ELSEVIER

Iterative patching and the asymmetric traveling salesman problem

Marcel Turkensteen [a,*], Diptesh Ghosh [c], Boris Goldengorin [b], Gerard Sierksma [b]

[a] *Faculty of Economics, University of Groningen, P.O. Box 800, 9700 AV Groningen, The Netherlands*
[b] *Faculty of Economics, University of Groningen, The Netherlands*
[c] *P&QM Area, Indian Institute of Management, Ahmedabad, India*

Abstract

Although Branch-and-Bound (BnB) methods are among the most widely used techniques for solving hard problems, it is still a challenge to make these methods smarter. In this paper, we investigate *iterative patching*, a technique in which a fixed patching procedure is applied at each node of the BnB search tree for the Asymmetric Traveling Salesman Problem. Computational experiments show that iterative patching results in general in search trees that are smaller than the classical BnB trees, and that solution times are lower for usual random and sparse instances. Furthermore, it turns out that, on average, iterative patching with the Contract-or-Patch procedure of Glover, Gutin, Yeo and Zverovich (2001) and the Karp–Steele procedure are the fastest, and that 'iterative' Modified Karp–Steele patching generates the smallest search trees.
© 2005 Elsevier B.V. All rights reserved.

Keywords: Asymmetric traveling salesman problem; Branch-and-Bound; Upper bounding; Patching

1. Introduction

The Asymmetric Traveling Salesman Problem (ATSP) is usually solved exactly by means of Branch-and-Bound (BnB) algorithms and Branch-and-Cut (BnC) algorithms; see [8]. In BnB type algorithms, an Assignment Problem (AP) is solved at every node of this tree, and the value of the optimal AP solution serves as a lower bound of the ATSP solution. A part of the search tree can be discarded when its lower bound exceeds an upper bound. This upper bound is usually the value of a shortest complete tour found so far. A class of heuristics applied to construct such a tour is *patching*. The question is: at which nodes of the search tree should such a tour be constructed? Patching at a node may reduce the search tree and the solution time, but if the reduction is too small, the overall solution time is increased due to the time invested in patching.

In the literature, the most effective BnB methods do not patch at each node; see for example, [13,3]. These methods use a best first search strategy, i.e., the subproblem with the smallest lower bound is solved first. According to these studies, patching at every node is too time-consuming.

In this paper, we consider a BnB algorithm that applies depth first search, which means that the most recently generated subproblem is solved first. This strategy requires algorithms to use much less computer memory than do

Fig. 1. ATSP instance.

Fig. 2. Minimum cycle cover.

best first strategies. Hence, it is useful for solving large problems. We apply *iterative patching*, in which a fixed patching procedure is applied at every node of the BnB depth first search tree. Four iterative patching procedures are considered in our computational experiments. These procedures are described in [5].

Given a set of locations and the distance between any pair of locations, the ATSP is the problem of finding a shortest Hamiltonian tour; i.e., a shortest round trip visiting each location exactly once. Fig. 1 is an example of an underlying graph that defines an instance of an ATSP. The nodes of the graph represent locations, and the arcs the connections between the locations. A number next to an arrowhead denotes the cost of traveling along that arc.

General instances of the ATSP are often solved to optimality by means of enumeration algorithms, in which a fraction of all feasible solutions are checked. BnB methods explore the solution space by using a search tree. We discuss BnB algorithms that solve an Assignment Problem (AP) at each node of the corresponding search tree. After solving the AP a minimum cycle cover F is obtained, say, consisting of k cycles ($k \geq 1$). In the example of Fig. 2, three cycles are generated. If $k > 1$, the subcycles in F can be *patched* into a complete tour. BnB algorithms use the value of a patching solution as an upper bound by which nodes of the search tree are fathomed.

A *patching operation* is the simultaneous deletion of two arcs from a cycle cover and the insertion of two other arcs, such that the number of cycles is reduced by one. In our example, two patching operations are needed for the generation of a complete tour (see Fig. 3), namely first arcs (2,4), (5,6) are deleted and (2,6) and (5,4) are inserted, and then we delete (12,9) and (2,6) and insert (2,9) and (12,6). The resulting tour is generally feasible but not optimal.

Fig. 3. Obtaining a tour by means of two patching operations.

Fig. 4. Best patching solution is not a shortest tour.

In [10], patching is defined as a sequence of $k - 1$ patching operations on a cycle cover of k cycles, $k \geq 1$. Recall that even a best possible patching procedure consisting of $k - 1$ patching operations does not always yield a shortest complete tour. For example, consider the sparse network in Fig. 4. The minimum cycle cover consists of the $k = 2$ cycles (1, 2, 3, 4, 5, 1) and (6, 7, 8, 9, 6) with total length 29. The unique shortest complete tour is (1, 2, 8, 9, 6, 7, 4, 3, 5, 1) with length 31. Since four arcs need to be inserted and deleted, this tour cannot be constructed from the cycle cover by means of one patching operation. Different patching procedures are introduced in the literature; see [5,10, 11,16]. These patching procedures are discussed in Section 3.

Most heuristics for the ATSP apply patching procedures only once, such as to obtain approximations to optimal solutions; see e.g. [5–7]. BnB algorithms apply patching procedures in order to obtain good feasible solutions with which parts of the search tree can be discarded. Any heuristic may be used to generate such solutions, but patching procedures are the most natural choices, since they use the structure of the already constructed minimum cycle cover. If a fixed patching procedure is applied at every node in a BnB algorithm, we call it *iterative patching*. (See Fig. 5.)

The currently best BnB algorithms for the ATSP are introduced in [3,13]. We call these the CDT algorithm and the MP algorithm, respectively. The CDT algorithm uses the patching procedure from Karp and Steele [11] at the top node of the search tree. Only if the number of zeros in the reduced matrix at the top node exceeds a threshold value β, then a subtour-merging procedure is carried out at each node of the search tree.

