further work are discussed in Section 8.

2 Problem description

To be able to define the Hamiltonian Completion Problem (HCP) on undirected graphs [17], it is necessary to introduce the notion of a Hamiltonian cycle and a Hamiltonian graph. A Hamiltonian cycle is a cycle that visits each node exactly once, except for the first and the last node, which are visited twice. A Hamiltonian graph is a graph that has at least one Hamiltonian cycle. HCP on undirected graphs is then the following problem: given an undirected graph $G = (V, E)$, find a set of edges E' such that $G' = (V, E \cup E')$ is a Hamiltonian graph and the size of E' is as small as possible.

As an example, consider Fig. 1. In this figure, the graph on the left is the graph G for which HCP has to be solved. The graph in the middle shows G' , which is obtained by adding 1 edge to G (shown in red). The graph on the right illustrates that G' is indeed a Hamiltonian graph by showing the existence of a Hamiltonian cycle (the cycle $1 - 2 - 3 - 5 - 4 - 7 - 9 - 8 - 6 - 1$).

Figure 1: Graph G , graph G' and a Hamiltonian cycle in G'

3 Applications

The need for solving HCP arises naturally in many applications. A first example application involving frequency assignment to transmitters has first been identified in [13] and has later been thoroughly described in [21]: “Assume that one has n radio transmitters that each have to send a signal by using frequency modulation (FM). To avoid interference, it is important that the frequencies that are used by the radio transmitters to send their signal are sufficiently dissimilar. To do this, the continuous frequency spectrum can be divided into several equally large, non-overlapping zones. We will number the k subsequent frequency zones from 1 to k . The goal is now to assign a frequency zone to every radio transmitter such that every frequency zone is assigned at most once to a radio transmitter. Every radio transmitter can only send signals with a frequency that is within the range of its assigned frequency zone. Furthermore, there is a constraint stemming from the fact that some radio transmitters are physically near each other: for all such pairs of radio transmitters, the assigned frequency zones cannot be neighbouring zones.

More formally the problem is as follows: let $f(i)$ be the number (of the frequency zone) that is assigned to radio transmitter i . If $i \neq j$, we demand that also $f(i) \neq f(j)$. Furthermore, if radio transmitter i is physically close to radio transmitter j , we demand that $|f(i) - f(j)| > 1$. We assume that k is large enough such that at least one valid assignment exists. The problem is now to assign frequency zones to radio transmitters such that:

$$\max_{1 \leq i \leq n} f(i) - \min_{1 \leq i \leq n} f(i) \quad (1)$$

is minimised (i.e. to use the least possible amount of space of the available frequency spectrum).

Let the assignment defined by f^* be any optimal assignment and let $(p_1, p_2, \dots, p_n)$ be a permutation of the numbers from 1 to n such that $f^*(p_1) < f^*(p_2) < \dots < f^*(p_n)$. It is clear that the shifted assignment g^* defined by $g^*(p_i) = f^*(p_i) - f^*(p_1) + 1$ is also an optimal assignment. Hence, for this assignment it holds that $g^*(p_1) = 1$. Furthermore, in the optimal solution it is impossible that $g^*(p_{i+1}) - g^*(p_i) > 2$, because then we could obtain a better solution by subtracting 1 from $g^*(p_j)$ for all $j > i$.

These insights allow us to reduce the frequency assignment problem to a problem instance of HCP as follows: let $G = (V, E)$ be a graph consisting of n nodes (one node for every radio transmitter). Draw an (undirected) edge between every pair of nodes for which the corresponding radio transmitters are not physically close to each other. The goal is now to find a graph $G' = (V, E \cup E')$ such that G' contains a Hamiltonian path and such that the amount of edges in E' is as small as possible. Let $(x_1, x_2, \dots, x_n)$ be the n subsequent nodes in some Hamiltonian path of G' . Now the frequency zones can be assigned as follows: $g^*(x_1) = 1$. If there is an edge $e \in E$ between x_i and x_{i+1} , we define $g^*(x_{i+1}) = g^*(x_i) + 1$ and otherwise we define $g^*(x_{i+1}) = g^*(x_i) + 2$. Note that the amount of edges in E' is as small as possible, such that the amount of consecutively assigned frequency zones is as large as possible and thus Equation (1) is indeed minimised. This new formulation is already very similar to HCP, but note that for HCP one is interested in obtaining a graph that contains a Hamiltonian cycle instead of a Hamiltonian path. This can be solved quite simply: let H be the graph that one obtains by adding a

node s to G and drawing an edge between s and all other nodes. Let H' be the Hamiltonian graph that one obtains by solving HCP for H . The problem of assigning frequencies to radio transmitters can now indeed be reformulated as a problem instance of HCP, because there is a bijection between Hamiltonian cycles in H' and Hamiltonian paths in G' . A Hamiltonian path in G' can be used to obtain a Hamiltonian cycle in H' by adding two (undirected) edges from both ends of the path to s and vice-versa.”

HCP can also be used to solve a special case of the Travelling Salesman Problem (TSP) [4]. A problem instance of TSP on a complete edge-weighted graph G with only two different edge weights (w_1 and w_2 , with $w_1 < w_2$) can be solved by solving HCP, as discussed in [27]. Since there are only two different edge weights in G , a least weight Hamiltonian cycle in G is equivalent with a Hamiltonian cycle that contains edges with weight w_2 as few times as possible (over all Hamiltonian cycles). This can in turn be found by solving HCP on a graph G' that is obtained by removing all the edges with weight w_2 from G . This is the case, because for HCP one is interested in finding the smallest possible amount of edges to add to G' to obtain a Hamiltonian graph and these added edges correspond precisely to edges with weight w_2 in G .

Yet another application comes from a generalisation of the Bottleneck Travelling Salesman Problem [16]. In this generalisation, we are given an integer k and an edge-weighted undirected graph $G = (V, E)$. The goal is to find a set of at most k paths in G such that every node in V belongs to exactly one path and such that the maximum edge weight amongst all edges in the paths is minimised. We assume k is chosen such that at least one solution exists. This problem can be solved by doing a binary search for the smallest maximum edge weight. To determine for a specific weight w whether a solution exists such that the maximum edge weight is at most w , one has to solve an instance of HCP. More specifically, let $f(G, w)$ denote an unweighted, undirected graph obtained by removing from G those edges that have a larger weight than w . The weights of the remaining edges are ignored to make the graph unweighted. The set of nodes in $f(G, w)$ is the same as the set of nodes in G . Let E' be the set of edges that is obtained by solving HCP on $f(G, w)$ (i.e. adding the edges in E' to $f(G, w)$ results in a Hamiltonian graph). Now, there exists a solution with maximum edge weight at most w if and only if E' contains at most k edges. Let w^* be the smallest value for which E' contains at most k edges and let H be the graph obtained by adding the edges in E' to $f(G, w^*)$ (so H is Hamiltonian). Now the set of at most k paths in G (the answer to the problem) can be found by taking any Hamiltonian cycle in H and dropping those edges that are in E' . The required set of paths is simply the obtained set of connected components. All edges in the obtained paths will have a weight that is at most equal to w^* .

4 Literature overview

It is clear that HCP on undirected graphs is NP-hard, because it can be used to determine whether a graph has a Hamiltonian cycle. Determining whether a graph has a Hamiltonian cycle, known as the Hamiltonian cycle problem, is NP-complete [15]. This means that HCP on undirected graphs cannot be solved exactly in polynomial time, unless $P = NP$. Hence, most of the existing literature concentrates on efficiently solving HCP for special cases of undirected graphs. Amongst those special

cases are algorithms that solve HCP in polynomial time for trees and unicyclic graphs [17] [18] [13], line graphs of trees [2] [26], line graphs of cacti [10] and series-parallel graphs [28]. HCP has also been studied in the context of sparse random graphs [14]. The amount of literature that concentrates on heuristically solving HCP is more limited [27] [11] [24].

In the remainder of this section, we will introduce some definitions and results from existing literature that will be needed to understand the rest of the paper.

