Algorithmic Operations Research Vol.2 (2007) 83–93

On the Approximability of TSP on Local Modifications of Optimally Solved Instances

Hans-Joachim Böckenhauer [a]

[a]Department of Computer Science, ETH Zurich, Switzerland.

Luca Forlizzi [b]

[b]Departimento di Informatica, Università di L'Aquila, Italy.

Juraj Hromkovič [c]

[c]Department of Computer Science, ETH Zurich, Switzerland.

Joachim Kneis [d]

[d]Department of Computer Science, RWTH Aachen University, Germany.

Joachim Kupke [e]

[e]Google Inc., Mountain View, USA.

Guido Proietti [b,f]

[f]Istituto di Analisi dei Sistemi ed Informatica "A. Ruberti", CNR, Roma, Italy

Peter Widmayer [g]

[g]Department of Computer Science, ETH Zurich, Switzerland.

Abstract

Given an instance of TSP together with an optimal solution, we consider the scenario in which this instance is modified locally, where a local modification consists in the alteration of the weight of a single edge. More generally, for a problem U, let LM-U *(local-modification-U) denote the same problem as U, but in* LM-U*, we are also given an optimal solution to an instance from which the input instance can be derived by a local modification. The question is how to exploit this additional knowledge, i.e., how to devise better algorithms for* LM-U *than for U. Note that this need not be possible in all cases: The general problem of* LM-TSP *is as hard as* TSP *itself, i.e., unless $P = NP$, there is no polynomial-time $p(n)$-approximation algorithm for* LM-TSP *for any polynomial p. Moreover,* LM-TSP *where inputs must satisfy the β-triangle inequality* (LM-Δ_β-TSP *) remains NP-hard for all $\beta > \frac{1}{2}$. However, for* LM-Δ-TSP *(i.e., metric* LM-TSP*), we will present an efficient 1.4-approximation algorithm. In other words, the additional information enables us to do better than if we simply used Christofides' algorithm for the modified input. Similarly, for all $1 < \beta < 3.34899$, we achieve a better approximation ratio for* LM-Δ_β-TSP *than for* Δ_β-TSP*. For $\frac{1}{2} \leq \beta < 1$, we show how to obtain an approximation ratio arbitrarily close to 1, for sufficiently large input graphs.*

Key words: TSP, approximation algorithms, local modifications, relaxed triangle inequality, sharpened triangle inequality

Email: Hans-Joachim Böckenhauer [hjb@inf.ethz.ch], Luca Forlizzi [forlizzi@di.univaq.it], Juraj Hromkovič [juraj.hromkovic@inf.ethz.ch], Joachim Kneis [joachim.kneis@cs.rwth-aachen.de], Joachim Kupke [joachimk@google.com], Guido Proietti [proietti@di.univaq.it], Peter Widmayer [widmayer@inf.ethz.ch].

1. Introduction

Traditionally, optimization theory has been concerned with the task of finding good feasible solutions to (practically relevant) input instances, little or nothing about which is known in advance. Many applications, however, demand good, sometimes optimal, solutions to a limited set of input instances which reflect a supposedly-constant environment (imagine, e.g., an existing railway system or communications network). When this environment does change, maybe only slightly and maybe only locally, do we have no choice but to recompute some good feasible solution, effectively forgetting about the old one?

We will analyze *local* modifications here. This means, we do not consider small perturbations of many parts of the input, but only one local change, which might on the other hand be arbitrarily large. In a graph problem, for example, the cost of a single edge might essentially change, an edge might be removed or added, or some other local parameter might be adjusted. Results related to this work pertain to the question by how much a given instance of an optimization problem may be varied if it is desired that optimal solutions to the original instance retain their optimality [10,16,17,19,12]. In contrast with this so-called "postoptimality analysis" or "sensitivity analysis," our approach here is to ask, if we cannot avoid to lose the optimality of a given solution when an instance is varied arbitrarily, what can we do to *restore* the quality of a solution, maybe in an approximative sense?

Surely, for some problems, knowing an optimal solution to the original instance trivially makes their local-modification variants easy to solve because the given optimal solution is itself a very good solution to the modified instance. For example, adding an edge in the instance of a coloring problem will increase the cost of an optimal solution by at most the amount of one—an excellent approximation, but certainly not the object of our interest.

Our goal is to show that while LM-TSP is as hard as TSP itself in terms of inapproximability, LM-TSP admits better approximation algorithms than TSP whenever input instances are either guaranteed to be metric or to be near-metric at a certain (generous, but not arbitrary) relaxation factor.