The subtour-merging procedure constructs first an admissible graph of zero-cost elements in the reduced matrix and then tries to find a complete tour in the admissible graph. The subtour-merging procedure patches cycles together,
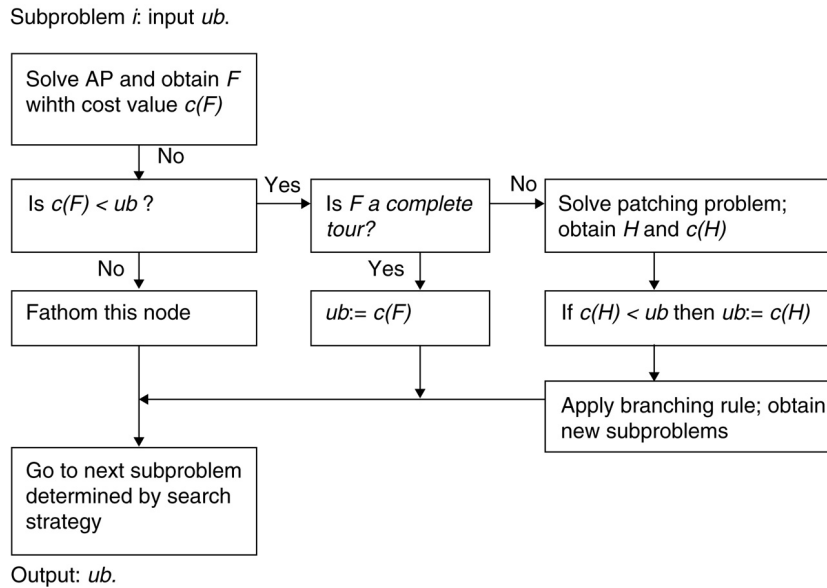
Subproblem *i*: input *ub*.

Fig. 5. Flowchart of a subproblem of a BnB algorithm with iterative patching.

but only when a zero-cost patching operation is available. It usually does not return a complete tour. In Carpaneto et al. [3], it is found that if β is set to $2.5n$, the solution times are the shortest, where n is the dimension of the instance.

The MP algorithm applies the Karp–Steele patching procedure, but not at every node of the search tree. Nodes close to the top node are patched more often than nodes deep in the tree. This algorithm also applies a subtour-merging procedure at each node.

The CDT and the MP algorithm both use a best first search (BFS) strategy, which means that a node with the smallest lower bound value is expanded next. BFS is the fastest search strategy, but requires exponential memory space. As a consequence, BFS algorithms are generally restricted to small or easily solvable problems [17]. In depth first search (DFS), the most recently generated subproblem is solved first, and it requires polynomial memory space. This makes it suitable for solving large and difficult instances. However, the search trees and solution times of DFS algorithms are usually large.

Miller and Pekny [13] report that iterative patching is too time-consuming. This may be true for BFS algorithms, but our algorithms use DFS. DFS algorithms search through deep nodes of the search tree even at an early stage; lower bounds of such nodes are generally high. A tight upper bound obtained early enables the algorithm to discard a large fraction of these nodes. Therefore, a DFS algorithm is more likely to benefit from a good upper bounding procedure, such as iterative patching, than a BFS algorithm.

The computational experiments in Section 4 compare the search tree sizes and the running times of BnB algorithms that apply iterative patching with a DFS implementation of the CDT algorithm. We apply four patching procedures, namely the ones discussed in Glover et al. [5]. The main questions that we answer on iterative patching in this paper are as follows. Is iterative patching effective for DFS algorithms? Is it true that if a patching procedure returns on average shorter tours than some other one, then, again on average, the search tree sizes are smaller and the running times are shorter? Hence, does better patching lead to the smaller search trees and shorter running times?

2. The quality of patching procedures

Let $G(V, A)$ be a graph with vertex set V and arc set A. A minimum cycle cover $F \subset A$ can be determined in $O(n^3)$ time by means of the Hungarian algorithm; see for example [9]. The speed of the Hungarian algorithm can be increased in successor nodes j to $O(n^2)$ by starting from the optimal solution in the parent node, i.e., the node in which subproblem j is generated; see for example Fischetti et al. [8].

Patching procedures delete pairs of arcs from F and insert pairs of arcs from $A \setminus F$ in such a way that a Hamiltonian cycle $H \subset A$ is obtained. The patching cost of any patching procedure P is then denoted by $c_P(F)$ and defined as

$$c_P(F) = \sum_{a \in H \setminus F} c(a) - \sum_{b \in F \setminus H} c(b), \tag{1}$$

where $c(a)$ denotes the cost of arc $a \in A$. The first term of (1) indicates the cost of the new arcs introduced by P, and the second term represents the cost of the arcs removed from the cycle cover. For any subset $Q \subset A$, $c(Q)$ denotes the sum of the cost of the arcs in Q.

Let $F_j \subset A$ denote a minimum cycle cover at node j of the BnB search tree in progress. By $BnB(Br, S, UBS)$ we denote a BnB algorithm for the ATSP that applies branching rule Br, search strategy S, and upper bounding strategy UBS. A branching rule Br partitions the current feasible regions into subsets. We consider branching rules that only depend on the current minimum cycle cover. The search strategy S in this paper is DFS. The upper bounding strategy UBS consists of two components: the first component prescribes at which nodes an upper bounding procedure should be applied, and the second component specifies the upper bounding procedure to be used. Clearly, iterative patching is an upper bounding strategy, where a tour is generated at every node of the search tree by means of a fixed patching procedure. If no confusion is likely, we simply write $BnB(UBS)$, since S and Br are fixed in this study.

Note that, in case of DFS, the order of node expansion is independent of the bounds used at each subproblem. For instance, if both algorithms $BnB(P_1)$ and $BnB(P_2)$ explore two subproblems S_1 and S_2, and $BnB(P_1)$ explores S_1 before S_2, then $BnB(P_2)$ will explore S_1 before S_2 as well.

Let $ub_j(UBS)$ be the current *upper bound*, i.e. the shortest complete tour obtained until node j using upper bounding strategy UBS. Recall that, when the UBS is iterative patching, we obtain at each node of the search tree a complete tour, i.e. a candidate for the value of $ub_j(UBS)$.

Node k is called a *successor* of j in a search tree if j is an intermediate node of the shortest path between k and the top node of the search tree; we use the notation $k \propto j$. Since the feasible region of the AP at node k is a subset of the feasible region of the AP at node j, we have of course that $c(F_k) \geq c(F_j)$ if $k \propto j$; see e.g. [17].

In the case of iterative patching, one may expect that if patching costs are low, then upper bounds are tighter and a larger number of subproblems can be fathomed. Theorem 1 formalizes this assertion: if for each instance patching procedure P_1 is cheaper than patching procedure P_2, then the search tree of $BnB(P_1)$ will be smaller than the search tree of $BnB(P_2)$.

For any iterative patching procedure P, let $BnB(P)$ be the algorithm that uses P iteratively. Define $\#BnB(P)$ as the size of the solution tree of $BnB(P)$, i.e. the number of nodes in this tree. We assume in Theorem 1 that $BnB(P_1)$ and $BnB(P_2)$ use the same AP-solver implementation, since the choice of another AP solver may result another initial minimum cycle cover. The cycle cover is the starting point of the patching procedure; if the minimum cycle covers are different, then the resulting patching solutions and their costs may differ, even though the same patching procedure is applied.