4.1 Reduction to Travelling Salesman Problem

Lemma 1 [11] states that HCP on undirected graphs can be solved by solving an instance of the symmetric travelling salesman problem. Let $G = (V, E)$ be an undirected, unweighted graph. Let $G' = (V', E')$ be an undirected, weighted, complete graph such that $V = V'$ and E' consists of $\frac{|V|*(|V|-1)}{2}$ edges linking all pairs of nodes in V' . The weight of an edge $e \in E'$ is defined to be equal to 0 if and only if there is also an edge in G that links the same pair of nodes and otherwise the weight is equal to 1.

Lemma 1 *HCP on G can be solved by determining the minimum weight Hamiltonian cycle in G' and adding to G those edges of this cycle that have weight 1. Determining the minimum weight Hamiltonian cycle in G' is an instance of the symmetric travelling salesman problem.*

Using this lemma, HCP can be solved by reducing it to an instance of the symmetric TSP problem and then using state-of-the-art TSP solvers like Concorde [3] (an exact solver) or LKH [19] (a heuristic solver). When the graphs for which we want to solve HCP become bigger however, the computation time needed to produce high-quality solutions for TSP with these solvers might no longer be acceptable.

4.2 Minimum path partition

In this subsection the minimum path partition problem is introduced. This problem is closely related to HCP and will be an important concept needed to understand the algorithm in Section 5. First, some terminology is introduced.

Definition 1 *Two paths P_1 and P_2 are called vertex-disjoint if they do not have a vertex in common.*

Definition 2 *A path partition of a graph $G = (V, E)$ is a set of pairwise vertex-disjoint paths $\{P_1, P_2, \dots, P_k\}$ such that every node in V belongs to exactly 1 path P_i .*

Definition 3 *A minimum path partition of graph $G = (V, E)$ is a path partition of G , which consists of the least amount of paths over all path partitions of G .*

Definition 4 *The path partition number of a graph G (denoted $PPN(G)$) is defined as the amount of paths in a minimum path partition of G .*

Definition 5 The Hamiltonian completion number of a graph G (denoted $HCN(G)$) is defined as the minimum amount of edges that need to be added to G to obtain a Hamiltonian graph.

The minimum path partition problem on undirected graphs asks to find a minimum path partition on a given undirected graph. The following lemma [17] shows that this problem is closely related to HCP on undirected graphs:

Lemma 2 Let $G = (V, E)$ be an undirected graph. If G is Hamiltonian, then $HCN(G) = 0$ and $PPN(G) = 1$. Else, $HCN(G) = PPN(G)$.

In case $PPN(G) > 1$, a minimum path partition of G can be used to determine which edges to add to G to make it Hamiltonian (by cyclically linking the endpoints of the paths of the path partition). This is illustrated in Fig. 2. The graph on the left of the figure is the graph G for which HCP needs to be solved. The minimum path partition is indicated on this graph. Nodes with an equal number are part of the same path. The edges in green indicate the links between nodes that belong to the same path. Hence, the path partition number of the shown graph is equal to 3. The graph on the right of the figure illustrates that a Hamiltonian graph can be obtained by cyclically linking the endpoints of the paths of the path partition. The paths of the path partition are shown in green. The added edges are shown in red. A Hamiltonian cycle in this graph is indicated by the red and green edges.

Figure 2: Graph G and a Hamiltonian cycle in graph G'

In case $PPN(G) = 1$, either $HCN(G) = 0$ or $HCN(G) = 1$, but there is no simple rule to deduct from a given minimum path partition which of these two equalities is the correct one. Note that $HCN(G) = 0$ implies that $PPN(G) = 1$ and that there exists at least one minimum path partition of G such that there is an edge between the endpoints of the single path in that minimum path partition. However, not necessarily all minimum path partitions of G have this property. Hence, one cannot use the (non-)existence of an edge between the endpoints of a path to determine whether $HCN(G) = 0$ or $HCN(G) = 1$.

An immediate consequence of Lemma 2 and the corresponding construction is the following lemma:

Lemma 3 *Let $G = (V, E)$ be an undirected graph that consists of multiple components (i.e. G is a disconnected graph). HCP for G can be solved by computing a minimum path partition for every component of G separately.*

Using this lemma, it is clear that HCP for a graph G can be solved by solving HCP for every component of G separately. Hence, in the rest of this paper we will assume that the graph G for which HCP has to be solved, is a connected graph, unless stated otherwise.

4.3 Special properties for trees

If we restrict the input graph G , for which HCP has to be solved, to be a tree, we can efficiently solve HCP [18] [13]:

Lemma 4 *Let $G = (V, E)$ be a tree. It is possible to compute a minimum path partition of G with a time complexity of $O(|V|)$.*

The minimum path partition can then be used to solve HCP, as seen in the previous subsection. The algorithm to compute the minimum path partition of a tree is not further discussed in this paper, but interested readers are referred to [18] and [13].

Furthermore, there is a relation between the PPN of a graph and the PPN of its spanning trees [17]:

Lemma 5 *Let $G = (V, E)$ be an undirected graph. For all spanning trees T of G , the following inequality holds: $PPN(T) \geq PPN(G)$. Furthermore there is at least one spanning tree T of G such that $PPN(T) = PPN(G)$.*

By combining Lemma 4 and 5, it is clear that finding a minimum path partition for a graph can be done by finding a minimum path partition for all of its spanning trees. Unfortunately, this approach is not always feasible, because the amount of spanning trees of a graph can be very large (e.g. a complete graph with n vertices has n^{n-2} spanning trees [9]). However, note that the fact that a graph has many spanning trees does not necessarily imply that it is hard to find a minimum path partition for this graph. As can be seen in the results section, the algorithm that is proposed in this paper often is able to find the minimum path partition for graphs that have many spanning trees.

5 Multi-start local search algorithm

The algorithm that we propose to solve HCP on an undirected graph G will do so by attempting to find a minimum path partition of G (see Lemma 2 for the relationship between HCP and minimum path partition). It does this by trying to search for an optimal spanning tree T of G (i.e. such that $PPN(T) = PPN(G)$, see Lemma 5). Note that any path partition of a spanning tree T of G is also a path partition of G itself, because T and G have the same set of nodes. The algorithm uses the multi-start local search paradigm: multiple initial spanning trees will be perturbed in order to attempt to decrease

the path partition number of the resulting spanning trees. The different algorithm components are first explained in Subsection 5.1 and 5.2, after which the complete algorithm is explained in Subsection 5.3.

5.1 Initial spanning tree

Several steps are taken to obtain an initial spanning tree of the graph $G = (V, E)$:

- Step 1: HCP is first transformed into a symmetric TSP instance, as discussed in Section 4.1. A heuristic solution for this symmetric TSP instance is then computed by the heuristic TSP solver LKH in order to obtain an approximation for the minimum weight Hamiltonian cycle in the complete graph G' (where G' is the graph as described before Lemma 1). In order to keep the running time low, the parameters that are supplied to LKH in this step are purposely chosen in such a way that only a fraction of the full power of the LKH solver is being used (e.g. the supplied parameter *RUNS* is much smaller than the default amount of runs, see Section 7.1).
- Step 2: This Hamiltonian cycle in G' is used to obtain a path partition of G (by omitting those edges of the cycle that have weight 1).
- Step 3: This path partition is modified to obtain another path partition, by applying a rotation move [25] to all of its paths (in case it is possible to apply such a move). A rotation move modifies a path in the following way: let $P = (v_1, v_2, \dots, v_k)$ be a path consisting of k vertices. If there is a vertex v_i (with $1 < i < k - 1$) such that there is an edge between v_i and v_k in G , a rotation move can be applied to obtain the path $P' = (v_1, v_2, \dots, v_i, v_k, v_{k-1}, v_{k-2}, \dots, v_{i+1})$. If there are multiple choices for v_i , an arbitrary one is chosen. Fig. 3 illustrates the situation before step 3 is executed (on the left) and afterwards (on the right). A rotation move is applied for the first path (indicated with number 1).
- Step 4: The obtained path partition is transformed into a spanning tree of G by adding a set of edges of G , that do not already belong to the paths in the path partition. If the path partition consists of k paths, $k - 1$ edges have to be added in order to obtain a spanning tree. Not all edges are added with equal probability. The edges are grouped in two categories: preferred edges and normal edges. Preferred edges are edges that have at least one endpoint that is also an endpoint of some path in the given path partition. The other edges are called normal edges. The edges are added one by one. Every time an edge needs to be added, a decision is made whether to add an edge from the category of preferred edges or normal edges. If one of the categories does not contain any edges, a random edge from the other category is selected. Else, a random edge is chosen from the category of preferred edges with probability $\frac{\text{preferredRatio}}{\text{preferredRatio}+1}$ (and from the category of normal edges with probability $1 - \frac{\text{preferredRatio}}{\text{preferredRatio}+1} = \frac{1}{\text{preferredRatio}+1}$). By defining the probabilities in this way, we have the property that the parameter *preferredRatio* indicates the ratio between the probability that an edge is selected from the category of preferred edges and the probability that an edge is selected from the

Figure 3: The paths before and after applying rotation moves. The edges in green indicate the links between neighbouring nodes in a path.

category of normal edges. This parameter can be used to indicate how many times more likely it is to select an edge from the category of preferred edges than an edge from the category of normal edges. The spanning tree that is obtained, tends to have a lower path partition number when more edges from the category of preferred edges are added. Hence, the intended use of this parameter is to choose *preferredRatio* such that it is larger than 1, but this is not explicitly necessary.