Let Δ-TSP denote metric TSP, and, for all $\beta \geq \frac{1}{2}$, let Δ_β-TSP denote the special case of TSP where all instances satisfy the β-triangle inequality

$$c(\{x,z\}) \leq \beta \cdot \big(c(\{x,y\}) + c(\{y,z\})\big)$$

for all vertices x, y, and z. If $\frac{1}{2} \leq \beta < 1$, we call this the *strengthened* triangle inequality; and if $\beta > 1$, we call it the *relaxed* triangle inequality.

For an optimization problem U, we denote our local-modification variant of U by LM-U. For the aforementioned special cases of TSP, we regard it as a local modification to change the cost of exactly one edge.

Our main results are as follows:

(i) It is well-known that TSP is not approximable in polynomial time with a polynomial approximation ratio (unless $P = NP$). We show that this holds for LM-TSP, too. Thus, in terms of a worst-case analysis, LM-TSP is as hard as TSP, and we do not have anything to gain from knowing an optimal solution to a close problem instance. By parameterizing TSP with respect to the β-triangle inequality [1–3,5,6] and by introducing the concept of stability of approximation [14,6], it was shown that TSP is not as hard as it may look like in the light of worst-case analyses. For any $\beta > \frac{1}{2}$, we have a constant polynomial-time approximation ratio, depending on β only. Böckenhauer and Seibert [7] proved that Δ_β-TSP is APX-hard for every $\beta > \frac{1}{2}$ (note that for $\beta = \frac{1}{2}$, the problem becomes trivially solvable in polynomial time). Here, we prove that LM-Δ_β-TSP is NP-hard for every $\beta > \frac{1}{2}$. This implies in particular that LM-Δ-TSP, too, is NP-hard. We conjecture that, for $\beta \geq 1$, this problem is also APX-hard, which, so far, we have been unable to prove and thus leave as an open research problem.

(ii) For many years, Christofides' algorithm [8] with its approximation ratio of 1.5 has been the best known approximation algorithm for attacking Δ-TSP. It remains a grand challenge to improve on Christofides' algorithm. We will show that, intriguingly enough, LM-Δ-TSP admits an efficient 1.4-approximation algorithm. This result can be generalized to LM-Δ_β-TSP , and the resulting approximation guarantee beats all previously-known approximation algorithms for Δ_β-TSP for all $1 < \beta < 3.34899$, which includes the practically most relevant TSP instances. Furthermore, for $\frac{1}{2} \leq \beta < 1$, we show how to obtain an approximation ratio arbitrarily close to 1, for sufficiently large input graphs.

So, on the one hand, additional information about an optimal solution to a related input instance may be useful to some extent, and on the other hand, the local-modification problem variant may remain exactly as hard as the original problem. Yet, the final aim of our paper is to call forth the investigation of the hardness of local-modification optimization problems in order to develop approaches to handle situations where multiple (and, potentially, dynamically determined) local modifications may arise.

The paper is subdivided into four sections. In Section 2, we will present our hardness results. In Section 3, we will present a 1.4-approximation algorithm for the local-modification metric TSP, Section 4 is devoted to approximability results for the case of the relaxed triangle inequality, and Section 5 contains our approximation results for the case of the sharpened triangle inequality.

2. Hardness Results

We start off with a formal definition of TSP and its local-modification variants.

Definition 1. Let $G = (V, E, c)$ be a weighted complete graph, and let $\beta \geq \frac{1}{2}$ be a real value. We say that G obeys the Δ_β-inequality iff for all vertices x, y, $z \in V$, we have

$$c(\{x,z\}) \leq \beta \cdot \big(c(\{x,y\}) + c(\{y,z\})\big) \quad . \qquad (1)$$

By TSP, we denote the following optimization problem. For a given weighted complete graph $G = (V, E, c)$, find a minimum cost Hamiltonian cycle, *i.e.*, a tour on all vertices of cost

$$OT_G := \min\Big\{ \sum_{e \in C'} c(e) \;\Big|\; (V, C') \text{ is a Hamiltonian cycle} \Big\}.$$

Restricting, for some value of β, the set of admissible input instances to those which obey the Δ_β-inequality yields the problem Δ_β-TSP. Besides, we define Δ-TSP $:= \Delta_1$-TSP.