Theorem 1. *Let $\mathcal{F}$ be the set of minimum cycle covers of a given instance of the ATSP, and let P_1 and P_2 be two patching procedures such that their respective patching costs satisfy $c_1(F) \leq c_2(F)$ for each $F \in \mathcal{F}$. It then follows that $\#BnB(P_1) \leq \#BnB(P_2)$.*

Proof. For any given instance of the ATSP, let $T(Br)$ be the complete search tree based only on branching rule Br, i.e. the search tree in which all possible solutions are enumerated. Usual BnB procedures apply the following pruning operations:

(1) If at a certain node of $T(Br)$ F is a complete tour, then all successor nodes are deleted from $T(Br)$.
(2) If at a certain node of $T(Br)$, say j, it holds that $c(F_j) \geq ub_j(P)$, then this node and all its successors are fathomed.

For any patching procedure P, $BnB(P)$ deletes nodes from the complete search tree $T(Br)$ until the usual BnB tree remains, which we denote by $T(P)$. Clearly, pruning operation (1) is independent of the patching procedure used, since the AP-solver implementation is taken fixed. Actually, at each node the same minimum cycle cover is found.

We now show that $T(P_1) \subseteq T(P_2)$ by showing that if node j is fathomed under P_2, then it is also fathomed under P_1. This is the case, if for each node j, it holds that $c(F_j) \geq ub_j(P_2) \implies c(F_j) \geq ub_j(P_1)$. So we need to show

that $ub_j(P_1) \leq ub_j(P_2)$ for each node j on the path obtained by the search strategy S. Thus, $BnB(P_2)$ is only able to discard nodes if $BnB(P_1)$ discards them, which implies that $\#BnB(P_1) \leq \#BnB(P_2)$.

Obviously, for the first node $j = 0$, it holds that $ub_0(P_1) \leq ub_0(P_2)$. Now assume that $ub_j(P_1) \leq ub_j(P_2)$ at node j. Let k be the next unsolved subproblem after node j according to the search strategy S. We show that $ub_k(P_1) \leq ub_k(P_2)$. Let $H_P(F)$ be the patching solution of procedure P given minimum cycle cover F.

After solving the AP at node j, both algorithms compare $c(F_j)$ with their current upper bounds. Three scenarios are possible:

(1) If $ub_j(P_1) \leq ub_j(P_2) \leq c(F_j)$, then both algorithms fathom node j and both procedures proceed to node k. Clearly, $ub_k(P_1) = ub_j(P_1) \leq ub_j(P_2) = ub_k(P_2)$.
(2) If $c(F_j) < ub_j(P_1) \leq ub_j(P_2)$, then both algorithms execute patching at node j. Since $c_1(F_j) \leq c_2(F_j)$, it follows that $c(H_1(F_j)) = c(F_j) + c_1(F_j) \leq c(F_j) + c_2(F_j) = c(H_2(F_j))$. Since $ub_k(P_i) = \min\{ub_j(P_i), c(H_i(F_j))\}$ for $i = 1, 2$, we have that $ub_k(P_1) \leq ub_k(P_2)$.
(3) If $ub_j(P_1) \leq c(F_j) < ub_j(P_2)$, then $BnB(P_1)$ fathoms node j, and $ub_k(P_1) := ub_j(P_1)$. $BnB(P_2)$ solves an additional patching problem at node j and possibly at the successor nodes of j. Let q be the successor node of j in which the best patching solution is obtained, i.e. $q = \arg\min_l\{c(H_2(F_l)); l \propto j, l = j\}$. After searching through all successors of j, or after discarding them, $BnB(P_2)$ arrives at node k with $ub_k(P_2) \geq \min\{ub_j(P_2), c(H_2(F_q))\}$. Clearly, $ub_k(P_1) = ub_j(P_1)$. Furthermore, it holds that $ub_j(P_2) \geq ub_j(P_1) = ub_k(P_1)$, and that $c(H_2(F_q)) \geq c(F_q) \geq c(F_j) \geq ub_j(P_1) = ub_k(P_1)$. Hence, $ub_k(P_2) \geq ub_k(P_1)$.

Hence, for all nodes j on the path according to S through $T(Br)$, we have that $ub_j(P_1) \leq ub_j(P_2)$. Therefore, $\#BnB(P_1) \leq \#BnB(P_2)$. $\square$

Theorem 1 can be extended to upper bounding strategies *UBS* for which the upper bound generated at node j is at least $c(F_j)$. In that case, upper bounds are only obtained at nodes at which a complete tour is constructed; elsewhere, the patching costs are infinite. For example, consider a BnB algorithm $BnB(P; ni)$ that applies patching procedure P not iteratively. It follows from Theorem 1 that its search tree is always at least the size of the search tree of the algorithm $BnB(P)$ that applies P iteratively.

In general, there are few iterative patching procedures that always return better patching solutions than some other one. Therefore, it makes more sense to consider the average performance of patching procedures. To this end, we conduct computational experiments in Section 4.

The most important measure of the quality of algorithms are solution times. Actually, high quality patching solutions may lead to long solution times of subproblems. So usually, a trade-off is made between the quality of the patching and time invested in patching. For instance, if patching procedure P is only applied at the top node, the search tree is larger than the tree with iterative patching procedure P. However, the average solution time at the nodes is smaller. In Section 4, solution times are taken into account more explicitly.

The following observation allows to increase the speed of iterative patching without losing quality. Recall that, if a cycle cover F consists of k cycles, patching is a sequence of $k - 1$ patching operations. Call the cycle cover after the i-th patching operation F_i, and denote its cost by $c(F_i)$, $i = 1, \ldots, k - 1$. If $c(F_i)$ exceeds the cost of the current best solution ub, the patching procedure will certainly not lead to a better solution, since the cost of each patching operation is nonnegative. Hence, we can abort the patching after i steps and save running time.

3. Patching procedures

We now compare the performance of four iterative patching procedures based on the four most well-known patching algorithms. We start with a short description of these four patching procedures. All these procedures have a worst-case time complexity of $O(n^3)$; see [5].

Karp–Steele patching (KSP) was introduced in [11]. Starting with the minimum cycle cover F, KSP patches the two longest subcycles successively by using a cheapest patching operation. In our example, KSP patches cycles 1 and 3 by deleting (10,2) and (9,8), and adding (10,8) and (9,2); see Fig. 6. The new cycle is then patched with cycle 2 by removing (12,9) and (5,6), and inserting (5,9) and (12,6).