Algorithm 1 illustrates the pseudocode to obtain an initial spanning tree. Here, *obtainInitialSpanningTree* is a function that accepts as input a graph G and produces as output a spanning tree of G . The function *solveTSP* is a function that accepts as input a graph and produces as output an approximation of the minimum weight Hamiltonian cycle in that graph. The function *makeTree* executes step 3 and step 4 from above. It accepts as input a graph G and a path partition of G . It produces as output a spanning tree of G . The pseudocode of this function is given in algorithm 2. The function *applyRotationMove* accepts as input a path p and returns as output the path that is obtained by applying a rotation move on p . Finally, the function *addEdges* corresponds to step 4 from above. It accepts as input a graph G and a path partition of G . It produces as output a spanning tree of G .

To analyse the time complexity of obtaining an initial spanning tree, we will analyse the time complexity of the different steps separately. The time complexity of step 1 heavily depends on the parameters that are supplied to LKH. The parameters for LKH can be chosen such that the preprocessing step of this solver runs in nearly linear time (by choosing the *POPMUSIC* option for candidate set generation [20]) and the amount of local search steps is constant. For a more thorough treatment

Algorithm 1: Obtain an initial spanning tree of the graph

```
Function obtainInitialSpanningTree(G)  
  Data:  
  G: an undirected graph  
  Result:  
  A spanning tree of G  
  /* Start of code */  
  Let  $G' = (V', E')$  be a complete graph such that  $V = V'$  and the weight of an edge  $e \in E'$   
  is defined to be equal to 0 if there is also an edge in G that links the same pair of nodes  
  and 1 otherwise;  
  /* step 1 */  
   $\text{minWeightHamiltonianCycle} \leftarrow \text{solveTSP}(G')$ ;  
  /* step 2 */  
   $\text{pathPartition} \leftarrow \emptyset$ ;  
  for edge in  $\text{minWeightHamiltonianCycle}$  :  
    if  $\text{weight}(\text{edge}) = 0$  then  
       $\text{pathPartition} \leftarrow \text{pathPartition} \cup \text{edge}$ ;  
  /* step 3 and 4 */  
  return  $\text{makeTree}(G, \text{pathPartition})$ 
```

Algorithm 2: Transform a path partition into a spanning tree

```
Function makeTree(G,  $\text{pathPartition}$ ) /* Step 3 and 4 */  
  Data:  
  G: an undirected graph  
   $\text{pathPartition}$ : a path partition of G  
  Result:  
  A spanning tree of G  
  /* Start of code */  
   $\text{newPathPartition} \leftarrow \text{pathPartition}$ ;  
  for path in  $\text{newPathPartition}$  :  
    if rotation move can be applied on path then  
       $\text{path} \leftarrow \text{applyRotationMove}(\text{path})$ ;  
  return  $\text{addEdges}(G, \text{newPathPartition})$ 
```

of the time complexity of this solver, the reader is referred to [19]. It is clear that step 2 and step 3 can each be implemented in $O(|V| + |E|)$. In order to efficiently execute step 4, one has to be able to quickly verify whether adding a certain edge would result in a cycle (if this is the case, the edge should not be added). This is a well-known problem that can be efficiently solved by using the union-find data

structure [29]. By using this data structure, one can verify in $O(\alpha(|V|))$ time whether adding an edge would result in a cycle. Here, $\alpha(|V|)$ denotes the inverse of the Ackermann function [1]. Because of this, step 4 can be implemented in $O(|E| * \alpha(|V|))$. Hence, the total time complexity to obtain an initial spanning tree is $LKHTime(V, E) + O(|V| + |E| * \alpha(|V|))$, where $LKHTime(V, E)$ denotes the time that the LKH solver spends in step 1.

5.2 Perturbing the spanning tree

The spanning tree that is obtained, is repeatedly perturbed in order to attempt to decrease its path partition number. In this subsection, we will denote by G the graph for which HCP needs to be solved and by T a spanning tree of G that will be perturbed. We associate with every edge of G one of two possible states: an edge can be active or inactive. Initially, all edges in G are inactive, except for the edges in T (which are active). The effect of the perturbation is that the state of some edges will change (edges go from inactive to active or vice-versa). After performing the perturbation, the set of active edges forms a (new) spanning tree of G . A single perturbation consists of multiple steps, which are explained below.

Step 1: Compute a minimum path partition of T . Mark all the edges that are part of T , but not part of the path partition as inactive. Note that any path partition of T is also a path partition of G .

Step 2: The main idea of this step is that we try to alter the path partition from the previous step to obtain a new path partition of G that consists of fewer paths. One possible way to approach this problem would be to try to link together endpoints of the paths from the path partition by marking the connecting edges as active. However, this idea can be generalised, without resulting in a higher time complexity: let $PP = \{P_1, P_2, \dots, P_k\}$ be the path partition obtained from step 1, which consists of k paths. For all of the paths in PP , it is verified whether the path can be transformed into a cycle by linking together the two endpoints of the path (i.e. it is verified whether there exists an edge in G such that the endpoints of this edge are equal to the endpoints of the path). In case the path can indeed be transformed into a cycle, this transformation is performed (the edge that is needed for this transformation is marked as active). After doing this, we have a set of paths S_1 (of size k_1) and a set of cycles S_2 (of size k_2) such that $k_1 + k_2 = k$.

These two sets will be iteratively updated by joining either a path with another path, a cycle with a path or a cycle with another cycle. To do this, the algorithm goes over all edges in G and verifies whether the current edge can be used to join the structures (paths or cycles) at both endpoints of the edge. In case this is possible, the structures are joined and the sets S_1 and S_2 are updated. The three cases in which an edge can be used to join two structures are discussed next.

Case 1: The edge is between two paths P_i and P_j . The edge can be used to join P_i and P_j if and only if the edge is between an endpoint of P_i and an endpoint of P_j (and the paths P_i

and P_j are not the same paths). If P_i and P_j can indeed be joined, the edge is marked as active, resulting in a longer path P' . The set S_1 is updated by removing P_i and P_j and adding P' .

Case 2: The edge is between a path P_i and a cycle C_j . The edge can be used to join P_i and C_j if and only if exactly one endpoint of the edge corresponds to an endpoint of P_i . If P_i and C_j can indeed be joined, the edge is marked as active. Furthermore, an edge of C_j is marked as inactive such that the joined structure forms a path. There are two possible such edges that could be marked as inactive. An arbitrary one of these two edges is chosen. By joining P_i and C_j a new path P' is obtained. The set S_1 is updated by removing P_i and adding P' . The set S_2 is updated by removing C_j .

Fig. 4 illustrates Case 2. The graph at the top of the figure consists of the path $1 - 2 - 3$ and the cycle $4 - 5 - 6 - 7 - 4$. The green and blue edges represent the links between consecutive nodes of the path and the cycle respectively. The black edge represents an inactive edge. There is an edge between the cycle and an endpoint of the path: the edge that links node 3 and node 4. This edge is used to join the path and the cycle and will be marked as active. There are two choices for an edge that could be marked as inactive in order to obtain a new path, namely the edge between node 4 and node 5 and the edge between node 4 and node 7. In this figure, the edge between node 4 and node 5 is marked as inactive in order to obtain the path $1 - 2 - 3 - 4 - 7 - 6 - 5$. The resulting path is shown at the bottom of the figure (the links between neighbouring nodes in the path are shown in green).