Modified Karp–Steele patching (MKS), also called *Greedy Karp–Steele patching*, see [5], performs the cheapest patching operation among all pairs of cycles in the current cycle cover. The patching costs are then updated and the procedure is repeated until a complete tour is obtained. Since it compares in general more patching operations than

Fig. 6. Modified Karp–Steele patching in action.

Fig. 7. RPC patching solution.

KSP, MKS is more time-consuming. In our example, MKS joins cycles 2 and 3 by deleting arcs (5,6) and (12,9), and inserting (5,9) and (12,6). Cycle 1 is included by inserting (2,9) and (5,4) and removing (2,4) and (5,9); see Fig. 6.

Recursive Path Contraction (*RPC*) was introduced in [16]. From all, say k, cycles a most expensive arc is deleted and the remaining paths are contracted, so transformed into single nodes. On these k nodes an AP is solved. So every contracted path is connected to another contracted path. The procedure is carried out recursively until one cycle is obtained. The calculations of Section 4 use the implementation from Glover et al. [5]. In our example, the most expensive arc from every cycle is deleted, namely (3,1), (5,6) and (12,9). The end nodes 3, 5 and, 12 are assigned to nodes 9, 1, and 6, respectively. Finally, the tour depicted in Fig. 7 is obtained.

Contract-or-Patch (*COP*) is a two-stage procedure consisting of RPC in the first stage and, either MKS or KSP in the second stage; see [5,6]. All cycles with length less than a user-defined threshold value t are patched using RPC. In [6], it is shown that the threshold value $t = 5$ is the most robust choice for different types of instances. Given the cycle cover from Fig. 2, cycles 2 and 3 are patched using the RPC procedure. The long cycles in the current cycle cover are patched with either KSP or MKS. In Section 4, the faster procedure KSP is selected, since in [7] it is asserted that there is no significant difference in the patching cost of COP using either KSP or MKS.

Table 1
Patching strategies tested

Name	Patching strategy
BnB(*KSP*)	Iterative KSP
BnB(*MKS*)	Iterative MKS
BnB(*RPC*)	Iterative RPC
BnB(*COP*)	Iterative COP
BnB(*CDT*)	CDT algorithm

4. Computational experiments

In this section, we compare both the tree sizes and the running times of the algorithms presented in Table 1. Recall that the size of a BnB tree is the number of subproblems solved before the first optimal solution is determined, i.e. the number of nodes visited on the path followed through $T(Br)$ according to search strategy S. The results of iterative patching procedures are compared with the results of the DFS implementation of the CDT algorithm. The DFS implementation is of practical use, because it solves ATSPLIB and symmetric instances which a BFS approach cannot solve; see for example [3,13].

The experiments are performed on a Pentium 4 computer with speed 2 GHz and 256 MB RAM under Windows 2000. The programming language is C and the compiler is GNU with speed -o2. Our branching rule branches by a largest cost arc in the shortest subcycle of a minimum cycle cover. In a forthcoming study we will apply tolerance-based branching rules, where branching is performed on an arc with the smallest tolerance value (the amount at which the cost can be changed without changing the solution at hand). The iterative patching procedures are tested for the following types of instances:

(1) Asymmetric TSPLIB instances (see [15]).
(2) Randomly generated instances with varying degree of symmetry.
(3) Randomly generated instances with varying degree of sparsity.
(4) Random instances with a large number of different intercity distances.
(5) Almost symmetric Buriol instances (see [2]).

From all asymmetric TSPLIB instances we have selected 16 instances that are solvable within reasonable time limits. The random instances have degree of symmetry 0, 0.33, 0.66, and 1, where the *degree of symmetry* is defined as the fraction of off-diagonal entries in the cost matrix $\{c_{ij}\}$ that satisfy $c_{ij} = c_{ji}$. The third class of instances consists of instances with varying *degree of sparsity*, defined as the fraction of the total possible number of arcs that are missing. We study instances with degree of sparsity of 0, 0.25, 0.5, and 0.75.

If both the degree of symmetry and degree of sparsity of an instance is unrestricted, we call such an instance *usual random*. The usual random instances have problem size 60, 70, 80, 100, 200, 300, 400, and 500. Random instances with degree of symmetry larger than 0 have problem size 60, 70, and 80. Only these samples of (quasi-)symmetric instances are considered, since computation times for larger symmetric instances tend to be extremely long. The instances with varying degree of sparsity have problem size 100, 200, and 400. The arc costs are drawn from a discrete uniform distribution supported on $\{1, 2, \ldots, 10^4\}$; for each problem set and for all problem sizes, 10 instances are generated. In comparison with other studies, namely [3,13], our random instances are relatively small, whereas our symmetric instances are relatively large. For example, the MP algorithm given by Miller and Pekny [13] solves random instances of size 500 000, but solves symmetric instances of size less than 30 only.

In addition to the usual random instances, we generate random instances with a large number of different intercity distances. The reason for considering these instances is given in [19,20], where it is shown that if the number of different intercity distances exceeds a threshold value, the instance becomes relatively hard to solve. Suppose the arc costs of an instance are uniformly distributed on $\{1, \ldots, R\}$, where R is the range of the distribution. It is shown in [19] that, as the range increases, the number of intercity distances increases as well. Moreover, uniform random instances are hard to solve if the range R is at least n^2; see [20]. This result implies that our randomly generated instances with size larger than 100 are relatively easy to solve. Therefore, we use additional 'hard' random instances with arc costs drawn from a uniform distribution supported on $\{1, \ldots, 10^5\}$ for $n = 200, 300$, and on $\{1, \ldots, 2 \times 10^5\}$ for $n = 400$.

Table 2
Normalized size of search tree for usual BnB (CDT = 100)

	CDT	KSP	MKS	RPC	COP
ATSPLIB	100.00	95.03	94.27	101.40	95.37
Buriol, $k' = 5$	100.00	98.37	97.39	98.97	98.01
Buriol, $k' = 50$	100.00	78.64	39.95	102.95	95.70
Usual random	100.00	47.27	43.97	129.98	47.27
Random, large range	100.00	48.99	48.99	167.83	48.99
Degree of symmetry 0.33	100.00	50.81	50.65	106.75	51.16
Degree of symmetry 0.66	100.00	74.52	73.66	101.45	75.44
Full symmetry	100.00	99.79	99.77	99.97	99.80
Degree of sparsity 0.25	100.00	51.66	51.20	113.26	51.66
Degree of sparsity 0.50	100.00	56.13	56.13	126.68	56.13
Degree of sparsity 0.75	100.00	56.43	56.35	129.98	56.43

Table 3
Normalized running times

	CDT	KSP	MKS	RPC	COP
ATSPLIB	100.00	114.81	139.56	114.44	116.01
Buriol, $k' = 5$	100.00	111.31	158.72	130.00	125.56
Buriol, $k' = 50$	100.00	91.25	57.53	116.91	112.69
Usual random	100.00	55.81	60.24	140.44	54.45
Random, large range	100.00	61.05	81.07	191.03	64.88
Degree of symmetry 0.33	100.00	72.22	72.22	170.83	55.56
Degree of symmetry 0.66	100.00	93.33	103.70	132.22	85.00
Full symmetry	100.00	108.24	126.76	114.98	111.54
Degree of sparsity 0.25	100.00	62.64	73.57	125.13	62.33
Degree of sparsity 0.50	100.00	69.05	90.16	144.79	77.44
Degree of sparsity 0.75	100.00	73.79	85.33	153.29	73.88

Finally, the almost symmetric Buriol instances are asymmetric TSP instances which are derived from symmetric instances from the TSPLIB. They are constructed as follows. Let σ be the average of all distances of the original symmetric TSPLIB instance, and let k' be a user-defined percentage. Then each entry in the lower diagonal of the cost matrix C of the symmetric instance is increased by $k\sigma$, where k is randomly drawn from the uniform distribution supported on $\{0, \ldots, k'\}$. However, the costs of edges belonging to a chosen optimal tour of the original symmetric instance are not modified. Note that the smaller the value of k', the higher the degree of symmetry of the instances generated. We construct eight instances for which $k' = 5$, and twelve for which $k' = 50$, respectively, using the instance generator from Buriol et al. [2]. These are instances which are solvable within reasonable time limits.

The average *size of the search tree* of the algorithms is shown in Table 2. In order to make the results more comparable, we have used *normalized* results, i.e., we have fixed the results of *BnB(CDT)* at 100. The number '50.65' in the MKS-column means that the *BnB(MKS)* generates on average about half the number of subproblems of *BnB(CDT)* for instances with degree of symmetry 0.33.

Table 2 shows that, except for the RPC procedure, iterative patching leads to smaller search trees. The search tree reductions of iterative patching are large for usual random and sparse instances; the sizes of the trees of *BnB(KSP)*, *BnB(MKS)* and *BnB(COP)* are half the size of the search tree of *BnB(CDT)*. The reductions of iterative patching are smaller for symmetric and ATSPLIB instances. *BnB(MKS)* and also *BnB(KSP)* have relatively large search tree reductions for the almost symmetric Buriol instances with $k' = 50$, but *BnB(COP)* is performing considerably worse. On average, the search trees generated by *BnB(MKS)* are the smallest, whereas *BnB(RPC)* only generates reasonably small search trees for symmetric instances.

In Table 3, we present the normalized *running times*. For usual random and sparse instances, iterative patching is clearly more effective; the search tree reduction outweighs the time invested in patching at nodes. Although

Table 4
Search tree sizes and solution times (seconds) of ATSPLIB instances

Instance	CDT		KSP		MKS		RPC		COP	
	Size	Time	Size	Time	Size	Time	Size	Time	Size	Time
ft53	21 189	2.20	20 111	2.31	20 111	2.64	21 189	2.36	20 111	2.42
ft70	26 025	3.57	25 831	3.85	25 831	4.40	26 025	4.01	25 831	4.07
ftv33	7 455	0.16	7 065	0.22	7 061	0.27	7 307	0.22	7 065	0.22
ftv35	7 305	0.16	6 945	0.16	6 939	0.22	8 267	0.22	6 951	0.22
ftv38	7 325	0.22	6 195	0.22	6 195	0.27	10 101	0.38	6 195	0.16
ftv44	3 753	0.11	619	0.01	619	0.05	3 753	0.16	3 083	0.16
ftv47	29 539	1.10	29 025	1.26	29 017	1.76	29 539	1.32	29 031	1.37
ftv55	114 403	4.73	92 447	4.51	92 447	5.82	114 785	5.44	103 839	5.55
ftv64	252 755	11.87	43 441	3.19	43 441	4.18	252 755	15.93	43 441	3.52
ftv70	326 827	23.41	253 873	24.95	206 195	27.36	410 545	35.60	261 199	24.73
ftv170	1 796 439	1073.63	1 796 149	1300.88	1 796 159	1614.56	1 796 459	1 198.96	1 796 149	1276.87
rbg323	3	0.05	3	0.05	1	0.05	9	0.05	3	0.01
rbg358	3	0.05	3	0.05	1	0.16	7	0.11	5	0.05
rbg403	3	0.05	3	0.05	1	0.11	7	0.11	3	0.05
rbg443	3	0.05	3	0.05	1	0.11	3	0.11	3	0.05
br17	3 674 829	16.59	3 674 829	24.23	3 674 829	32.69	3 674 829	24.51	3 674 829	24.40

Table 5
Search tree sizes and solution times (seconds) of usual random instances

n	CDT		KSP		MKS		RPC		COP	
	Size	Time	Size	Time	Size	Time	Size	Time	Size	Time
60	6 508	0.60	3 808	0.38	3 808	0.44	12 880	1.10	3 808	0.33
70	10 828	1.21	4 528	0.44	4 528	0.71	18 522	2.14	4 528	0.55
80	21 834	2.75	9 014	1.26	8 622	1.48	27 822	4.1	9 014	1.26
100	13 454	2.42	9 002	1.92	6 814	1.81	17 424	3.73	9 002	1.98
200	138 522	114.	36 390	33.	3 639	40.	172 054	151.	36 390	33.
300	412 930	798.	178 498	481.	178 498	551.	500 100	1 081.	178 498	424.
400	525 088	2142.	284 994	1410.	284 982	1746.	640 440	2 825.	284 994	1349.
500	951 188	6428.	434 576	3687.	432 000	5284.	1 456 440	10 868.	434 576	3889.

BnB(*MKS*) often requires the smallest search trees, *BnB*(*COP*) and *BnB*(*KSP*) mostly display smaller running times. This indicates that the speed of solving patching problems is relevant. Solution times of iterative patching are longer for instances from the ATSPLIB and for symmetric instances than of *BnB*(*CDT*), although in both cases the differences are small.