Case 3: The edge is between a cycle C_i and a cycle C_j . The edge can be used to join the cycles if and only if C_i is not the same cycle as C_j . In order to obtain a path by joining C_i and C_j , the current edge is marked as active and furthermore one edge from C_i and one edge from C_j are marked as inactive. Again (as in case 2), there are two possible edges that can be marked as inactive for a given cycle and an arbitrary one of these two is chosen for each cycle. The set S_1 is updated by adding the path P' that results from joining the cycles. The set S_2 is updated by removing C_i and C_j .

After having gone through all edges of G and joining the structures at the endpoints of the edges in case this is possible, we have obtained a new set of paths $S'_1 = \{P_1, P_2, \dots, P_{k'_1}\}$ and a new set of cycles $S'_2 = \{C_1, C_2, \dots, C_{k'_2}\}$ such that $k'_1 + k'_2 \leq k$. Finally, for all cycles in S_2 an arbitrary edge is chosen to mark as inactive (to transform the cycle back into a path). The resulting set of active edges forms a (new) path partition of G .

Step 3: This step is identical to step 3 from Section 5.1. The path partition that is obtained from step 2 is transformed into another path partition by applying rotation moves to all of the paths of the path partition for which this is possible.

Step 4: This step is identical to step 4 from Section 5.1. The path partition from step 3 is transformed into a spanning tree T' of G by marking a particular set of edges as active.

Figure 4: A path and a cycle are joined to obtain a new path.

To analyse the time complexity of perturbing a spanning tree, we will analyse the time complexity of the different steps separately. Recall from Lemma 4 that step 1 can be implemented in $O(|V|)$. Step 2 can be implemented in $O(|V| + |E|)$, because the state of an edge will only be changed at most once (from active to inactive or vice-versa). As we have seen in Section 5.1, step 3 can be implemented in $O(|V| + |E|)$ and step 4 can be implemented in $O(|E| * \alpha(|V|))$. Hence, the total time complexity of perturbing a spanning tree is $O(|V| + |E| * \alpha(|V|))$.

5.3 Complete algorithm

Now that the building blocks have been explained, we are ready to discuss the complete algorithm. The algorithm has three parameters: *preferredRatio*, *maxAmountInitialSpanningTrees* and *maxAmountBadPerturbations*. The first parameter (*preferredRatio*) has already been discussed in step 4 of Section 5.1, whereas *maxAmountInitialSpanningTrees* indicates how many initial spanning trees have to be generated. The third parameter (*maxAmountBadPerturbations*) requires a more detailed discussion of a crucial property of the perturbation operator.

For every initial spanning tree T of G , a number of perturbations will be performed. Let $perturbed(T, G)$ denote the spanning tree that is obtained by perturbing the spanning tree T (by making certain edges in G active and others inactive). The perturbation that was discussed before has the property that $PPN(perturbed(T, G)) \leq PPN(T)$, so perturbing a spanning tree always yields another

spanning tree that is at least as good. This property holds, because for every step of the perturbation, the path partition number of the graph that consists of the set of active edges does not increase. This property is the motivation for introducing the concept of a bad perturbation of a spanning tree: a perturbation such that $PPN(perturbed(T, G)) = PPN(T)$. The parameter *maxAmountBadPerturbations* indicates how many bad perturbations are applied to a spanning tree, before we move on to the next spanning tree.

For every spanning tree T' , that is obtained after performing perturbations on an initial spanning tree T , a minimum path partition is computed. This path partition is used to estimate $HCN(G)$ (see Lemma 2). In case $PPN(T') > 1$, the algorithm estimates $HCN(G)$ to be equal to $PPN(T')$. In the special case where the path partition consists of a single path (i.e. $PPN(T') = 1$), it is verified whether there is an edge in G between the endpoints of this path. The algorithm estimates $HCN(G)$ as 1 in case such an edge does not exist and 0 otherwise.

The pseudocode of the complete algorithm is shown in algorithm 3. Here, *estimateHCN* is a function that accepts as input a graph G for which HCP has to be solved and produces as output an estimate for $HCN(G)$. The function *obtainInitialSpanningTree* has been discussed in Section 5.1. The function *computeMinimumPathPartitionOfTree* accepts as input a tree and produces as output a minimum path partition of that tree. The function *perturbed* accepts as input a spanning tree of G as well as the graph G itself and produces as output a (new) spanning tree of G (see Section 5.2). Finally, the function *estimateHCNPathPartition* accepts as input a path partition of G as well as the graph G itself and produces as output an estimate of $HCN(G)$.

To analyse the worst-case time complexity of the complete algorithm, first note that the path partition number of any graph $G = (V, E)$ is at most equal to $|V|$, because of the existence of a trivial path partition in which every path consists of a single node. In every iteration of the while loop either the size of the path partition decreases by at least one or the size of the path partition remains the same, in which case the amount of bad perturbations is incremented. If we combine this insight with the time complexities from the previous subsections, we obtain for the complete algorithm a worst-case time complexity of $O(maxAmountInitialSpanningTrees) * (LKHTime(V, E) + (|V| + maxAmountBadPerturbations) * (|V| + |E| * \alpha(|V|)))$. If the parameters are regarded as constants, the time complexity becomes $LKHTime(V, E) + O(|V| * |E| * \alpha(|V|))$.

6 Benchmark set generation

To the best of our knowledge, up until now there was no standard benchmark set of problem instances already available for HCP on undirected graphs. In order to be able to test the quality of the proposed algorithm, we have constructed a benchmark set that consists of a variety of problem instances. A problem instance for HCP on undirected graphs is simply an undirected graph for which HCP has to be solved.

We have constructed a benchmark set of problem instances that vary with respect to size, density, degree distribution and several other features. It is composed both of graphs that arise from practice as well as graphs that were generated according to theoretical models. For the former case, a number

Algorithm 3: Estimate the HCN of a graph

```
Function estimateHCN( $G$ )
  Data:
   $G$ : an undirected graph
  Result:
  An estimate for  $HCN(G)$ 
  /* Start of code */
   $estimate \leftarrow \infty$ ;
  for  $times$  in  $range(maxAmountInitialSpanningTrees)$  :
     $T \leftarrow obtainInitialSpanningTree(G)$ ;
     $amountBadPerturbations \leftarrow 0$ ;
     $pathPartition \leftarrow computeMinimumPathPartitionOfTree(T)$ ;
    while  $amountBadPerturbations \leq maxAmountBadPerturbations$  :
       $T \leftarrow perturbed(T, G)$ ;
       $newPathPartition \leftarrow computeMinimumPathPartitionOfTree(T)$ ;
      if  $pathPartition$  and  $newPathPartition$  have the same amount of paths then
         $amountBadPerturbations \leftarrow amountBadPerturbations + 1$ ;
       $pathPartition \leftarrow newPathPartition$ ;
       $newEstimate \leftarrow estimateHCNPathPartition(pathPartition, G)$ ;
      if  $newEstimate < estimate$  then
         $estimate \leftarrow newEstimate$ ;
      if  $estimate = 0$  then /* Guaranteed global optimum */
        return 0;
  return  $estimate$ ;
```

of graphs were selected from the instances of the tenth DIMACS challenge about graph partitioning and graph clustering [5]. Some of these graphs are directed graphs. These graphs are transformed into undirected graphs (by ignoring the direction of the edges) to obtain the problem instances. The other problem instances were generated by using several graph generators. Some of these generators are part of the library of the Stanford Network Analysis Platform (SNAP) [23]. This library contains many graph generators that have an underlying theoretical model. A brief overview of the six generators that were used to devise the benchmark set follows.

Generator 1: This generator follows the Erdős-Rényi model [12]. This graph generator is a stochastic generator and has two parameters n and p . The generated graph consists of n nodes and for every pair of nodes there is an edge between them with probability p . Hence, each (simple, undirected) graph that consists of n nodes and m edges is equally likely to be generated by this model (with probability $p^m * (1 - p)^{\frac{n*(n-1)}{2} - m}$).