The following tables show the absolute search tree sizes and solution times in more detail. For most ATSPLIB instances, the search tree reductions of iterative patching are minor, and the solution times increase; see Table 4. For the usual random instances, the iterative patching procedures *BnB*(*KSP*), *BnB*(*MKS*), and *BnB*(*COP*) have clearly smaller search tree sizes and solution times than *BnB*(*CDT*); see Table 5. These benefits appear to be independent of the instance size. Finally, Tables 6 and 7 present the absolute tree sizes and solution times of sparse and symmetric instances.

The results for the almost symmetric instances from [2] are presented in Tables 8 and 9. They indicate that patching does not lead to shorter solution times for most instances with $k' = 5$, but it becomes worthwhile if the deviation k' is increased to 50.

Table 10 presents the results for the random instances with a large number of different intercity distances. We obtain similar results as for the usual random instances: the solution times of *BnB*(*KSP*), *BnB*(*MKS*), and *BnB*(*COP*) are clearly lower than those of *BnB*(*CDT*) and *BnB*(*RPC*). So iterative patching is not sensitive to changes in the range of the uniform distribution.

Symmetric and ATSPLIB instances can be considered 'hard', i.e., even small instances have large search trees and running times. For these instances, cycle covers often consist of many short cycles. Hence, tours obtained by patching are long, and only minor parts of the search tree can be discarded, so the small reductions of the search tree

Table 6
Search tree sizes of symmetric and sparse instances

Instance	CDT	KSP	MKS	RPC	COP
	Size	Size	Size	Size	Size
Degree of symmetry 0.33	122 520	58 878	58 724	129 914	59 458
Degree of symmetry 0.66	259 626	202 894	200 630	264 444	204 470
Full symmetry	114 984 046	114 912 026	114 908 592	109 843 207	114 915 850
Degree of sparsity 0.25	637 872	362 188	354 610	732 500	362 188
Degree of sparsity 0.50	653 016	368 736	368 736	801 526	368 736
Degree of sparsity 0.75	704 832	386 468	386 392	883 026	386 468

Table 7
Solution times (seconds) of symmetric and sparse instances

Instance	CDT	KSP	MKS	RPC	COP
	Time	Time	Time	Time	Time
Degree of symmetry 0.33	13	8	8	19	7
Degree of symmetry 0.66	33	33	38	45	32
Full symmetry	17 584	19 182	22 521	19 271	19 972
Degree of sparsity 0.25	1 935	1 451	1 801	2 386	1 434
Degree of sparsity 0.50	1 797	1 341	1 746	2 350	1 345
Degree of sparsity 0.75	1 998	1 467	1 909	2 857	1 472

Table 8
Running times and search tree sizes for almost symmetric Buriol instances, $k' = 5$

Instance	CDT		KSP		MKS		RPC		COP	
	Size	Time	Size	Time	Size	Time	Size	Time	Size	Time
ulysses16	10 607	0.05	10 165	0.05	10139	0.11	10 877	0.16	10 123	0.11
ulysses22	863 039	8.41	857 027	13.08	818 327	11.65	863 039	12.31	832 311	10.33
bayg29	12 443	0.27	6 555	0.16	5 563	0.16	12 443	0.44	6 555	0.22
bays29	17 283	0.33	13 931	0.33	12 355	0.33	16 471	0.49	14 687	0.33
eil51	1197 185	57.64	1193 015	64.89	1193 015	89.78	1197 185	79.34	1196 191	72.25
fri26	12 891	0.22	12 725	0.22	9 187	0.27	12 891	0.33	12 759	0.27
gr24	2 311	0.05	1 329	0.05	1 329	0.00	2 311	0.05	2 299	0.05
gr48	3400 347	155.55	3331 593	168.90	3322 435	250.88	3344 327	196.15	3331 611	195.82

Table 9
Running times and search tree sizes for almost symmetric Buriol instances, $k' = 50$

Instance	CDT		KSP		MKS		RPC		COP	
	Size	Time	Size	Time	Size	Time	Size	Time	Size	Time
ulysses16	2 553	0.00	2 481	0.05	2 481	0.00	12 025	0.11	7 655	0.05
ulysses22	94 831	0.88	35 577	0.44	17 711	0.27	99 963	1.21	89 035	1.10
att48	181 943	7.53	147 201	7.42	5 967	0.38	186 547	9.45	147 201	7.25
bayg29	217	0.00	177	0.00	13	0.05	217	0.00	211	0.00
bays29	2 289	0.11	323	0.00	223	0.00	2 613	0.05	521	0.05
eil51	21 689	1.10	16 869	0.99	951	0.05	26 399	1.65	20 497	1.26
fri26	2 495	0.05	2 481	0.05	2 481	0.05	2 495	0.05	2 481	0.05
gr24	141	0.00	49	0.00	49	0.00	141	0.00	141	0.00
gr48	5 213	0.33	4 503	0.33	2 405	0.16	5 213	0.33	5 177	0.38
pr76	1 659	0.27	1 643	0.27	415	0.05	17 229	2.64	5 293	0.88
eil76	567	0.05	243	0.05	239	0.00	567	0.11	243	0.00
gr96	1038 095	262.91	851 473	239.73	507 019	156.15	1038 095	303.85	1015 109	296.87

Table 10
Search tree sizes and solution times (seconds) of random instances with a large number of different intercity distances

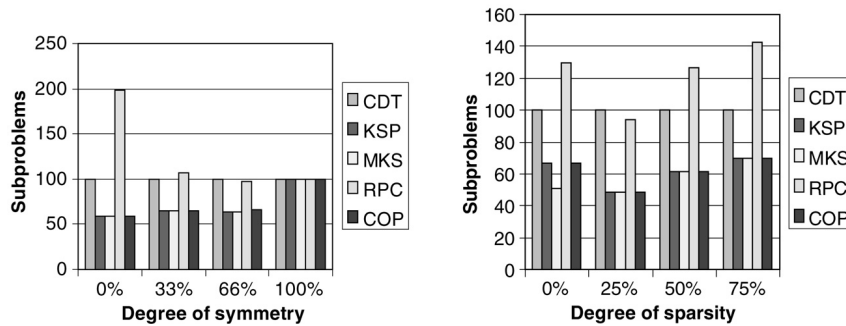| n	CDT		KSP		MKS		RPC		COP	
	Size	Time	Size	Time	Size	Time	Size	Time	Size	Time
200	153 792	146.32	42 358	43.85	42 358	52.36	42 358	54.73	185 116	174.40
300	363 654	824.01	184 752	470.33	184 752	608.13	184 752	600.93	454 114	1065.38
400	589 564	2640.22	271 034	1394.23	271 034	1815.71	271 034	1739.07	1035 296	4775.33

Fig. 8. Normalized search tree sizes of instances with varying degree of symmetry ($n = 60$) and sparsity ($n = 100$), CDT = 100.

do not compensate for the time invested in patching at each node. This explains the special behavior of symmetric and ATSPLIB instances. The same holds for the almost symmetric Buriol instances.