The problem instances that were generated by this generator contain a varying amount of nodes (powers of 2 between 2^8 and 2^{13}). For a fixed amount of nodes n , several problem instances were generated with varying probabilities (such that the expected average degree of a node varies between $2^0, 2^1, 2^2, \dots, n$).

For graphs generated with this model, many theoretical results are known. We highlight two relevant theoretical results from [7]:

Theorem 1 *Let $0 < r < 1$ be a real number and let $G(n, p) = (V, E)$ be a graph obtained by the Erdős-Rényi random model with parameters n and p . If n tends to infinity and p is chosen such that $|E| = o(|V|^r)$, the probability that all components of the graph are trees, tends to 1.*

Because of the previous theorem, for sufficiently small values of p , the probability that the algorithm proposed in this paper computes the global optimum of HCP for Erdős-Rényi graphs tends to 1 if n tends to infinity. This is the case because for sufficiently small values of p , the probability that all components of the graph are trees tends to 1 if n tends to infinity and the proposed algorithm solves HCP exactly for trees (see Lemma 4).

The following theorem indicates for which graphs the Hamiltonian completion number is likely to be equal to 0:

Theorem 2 *Let $f(n)$ be a function that tends to infinity if n tends to infinity and let $G(n, p) = (V, E)$ be a graph obtained by the Erdős-Rényi random model with parameters n and p . Set $p_l = \frac{\log(n) + \log(\log(n)) - f(n)}{n}$ and $p_u = \frac{\log(n) + \log(\log(n)) + f(n)}{n}$. If n tends to infinity, the probability that $G(n, p_l)$ contains a Hamiltonian cycle tends to 0 and the probability that $G(n, p_u)$ contains a Hamiltonian cycle tends to 1.*

Hence, for sufficiently large values of p the generated graphs are very likely to have a Hamiltonian completion number equal to 0.

Generator 2: This generator generates circular graphs and has two parameters n and k . The generated graph consists of n nodes, labelled $0, 1, \dots, n-1$. For every node with label i , there is an edge to the k subsequent nodes $(i+1) \bmod n, (i+2) \bmod n, \dots, (i+k) \bmod n$. For all of these graphs, the Hamiltonian completion number is equal to 0.

The parameters that were used to generate problem instances by this generator are varying: n is between 100 and 30000 and k is between 2 and 10.

Generator 3: This generator generates two-dimensional grid graphs and has two parameters n and m . The generated graph consists of $n * m$ nodes, which can be thought of as n rows of m nodes which are laid out in a two-dimensional grid. There is an edge between a node in row r_1 and column c_1 and a node in row r_2 and column c_2 if and only if

$|r_1 - r_2| + |c_1 - c_2| = 1$. For these graphs, the Hamiltonian completion number is equal to 0 if either n or m are even and otherwise equal to 1.

The problem instances that were generated by this generator contain between 100 and 10000 nodes.

Generator 4: This generator generates preferential attachment graphs, following the Barabási-Albert model [6]. In this model, the degree distribution of the nodes follows a power-law: the probability that a node is connected to k other nodes is proportional with k^{-3} .

The graphs that were generated according to this model contain between 100 and 3000 nodes and the nodes have an average degree varying between 6 and 18.

Generator 5: This generator starts from a star graph consisting of $n + 1$ nodes and adds n random edges to obtain a problem instance. A star graph consisting of $n + 1$ nodes is the complete bipartite graph $K_{1,n}$ (a tree with 1 internal node and n leaves).

The problem instances that were generated by this generator contain between 100 and 4000 nodes.

Generator 6: The final generator generates structured trees and has two parameters l and c . The generated graphs are trees that consist of l levels of nodes and every node (except for the leaf nodes) has c child nodes. Note that for these graphs, the algorithm proposed in the paper is guaranteed to compute the global optimum, because our algorithm solves HCP exactly for trees.

The parameters that were used to generate problem instances according to this generator has values of c varying between 2 and 9, while l was chosen such that the total amount of nodes varies between 7 and 9841.

7 Results and experiments

In total, the benchmark set consists of 285 problem instances. This benchmark set is split into two parts, consisting of 40% and the remaining 60% of all problem instances respectively: a tuning set (consisting of 113 randomly chosen problem instances) and a test set (consisting of the remaining 172 problem instances). Table 1 shows the distribution of the problem instances over each of the two sets. The tuning set will be used to find appropriate parameter settings for the algorithm proposed in this paper. Furthermore, we will demonstrate the effect of varying these parameters and how this influences the performance of the algorithm. The test set will be used to show the added value of the algorithm components that were discussed in this paper. Furthermore, we compare our algorithm with the performance of state-of-the-art algorithms and we attribute special attention to when our algorithm performs well and when it does not. For the sake of completeness we will also compare the performance of the algorithms on the tuning set and demonstrate that we obtain similar results for the tuning and test set. From these experiments it becomes clear that the algorithm proposed in this

paper is the best option for the hardest problem instances. All experiments were performed on the ThinKing cluster of the Flemish Supercomputer Center (VSC), using CPU's with a clock rate of 2.5 GHz. The test data is available online at <https://set.kuleuven.be/codes/hamiltoniancompletionproblem> and at <https://github.com/JorikJooker/ProblemInstancesHamiltonianCompletionProblem>.

Table 1: Amount of problems in tuning set and test set

	Tuning set	Test set
DIMACS	6	9
Erdős-Rényi	27	40
Circular	16	29
Grid	14	25
Preferential attachment	16	26
Star	18	22
Structured tree	16	21

7.1 Tuning

As discussed before, the algorithm has three parameters: *preferredRatio*, *maxAmountInitialSpanningTrees* and *maxAmountBadPerturbations*. For the parameter *preferredRatio*, it is not immediately clear without any experiments which setting would lead to the best performance. For the other two parameters, however, it is clear that higher values correspond to a better performing algorithm, because they indicate the amount of local search that has to be done. Hence, for these two parameters we will try to find settings that yield a good trade-off between execution time and performance. Note that the LKH solver that is used in step 1 of the algorithm also has parameters itself. These parameters will not be subject to tuning, but instead they are chosen in such a way that only a fraction of the full power of the LKH solver is being used. The precise parameters that were supplied to this solver can be found in Table 2.

Table 2: Parameter settings for LKH for step 1 of constructing an initial spanning tree

Parameter	Value
<i>MOVE_TYPE</i>	5
<i>PATCHING_C</i>	3
<i>PATCHING_A</i>	2
<i>RUNS</i>	1
<i>CANDIDATE_SET_TYPE</i>	<i>POPMUSIC</i>
<i>MAX_TRIALS</i>	2

In the following sections, the multi-start local search algorithm from this paper will be used with various parameter settings. For the sake of notational convenience, we will abbreviate the multi-start

local search algorithm with a specific parameter setting as *MSLS*_{*a*_*b*_*c*}, where *a*, *b* and *c* correspond to the chosen values for *preferredRatio*, *maxAmountInitialSpanningTrees* and *maxAmountBadPerturbations* respectively. For every problem instance, the algorithm is interrupted after 1000 seconds if the algorithm has not finished within this time. If the algorithm finishes in less than 1000 seconds, no further search is performed to fill the remaining time. In order to obtain fair comparisons, in all of the experiments the reported averages will be computed only over those problem instances for which all algorithms (in the scope of the specific experiment) were able to produce an answer within 1000 seconds. Hence, the reported averages can vary significantly between different experiments (depending on which problem instances the averages are computed over). This makes the reported averages only comparable within one specific experiment. We will also report about the problem instances for which an algorithm did not produce an answer within 1000 seconds.

To do the tuning, we have used a full factorial design of experiments, because the amount of parameters of the algorithm is low. The values that were used for the parameters are as follows: *preferredRatio* could take the values 1, 5, 25 and 125, *maxAmountInitialSpanningTrees* could take the values 1, 10 and 100 and finally *maxAmountBadPerturbations* could take the values 100, 300, 1000 and 3000. This results in a total of $4 * 3 * 4 = 48$ combinations of parameter settings. The summary of the results of this experiment for every individual parameter setting can be found in Table 3. The first column of the table indicates the parameter settings that were used. The second column of the table indicates on how many problem instances the algorithm did not finish within 1000 seconds, whereas the third column of the table shows the average amount of time needed to solve a problem instance. The fourth column displays the average amount of edges that the algorithm has added in order to obtain a Hamiltonian graph (the smaller the better). In addition, we show in Table 4 the averaged results for every specific parameter value. From these tables we can see that, by varying the parameters, the average amount of added edges shows variations of several percentages. However, the influence on the execution time and the amount of problem instances for which the algorithm was not finished within 1000 seconds is quite large. There are four parameter settings (*MSLS*_{1_100_3000}, *MSLS*_{5_100_3000}, *MSLS*_{25_100_3000} and *MSLS*_{125_100_3000}) for which the average amount of added edges is the lowest (67.59) amongst all parameter settings. However, for these four parameter settings there were 19 problem instances for which the algorithm did not finish within 1000 seconds and the average execution time is relatively high. There are four other parameter settings (*MSLS*_{1_10_3000}, *MSLS*_{5_10_3000}, *MSLS*_{25_10_3000} and *MSLS*_{125_10_3000}) for which the average amount of added edges is slightly higher (67.60), but the average execution time (and amount of interrupted instances) is much lower. For these parameter settings, the algorithm finished within 1000 seconds for all problem instances and the average execution time is lowest (16.58 seconds) for *MSLS*_{25_10_3000}. These experiments lead us to use the following parameter configuration for testing purposes: *preferredRatio* = 25, *maxAmountInitialSpanningTrees* = 10 and *maxAmountBadPerturbations* = 3000.

Table 3: Results of full factorial experiment on tuning set for every individual parameter setting

Parameter setting	Amount interrupted	Average time (in seconds)	Average amount edges added
<i>MSLS</i> _1_1_100	0	2.47	69.41
<i>MSLS</i> _1_1_300	0	2.68	68.40
<i>MSLS</i> _1_1_1000	0	3.56	67.89
<i>MSLS</i> _1_1_3000	0	4.99	67.69
<i>MSLS</i> _1_10_100	0	5.12	68.90
<i>MSLS</i> _1_10_300	0	7.99	68.05
<i>MSLS</i> _1_10_1000	1	11.54	67.73
<i>MSLS</i> _1_10_3000	0	17.18	67.60
<i>MSLS</i> _1_100_100	3	32.38	68.72
<i>MSLS</i> _1_100_300	5	45.16	67.95
<i>MSLS</i> _1_100_1000	9	55.57	67.61
<i>MSLS</i> _1_100_3000	19	107.14	67.59
<i>MSLS</i> _5_1_100	0	2.62	69.25
<i>MSLS</i> _5_1_300	0	2.77	68.34
<i>MSLS</i> _5_1_1000	0	3.72	67.87
<i>MSLS</i> _5_1_3000	0	4.90	67.69
<i>MSLS</i> _5_10_100	0	5.51	68.78
<i>MSLS</i> _5_10_300	0	6.88	68.07
<i>MSLS</i> _5_10_1000	1	10.01	67.72
<i>MSLS</i> _5_10_3000	0	16.60	67.60
<i>MSLS</i> _5_100_100	4	32.42	68.52
<i>MSLS</i> _5_100_300	5	44.12	67.83
<i>MSLS</i> _5_100_1000	9	62.82	67.64
<i>MSLS</i> _5_100_3000	19	98.89	67.59
<i>MSLS</i> _25_1_100	0	2.58	69.34
<i>MSLS</i> _25_1_300	0	3.11	68.39
<i>MSLS</i> _25_1_1000	0	3.54	67.89
<i>MSLS</i> _25_1_3000	0	4.86	67.67
<i>MSLS</i> _25_10_100	0	5.50	68.74
<i>MSLS</i> _25_10_300	0	7.16	68.03
<i>MSLS</i> _25_10_1000	1	10.27	67.70
<i>MSLS</i> _25_10_3000	0	16.58	67.60
<i>MSLS</i> _25_100_100	4	32.18	68.52
<i>MSLS</i> _25_100_300	5	43.88	67.84
<i>MSLS</i> _25_100_1000	7	57.02	67.61
<i>MSLS</i> _25_100_3000	19	97.50	67.59
<i>MSLS</i> _125_1_100	0	2.81	69.36
<i>MSLS</i> _125_1_300	0	2.98	68.37
<i>MSLS</i> _125_1_1000	0	3.80	67.84
<i>MSLS</i> _125_1_3000	0	4.92	67.72
<i>MSLS</i> _125_10_100	0	5.44	68.81
<i>MSLS</i> _125_10_300	0	6.85	68.00
<i>MSLS</i> _125_10_1000	1	10.17	67.73
<i>MSLS</i> _125_10_3000	0	16.61	67.61
<i>MSLS</i> _125_100_100	4	32.40	68.49
<i>MSLS</i> _125_100_300	5	43.77	67.85
<i>MSLS</i> _125_100_1000	9	54.91	67.62
<i>MSLS</i> _125_100_3000	19	94.00	67.59

Table 4: Results of full factorial experiment on tuning set, averaged for every specific parameter value

Parameter	Value	Average amount interrupted	Average time (in seconds)	Average amount edges added
<i>preferredRatio</i>	1	3.08	24.65	68.13
	5	3.17	24.27	68.08
	25	3.00	23.68	68.08
	125	3.17	23.22	68.08
<i>maxAmountInitialSpanningTrees</i>	1	0.00	3.52	68.32
	10	0.25	9.96	68.04
	100	9.06	58.39	67.91
<i>maxAmountBadPerturbations</i>	100	1.25	13.45	68.90
	300	1.67	18.11	68.09
	1000	3.17	23.91	67.74
	3000	6.33	40.35	67.63

7.2 Contribution of algorithm components

The algorithm proposed in this paper makes use of the TSP solver LKH (in step 1 of the process of obtaining an initial spanning tree). In this subsection, we show that the other algorithm components contribute significantly to the final result obtained by our algorithm. We make a comparison between the results obtained by on the one hand *MSLS_25_10_3000* and on the other hand an algorithm that only performs step 1 of the process of obtaining an initial spanning tree (i.e. *MSLS_25_10_3000* without all other steps). We will call the latter algorithm *MULTI_LKH*. The same parameter settings were used for LKH in both algorithms. These parameter settings can be found in Table 2.

The results of running these two algorithms on the problem instances of the test set can be found in Table 5. From this table, we can see that about 12.4% of the total computing time of *MSLS_25_10_3000* was spent on step 1 for obtaining initial spanning trees. By also performing the other steps, the amount of added edges could on average be reduced by approximately 6. Hence, the contribution of the other steps is significant. This is also confirmed by doing a paired samples t-test on the pairs obtained by pairing the results of both algorithms. The null hypothesis in this test states that the mean difference (the result of *MSLS_25_10_3000* minus the result of *MULTI_LKH*) is equal to 0 and the alternative hypothesis states that the mean difference is smaller than 0. A p-value of 0.000008 was obtained, such that at a significance level of 1%, we reject the null hypothesis in favor of the alternative hypothesis.

Table 5: Results on test set for *MSLS_25_10_3000* and for *MULTI_LKH*

	<i>MSLS_25_10_3000</i>	<i>MULTI_LKH</i>
Amount interrupted	1	1
Average time (in seconds)	59.82	7.39
Average amount edges added	241.50	247.73

7.3 Comparison with existing solvers

We will compare the results of our algorithm with the results obtained by the TSP solvers LKH and Concorde by regarding HCP as a specific TSP instance (see Subsection 4.1). Concorde is an exact solver and hence for this solver the most interesting aspect is its computation time. LKH on the other hand is a heuristic solver. For LKH, we have used the recommended parameter settings [19] in this experiment. Note that these parameter settings are much more time consuming than those used by LKH in step 1 of our own algorithm. The parameter settings used for LKH in this experiment can be found in Table 6.