Table 6 and Fig. 8 show that, as the degree of symmetry increases, the search trees of *BnB*(*CDT*) and *BnB*(*RPC*) converge to the size of the other trees. Hence, applying iterative patching makes no sense for symmetric instances. On the other hand, the degree of sparsity does not influence the relative search tree sizes of the algorithms; see Fig. 8. So sparsity does not influence the usefulness of iterative patching.

The major drawback of BnB algorithms is their time consumption: it may take very long before an optimal solution is obtained. When the BnB process is terminated and the best solution so far is taken, the procedure is called Truncated Branch-and-Bound (TBnB); see [18]. Usually, the TBnB algorithm is terminated if a predefined number of nodes in the search tree is expanded. Currently, TBnB uses KSP only. An interesting question for future research is whether TBnB can be improved by including other iterative patching procedures.

Premature termination of BnB is effective if good solutions are found at an early stage of the BnB process, and a large portion of the time is spent on proving optimality. In Table 11, we present the solution quality for difficult practical instances from [2] with $k' = 5$. The solution quality reported in the table is the relative gap between the optimal solution of the instance and the best solution found by the BnB algorithm after a fixed number of subproblems. The results indicate that, even after solving a large number of subproblems, the BnB solutions are still far from optimal. For example, solving problem instance pr76 takes about 12 h; about 50% of the time is spent on finding an optimal solution. So when our BnB methods are terminated in an early stage, the resulting solutions are not competitive with metaheuristic solutions for these instances. Table 11 also shows that iterative KSP finds good solutions in an early stage, whereas the top node patching algorithm has to solve a very large number of subproblems before achieving the same solution quality.

All BnB algorithms presented in this paper are able to solve small instances in more or less the same amount of time as most metaheuristics do. However, our computational experiments indicate that these BnB methods are not able to find optimal solutions within reasonable time limits for large, almost symmetric, Buriol instances with $k' = 5$, or for fully symmetric instances. On the other hand, metaheuristics generate solutions to these instances within a few per cent from the optimal solution value in fractions of seconds; see for example the survey paper Buriol et al. [2]. The errors of metaheuristics are in the range 0 to 0.44% for ATSPLIB instances, and below 0.6% for the almost symmetric Buriol instances. However, these results do not imply that metaheuristics are always preferable over BnB methods. Although the errors appear pretty small, a solution with 0.5% higher cost than optimal may be very costly in practice.

A new line of research is combining the force of BnB and metaheuristics, in particular *memetic algorithms* [14]. In evolutionary algorithms the elements, paths in case of the ATSP, inherited from the parents are recombined, but

Table 11
Solution quality after number of subproblems solved for almost symmetric instances with $k' = 5$

Subproblems	BnB(KSP)			BnB(CDT)		
	1000 (%)	10 000 (%)	100 000 (%)	1000 (%)	10 000 (%)	100 000 (%)
pr76	9.78	9.78	9.78	20.28	19.22	16.08
eil76	7.25	7.25	7.06	21.38	19.52	13.57
gr96	7.02	7.02	7.02	14.47	13.77	11.55
kroD100	13.17	13.17	13.17	39.34	38.85	37.11
rd100	13.50	13.50	13.50	40.09	39.12	30.78
lin105	10.60	10.60	10.60	26.64	25.88	24.14
ch130	15.04	15.04	15.04	32.09	31.93	28.82
ch150	16.56	16.56	16.56	36.57	36.57	36.17
brg180	16.77	16.41	14.46	18.62	18.51	15.03
Average time (s)	0.48	4.05	35.11	1.30	2.78	28.75

Table 12
Ordering of the top node solution quality and the number of iterations

		Average relative excess over AP lower bound		Normalized search tree size (CDT = 100)
ATSPLIB	MKS	3.36%	MKS	86.15
	KSP	4.29%	KSP	87.99
	COP	4.77%	COP	88.81
	RPC	18.02%	RPC	103.38
Usual random	COP	1.88%	MKS	43.97
	MKS	3.36%	COP	47.27
	KSP	3.11%	KSP	47.27
	RPC	106.65%	RPC	129.98
Full symmetry	COP	79.87%	MKS	99.77
	RPC	183.57%	KSP	99.79
	MKS	586.92%	COP	99.80
	KSP	744.22%	RPC	99.97

memetic algorithms use optimization methods to construct good feasible solutions. Similar to BnB, the question is how much time should be spent on the optimization of each agent's tour. Buriol et al. [2] use a type of patching algorithm for the optimization, and the memetic algorithm by Cotta and Troya constructs such tours by means of a Branch-and-Bound subroutine [4].

The results indicate that it is worthwhile to invest time in finding good upper bounds in BnB algorithms with a DFS strategy. The method for finding a good upper bound need not be patching, and need not be iterative as well. Another interesting question is whether it is worthwhile to invest more time in finding a tight upper bound at the top node of the BnB search tree. For that purpose, a metaheuristic can be applied instead of patching. Klau et al. apply a memetic algorithm initially for preprocessing and to obtain a good starting solution for the Biobjective Flowshop Scheduling Problem [12]. The same strategy is followed by Basseur et al. for the Prize-Collecting Steiner Tree [1]. A hybrid application of metaheuristics and Branch and Bound may form fertile area of future research.

In [5], the performance of patching heuristics on solution quality is studied. The results show that MKS returns the best patching solutions for ATSPLIB instances, and COP for random instances, both symmetric and asymmetric. In Table 12, the solution quality results from Glover et al. [5] are compared with our search tree sizes. The results show that the ordering with respect to solution quality of patching procedures differs from the ordering with respect to search tree sizes of the corresponding iterative patching procedure. This phenomenon may be caused by the following effect. Recall that, when iterative patching is applied, patching solutions are constructed at each node of the search tree. It may be misleading to take into consideration the patching quality only at the top node of the search tree, and expect that for all nodes in the search tree on average the same quality holds. Actually, it is more likely that good upper bounds are found deep in the search tree and that the average patching solution quality deep into the tree differs

from the average top node patching quality. In fact, top node cycle covers may consist of many short cycles, whereas subcycles tend to become longer as the BnB algorithm proceeds deeper into the search tree, because our branching rule attempts to break short cycles. This may explain the differences in the orderings according to the average patching quality and to the average search tree size of the iterative patching procedures.