Table 6: Recommended parameter settings for LKH

Parameter	Value
<i>MOVE_TYPE</i>	5
<i>PATCHING_C</i>	3
<i>PATCHING_A</i>	2
<i>RUNS</i>	10
<i>CANDIDATE_SET_TYPE</i>	<i>POPMUSIC</i>
<i>MAX_TRIALS</i>	$ V $

The results of running these three algorithms on the 172 problem instances of the test set are summarised in Table 7. All three algorithms were given a maximum computation time of 1000 seconds per test instance. There were 15 problem instances for which both Concorde and LKH were not able to compute an answer within 1000 seconds, while there was only 1 such problem instance for *MSLS_25_10_3000*. Concorde needed on average around 78.9% and 36.7% of the time needed by the multi-start local search algorithm proposed in this paper and LKH respectively (however, recall that the averages are only calculated over those problem instances for which all algorithms were able to finish within 1000 seconds, as indicated earlier). The average amount of edges added produced by the three algorithms are very close to each other (Concorde produced the best results, followed by *MSLS_25_10_3000*, followed by LKH). Recall that Concorde is an exact solver, so the result achieved by Concorde is the global optimum. From this comparison, we see that the other two algorithms are also often able to compute a solution that is (close to) the global optimum. For the sake of completeness, we also summarised the results of all algorithms on the 113 problem instances of the tuning set in Table 8. In terms of amount of interrupted problem instances and average amount of added edges, we get a very similar relationship between the three algorithms for both the tuning set and the test set. However, if we look at the average computation time for the tuning set, we get a somewhat different impression than for the test set. For the tuning set, *MSLS_25_10_3000* needed on average around 62.6% and 44.0% of the time needed by Concorde and LKH respectively, whereas for the test set Concorde was on average the fastest algorithm. This can be explained by the fact that this tuning set was explicitly used to pick the best performing parameters for the multi-start local search algorithm and hence we have to keep this mind while interpreting the results. However, by comparing the results

in Table 7 and 8 it becomes clear that the found parameters for the multi-start local search algorithm generalise quite well towards unseen problem instances from the same generators.

We can conclude that, somewhat surprisingly, the exact solver is on average faster than both heuristic solvers for the test set (at least for those instances for which all algorithms were able to produce an answer within 1000 seconds). Hence, the main application for the heuristic solvers is to solve those instances for which the execution time of Concorde becomes too high (i.e. the interrupted instances). We will now focus on these instances.

Table 7: Results for Concorde, *MSLS_25_10_3000* and LKH on the test set

	Concorde	<i>MSLS_25_10_3000</i>	LKH
Amount interrupted	15	1	15
Average time (in seconds)	32.97	41.80	89.75
Average amount edges added	245.84	245.90	246.11

Table 8: Results for Concorde, *MSLS_25_10_3000* and LKH on the tuning set

	Concorde	<i>MSLS_25_10_3000</i>	LKH
Amount interrupted	10	0	14
Average time (in seconds)	19.92	12.46	28.34
Average amount edges added	77.21	77.21	77.35

There were 140 problem instances (out of 172) for which all three algorithms were able to produce the global optimum. The results on the other problem instances are shown in Table 9. The values in the table are the amount of edges added by each algorithm, whereas the corresponding computation times are provided between brackets. *MSLS_25_10_3000* was able to compute an answer within 1000 seconds for almost all problem instances (171 out of 172), while both Concorde and LKH did not finish within this time limit for 15 problem instances (indicated by the dashes in the table). The results obtained by *MSLS_25_10_3000* were at least as good as those obtained by LKH for all problem instances except for two problem instances: a grid graph consisting of 2 rows and 3000 columns and a grid graph consisting of 2 rows and 5000 columns (*grid_graph_2_3000* and *grid_graph_2_5000* in Table 9). For both of these problem instances, *MSLS_25_10_3000* produced an answer of 2, while LKH produced the global optimum of 0.

On all problem instances for which Concorde was not able to compute an answer within 1000 seconds, the global optimum was equal to 0 (by construction, see Section 6). Amongst these instances, there were 4 circular graphs and 11 grid graphs. For the circular graphs, *MSLS_25_10_3000* was able to compute the global optimum of 0 for every problem instance, while LKH did not produce an answer within 1000 seconds. For the grid graphs, the deviations from the global optimum range

between 0 and 2 for *MSLS_25_10_3000* and between 0 and 11 for LKH. Amongst these 11 grid graphs, *MSLS_25_10_3000* produced the global optimum 7 times and LKH 3 times.

The problem instances *preferential_attachment_1000_4*, *preferential_attachment_2000_4* and *polblogs* deserve further attention. These first two problem instances are preferential attachment graphs consisting of 1000 and 2000 nodes respectively. The average degree of a node is 8 for both problem instances. The last problem instance (*polblogs*) is a DIMACS problem instance, representing a network of hyperlinks between political blogs. This graph consists of 1490 nodes and the average degree of a node is 22.44. These three problem instances are the only problem instances for which the results obtained by *MSLS_25_10_3000* and LKH differ significantly from the global optimum. The global optimum for *preferential_attachment_1000_4* is 0 and *MSLS_25_10_3000* and LKH produced an answer of 5 and 6 respectively. Similarly the global optimum is 0 for *preferential_attachment_2000_4*, while *MSLS_25_10_3000* and LKH produced an answer of 7 and 10 respectively. For the problem instance *polblogs*, the global optimum is 460, while *MSLS_25_10_3000* produced an answer of 467 and LKH was not able to compute an answer within 1000 seconds. We suspect that the reason why these particular instances are hard to solve for *MSLS_25_10_3000* and LKH is related to the degree distribution of the nodes. If the nodes in a graph all have a very low degree, the graph can be thought of as being quite similar to a tree, because the average degree of a node in a tree is approximately equal to 2. Hence, one could expect that for these graphs HCP can be solved (nearly) to optimality, because HCP can be solved exactly in linear time for trees. If, however, the nodes in a graph all have a very high degree, the graph can be expected to contain many Hamiltonian cycles, such that the Hamiltonian completion number of these graphs will most likely be close to 0. This implies that one could also expect that the HCP can be solved (nearly) to optimality for these graphs. The degree distribution for the nodes in the graphs of these three instances does, however, not belong to either of these two categories. For these graphs most nodes have a small degree while only a few nodes have a very large degree in comparison with the other nodes.

8 Conclusions and further work

In this paper, we have proposed a multi-start local search algorithm for HCP on undirected graphs. Furthermore, we have constructed a benchmark set of problem instances for this problem, which can hopefully help future researchers to be able to compare the performance of their algorithms against other algorithms. This benchmark set was used to demonstrate that the algorithm in the paper performs well in comparison with state-of-the-art TSP solvers, both in respect of time and quality of the produced answer.

An interesting idea that was not further explored in this paper is related to decomposing a problem instance into smaller problem instances. For HCP, the articulation points and bridges of the graph allow it to be decomposed, as discussed in [27]. It is, however, not immediately clear whether the overhead involved in decomposing the problem and reassembling the different parts of the solution is worthwhile.

Another interesting future research avenue is inspired upon the work that has been done for TSP.

Table 9: Results on problem instances for which at least one algorithm did not produce the global optimum.
A dash indicates that the algorithm was not able to find a solution within 1000 seconds.