Consider for example the iterative patching procedures RPC and COP. *BnB(RPC)* needs long running times and large search trees for random instances, because RPC deletes an arc from every cycle without calculating patching costs. Therefore, if cycles are long, bad patching operations are likely. COP, on the other hand, patches long cycles carefully, leading to smaller search trees.

5. Conclusion

We studied the performance of four iterative patching procedures, which are fixed patching procedures at every node of the search tree, which we compared with the performance of a depth first search implementation of the CDT algorithm given by Carpaneto et al. [3]. Our performance measures are the size of the search tree and the running times of the algorithms. Clearly, there is a trade-off between the quality of patching, leading to smaller search trees, and the speed of solving each patching problem. We conclude with an answer to the main questions.

Is it worthwhile to use iterative patching procedures? At least, search trees are always smaller. However, only for 'practical' instances are the solution times shorter when *BnB(CDT)* is applied. A side effect of iterative patching is that, if calculations are finished prematurely, a satisfactory solution is often at hand; see [17]. An interesting direction of future research is to study iterative patching procedures in Truncated BnB algorithms and other metaheuristics. Another interesting direction is finding intermediate strategies between full iterative patching and top node patching. To this end, the nodes of the search tree must be identified at which a good patching solution can be expected.

Which iterative patching procedure is the most efficient one? On the whole, the algorithm using MKS generates the smallest solution trees, and our COP and KSP implementations achieve the best solution times. The most important performance criterion is usually the solution time. However, if the memory of the computer is the restrictive factor, then it also becomes important to keep the search trees small.

Acknowledgement

References

[1] M. Basseur, J. Lemesre, C. Dhaenens, E.-G. Talbi, Cooperation between branch and bound and evolutionary algorithms to solve a biobjective flow shop problem, in: Workshop on Experimental and Efficient Algorithms, WEA'04, 2004, pp. 72–86.

[2] L. Buriol, P.M. França, P. Moscato, A new memetic algorithm for the asymmetric traveling salesman problem, Journal of Heuristics 10 (2004) 483–506.

[3] G. Carpaneto, M. Dell'Amico, P. Toth, Exact solution of large-scale, asymmetric traveling salesman problems, ACM Transactions on Mathematical Software 21 (4) (1995) 394–409.

[4] C. Cotta, J.M. Troya, Embedding Branch and Bound within evolutionary algorithms, Applied Intelligence 18 (2003) 137–153.

[5] F. Glover, G. Gutin, A. Yeo, A. Zverovich, Construction heuristics and domination analysis for the asymmetric TSP, European Journal of Operational Research 129 (2001) 555–568.

[6] G. Gutin, A. Zverovich, Evaluation of the contract-or-patch heuristic for the asymmetric traveling salesman problem, INFOR 43 (2005) 23–31.

[7] D.S. Johnson, G. Gutin, L.A. McGeoch, A. Yeo, W. Zhang, A. Zverovich, Experimental analysis of heuristics for the ATSP, in: G. Gutin, A.P. Punnen (Eds.), The Traveling Salesman Problem and its Variations, Kluwer Academic Publishers, 2002, pp. 445–487.

[8] M. Fischetti, A. Lodi, P. Toth, Exact methods for the asymmetric traveling salesman problem, in: G. Gutin, A.P. Punnen (Eds.), The Traveling Salesman Problem and its Variations, Kluwer Academic Publishers, 2002, pp. 169–194.

[9] R. Jonker, A. Volgenant, Improving the Hungarian assignment algorithm, Operations Research Letters 5 (1986) 171–175.

[10] R.M. Karp, A patching algorithm for the nonsymmetric traveling salesman problem, SIAM Journal of Computing 8 (1979) 561–573.

[11] R.M. Karp, J.M. Steele, Probabilistic analysis of heuristics, in: E.L. Lawler, J.K. Lenstra, A.H.G. Rinnooy Kan, D.B. Shmoys (Eds.), The Traveling Salesman Problem, Wiley, New York, 1990, pp. 181–205.

[12] G.W. Klau, I. Ljubic, A. Moser, P. Mutzel, P. Neuner, U. Pferschy, G. Reidl, R. Weiskircher, Combining a memetic algorithm with integer programming to solve the prize-collecting Steiner tree problem, in: K. Deb, R. Poli, W. Banzhaf, H.-G. Beyer, E. Burke, P. Darwen, D. Dasgupta, D. Floreano, J. Foster, M. Harman, O. Holland, P.L. Lanzi, L. Spector, A. Tettamanzi, D. Thierens, A. Tyrrell (Eds.), Genetic and Evolutionary Computation—GECCO 2004, in: Lecture Notes in Computer Science, vol. 3102, Springer Verlag, Berlin, 2004, pp. 1304–1315.

[13] D.L. Miller, J.F. Pekny, Exact solution of large asymmetric traveling salesman problems, Science 251 (1991) 754–761.

[14] P. Moscato, C. Cotta, A gentle introduction to memetic algorithms, in: F. Glover, G.A. Kochenberger (Eds.), Handbook of Metaheuristics, Kluwer Academic Publishers, 2003, pp. 105–144.

[15] G. Reinelt, TSPLIB, http://www.informatik.uni-heidelberg.de/groups/-comopt/software/TSPLIB95, 1995.

[16] A. Yeo, Large exponential neighbourhoods for the TSP, Department of Mathematics and Computer Science, Odense University, Odense, Denmark, 1997 (preprint).

[17] W. Zhang, Truncated Branch-and-Bound: a case study on the asymmetric TSP, in: Proc. AAAI-93 Spring Symposium on AI and NP-Hard Problems, Stanford, 1993, pp. 160–166.

[18] W. Zhang, Depth-first Branch-and-Bound versus local search: a case study, in: Proc. 17th National Conf. on Artificial Intelligence, AAAI-2000, Austin, Texas, 2000, pp. 930–935.

[19] W. Zhang, Phase transitions of the asymmetric traveling salesman, in: Proc. 18th Intern. Joint Conf. on AI, IJCAI-03, Acapulco, Mexico, 2003, pp. 1202–1207.

[20] W. Zhang, Phase transitions and backbones of the asymmetric traveling salesman problem, Journal of Artificial Intelligence Research 20 (2004) 471–497.