Name problem instance	Concorde	<i>MSLS_25_10_3000</i>	LKH
<i>circle_like_2000_10</i>	0 (17.18 s)	0 (9.29 s)	—
<i>circle_like_10000_5</i>	0 (63.04 s)	0 (45.71 s)	—
<i>circle_like_10000_7</i>	0 (130.70 s)	0 (89.41 s)	—
<i>circle_like_10000_10</i>	0 (58.38 s)	0 (83.75 s)	—
<i>circle_like_20000_3</i>	0 (164.66 s)	0 (29.91 s)	—
<i>circle_like_20000_7</i>	0 (239.18 s)	0 (126.88 s)	—
<i>circle_like_20000_10</i>	0 (168.38 s)	0 (252.13 s)	—
<i>circle_like_30000_2</i>	—	0 (20.55 s)	—
<i>circle_like_30000_3</i>	—	0 (56.68 s)	—
<i>circle_like_30000_7</i>	—	0 (842.42 s)	—
<i>circle_like_30000_10</i>	—	—	—
<i>er_11_2</i>	130 (43.23 s)	130 (104.59 s)	133 (126.26 s)
<i>er_12_12</i>	0 (58.86 s)	0 (411.74 s)	—
<i>er_12_2</i>	242 (65.17 s)	243 (254.22 s)	256 (568.11 s)
<i>er_13_11</i>	0 (130.07 s)	0 (367.39 s)	—
<i>grid_graph_10_200</i>	—	0 (23.12 s)	0 (94.74 s)
<i>grid_graph_20_70</i>	—	0 (16.51 s)	1 (52.93 s)
<i>grid_graph_20_80</i>	0 (10.33 s)	0 (3.26 s)	2 (61.24 s)
<i>grid_graph_2_3000</i>	—	2 (150.86 s)	0 (9.20 s)
<i>grid_graph_2_5000</i>	—	2 (242.06 s)	0 (27.05 s)
<i>grid_graph_30_60</i>	—	0 (19.88 s)	1 (78.34 s)
<i>grid_graph_30_80</i>	—	0 (5.47 s)	4 (145.56 s)
<i>grid_graph_40_80</i>	—	0 (79.24 s)	7 (259.40 s)
<i>grid_graph_50_60</i>	0 (19.82 s)	0 (17.69 s)	5 (224.91 s)
<i>grid_graph_50_70</i>	—	0 (10.35 s)	7 (315.56 s)
<i>grid_graph_60_60</i>	—	1 (150.12 s)	7 (370.33 s)
<i>grid_graph_60_70</i>	—	0 (65.74 s)	9 (452.53 s)
<i>grid_graph_60_80</i>	—	1 (212.39 s)	11 (619.89 s)
<i>polblogs</i>	460 (61.52 s)	467 (305.82 s)	—
<i>preferential_attachment_1000_4</i>	0 (17.94 s)	5 (56.87 s)	6 (90.67 s)
<i>preferential_attachment_2000_4</i>	0 (33.49 s)	7 (132.46 s)	10 (400.34 s)
<i>structured_tree_8_3</i>	4921 (94.88 s)	4921 (509.26 s)	—

The LKH TSP solver first converts the input graph to a new graph in which some edges, that are unlikely to be part of an optimal TSP tour, are removed. By doing this, the search space can be drastically reduced and the solver can focus its search on more promising parts of the search space. The most common measure that is used to estimate how likely it is that a certain edge is part of an optimal TSP tour is called the α -nearness of an edge [19]. Note that for HCP the α -nearness of all edges are equal and hence this measure is not really useful for HCP. It would be interesting to see whether some similar measure could be used for HCP.

Acknowledgements

We gratefully acknowledge the support provided by the ORDinL project (FWO-SBO S007318N, Data Driven Logistics, 1/1/2018 - 31/12/2021). Pieter Leyman is a Postdoctoral Fellow of the Research Foundation - Flanders (FWO) with contract number 12P9419N.

This is a post-peer-review, pre-copyedit version of an article published in Journal of Heuristics. The final authenticated version is available online at: <https://doi.org/10.1007/s10732-020-09447-9>.

References

- [1] W. Ackermann. Zum hilbertschen aufbau der reellen zahlen. *Mathematische Annalen*, 99(1):118–133, Dec 1928.
- [2] A. Agnetis, P. Detti, C. Meloni, and D. Pacciarelli. A linear algorithm for the hamiltonian completion number of the line graph of a tree. *Information Processing Letters*, 79(1):17 – 24, 2001.
- [3] D. Applegate. Concorde - a code for solving traveling salesman problems. <http://www.math.princeton.edu/tsp/concorde.html>, 2001.
- [4] D. L. Applegate, R. E. Bixby, V. Chvatal, and W. J. Cook. *The Traveling Salesman Problem: A Computational Study (Princeton Series in Applied Mathematics)*. Princeton University Press, Princeton, NJ, USA, 2007.
- [5] D. A. Bader, H. Meyerhenke, P. Sanders, and D. Wagner, editors. *Graph Partitioning and Graph Clustering, 10th DIMACS Implementation Challenge Workshop, Georgia Institute of Technology, Atlanta, GA, USA, February 13-14, 2012. Proceedings*, volume 588 of *Contemporary Mathematics*. American Mathematical Society, 2013.
- [6] A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.
- [7] B. Bollobás. *Modern Graph Theory*, volume 184. 01 1998.
- [8] M. A. Boschetti, V. Maniezzo, M. Roffilli, and A. Bolufé Röhler. Matheuristics: Optimization, simulation and control. In M. J. Blesa, C. Blum, L. Di Gaspero, A. Roli, M. Sampels, and A. Schaerf, editors, *Hybrid Metaheuristics*, pages 171–177, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg.
- [9] A. Cayley. *The Collected Mathematical Papers*, volume 10 of *Cambridge Library Collection - Mathematics*. Cambridge University Press, 2009.
- [10] P. Detti and C. Meloni. A linear algorithm for the hamiltonian completion number of the line graph of a cactus. *Discrete Applied Mathematics*, 136(2):197 – 215, 2004. The 1st Cologne-Twente Workshop on Graphs and Combinatorial Optimization.
- [11] P. Detti, C. Meloni, and M. Pranzo. Local search algorithms for finding the hamiltonian completion number of line graphs. *Annals of Operations Research*, 156:5–24, 09 2007.
- [12] P. Erdős and A. Rényi. On random graphs, I. *Publicationes Mathematicae (Debrecen)*, 6:290–297, 1959.
- [13] D. S. Franzblau and A. Raychaudhuri. Optimal hamiltonian completions and path covers for trees, and a reduction to maximum flow. *The ANZIAM Journal*, 44(2):193204, 2002.

- [14] D. Gamarnik and M. Sviridenko. Hamiltonian completions of sparse random graphs. *Discrete Appl. Math.*, 152(1-3):139–158, Nov. 2005.
- [15] M. R. Garey and D. S. Johnson. *Computers and Intractability; A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., New York, NY, USA, 1990.
- [16] P. C. Gilmore and R. E. Gomory. Sequencing a one state-variable machine: A solvable case of the traveling salesman problem. *Operations Research*, 12(5):655–679, 1964.
- [17] S. Goodman and S. Hedetniemi. On the hamiltonian completion problem. In R. A. Bari and F. Harary, editors, *Graphs and Combinatorics*, pages 262–272, Berlin, Heidelberg, 1974. Springer Berlin Heidelberg.
- [18] S. E. Goodman, S. T. Hedetniemi, and P. J. Slater. Advances on the hamiltonian completion problem. *J. ACM*, 22(3):352–360, July 1975.
- [19] K. Helsgaun. An effective implementation of the linkernighan traveling salesman heuristic. *European Journal of Operational Research*, 126(1):106 – 130, 2000.
- [20] K. Helsgaun. *Using POPMUSIC for Candidate Set Generation in the Lin-Kernighan-Helsgaun TSP Solver*. Roskilde Universitet, 7 2018.
- [21] J. Jookan. Combinatorische optimalisatie als een perturbatie. Master’s thesis, Leuven, 2019.
- [22] R. M. Karp. *Reducibility among Combinatorial Problems*, pages 85–103. Springer US, Boston, MA, 1972.
- [23] J. Leskovec and R. Sosič. SNAP: A general-purpose network analysis and graph-mining library. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 8(1):1, 2016.
- [24] H. P. Maretic and A. Grbic. A heuristics approach to hamiltonian completion problem (hcp). In *2015 38th International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO)*, pages 1607–1612, May 2015.
- [25] L. Pósa. Hamiltonian circuits in random graphs. *Discrete Mathematics*, 14(4):359 – 364, 1976.
- [26] A. Raychaudhuri. The total interval number of a tree and the hamiltonian completion number of its line graph. *Information Processing Letters*, 56(6):299 – 306, 1995.
- [27] V. J. Rayward-Smith. The hamiltonian completion problem and its solution. *Engineering Optimization*, 10(4):309–322, 1987.
- [28] K. Takamizawa, T. Nishizeki, and N. Saito. Linear-time computability of combinatorial problems on series-parallel graphs. *J. ACM*, 29(3):623–641, July 1982.
- [29] R. E. Tarjan and J. van Leeuwen. Worst-case analysis of set union algorithms. *J. ACM*, 31(2):245–281, Mar. 1984.