Definition 2. Let $U \in \{\text{TSP}, \Delta\text{-TSP}, \Delta_\beta\text{-TSP}\}$. The problem LM-$U$ is defined as follows.
Input:
- two complete weighted graphs $G_O = (V, E, c_O)$, $G_N = (V, E, c_N)$ such that G_O and G_N are both admissible inputs for U and such that c_O and c_N coincide, except for one edge;
- a Hamiltonian cycle $(V, \overline{C})$ such that $\sum_{e \in \overline{C}} c_O(e) = OT_{G_O}$.

Problem: Find a Hamiltonian cycle (V, C) that minimizes $\sum_{e \in C} c_N(e)$.

Before presenting approximation algorithms for LM-Δ-TSP, we start by proving some hardness results.

First, we will show that LM-TSP is as hard to approximate as "normal" (i. e., unaltered) TSP.

Theorem 1 *There is no polynomial-time $p(n)$-approximation algorithm for* LM-TSP *for any polynomial p (unless $P = NP$).*

Proof: We will give a reduction from the Hamiltonian cycle problem (HC): Given an undirected, unweighted graph G, decide whether G contains a Hamiltonian cycle or not. Let $G = (V, E)$ be an input instance for HC where $V = \{v_1, \ldots, v_n\}$.

In order to construct an input instance $(G_O, G_N, \overline{C})$ for LM-TSP, we employ a graph construction due to Papadimitriou and Steiglitz [18], who used the same construction in order to give examples of TSP instances which are hard for local search strategies: For each vertex v_i, we construct a so-called diamond graph D_i as shown in Figure 1 (a). We will refer to the corner vertices N_i, S_i, W_i, and E_i of D_i as to the north, south, west, and east vertex of D_i, respectively.

The main property of the diamond graph, which we will employ in our reduction, is the following. Assuming that a path may only enter or leave a diamond D_i at a corner vertex, there are only two distinct possibilities to traverse all vertices of D_i: either from west to east, as shown in Figure 1 (b), or from north to south, as shown in Figure 1 (c).

These diamonds are now connected as shown in Fig-

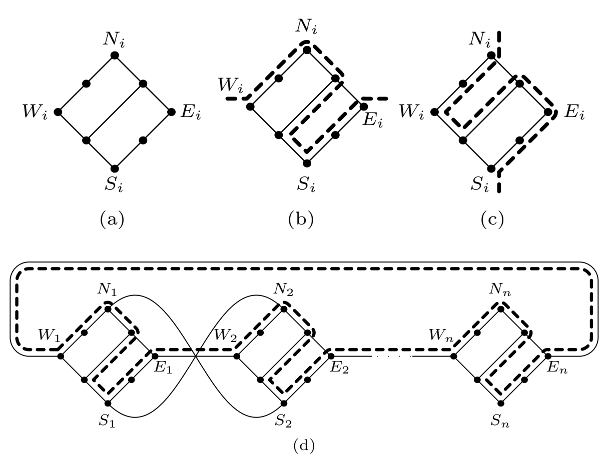

Fig. 1. The diamond construction in the proof of Theorem 1.

ure 1 (d). The edge costs in G_O are set as follows. Let $M := n \cdot 2^n + 1$. All diamond edges shown in Figure 1 (a) and the east-west-connections from E_i to W_{i+1} and from E_n to W_1 as shown in Figure 1 (d) are assigned a cost of 1 each. The north-south-edges $\{N_i, S_j\}$ are assigned a cost of 1 whenever $\{v_i, v_j\} \in E$ and a cost of M otherwise. All other edges receive a cost of M each. The choice of these edge costs assures that any Hamiltonian path in G_0 traverses the diamonds from north to south (as shown in Figure 1 (b)) or from east to west (as shown in Figure 1 (c)), unless it uses at least one expensive edge.

In G_N, the cost of the edge $\{E_n, W_1\}$ is changed from 1 to M. The given optimal Hamiltonian cycle $\overline{C}$ is the one shown in Figure 1 (d). This optimal solution for G_O has a cost of $8n$.

It is easy to see that if there is a Hamiltonian cycle H' in G, a corresponding Hamiltonian cycle H in G can traverse all diamonds in north-south direction. Hence, $c_N(H) = 8n$. All Hamiltonian cycles in G_N that do not correspond (in this way) to Hamiltonian cycles in G cost at least $M + 8n - 1$. Thus, the approximation ratio of any non-optimal solution is at least as bad as $1 + 2^{n-3}$. For the detailed description of similar diamond graph constructions, also see, for example, [15].

Now, we will show that LM-Δ-TSP remains a hard problem for any $\beta > \frac{1}{2}$.

Theorem 2 LM-Δ_β-TSP *is NP-hard for any* $\beta > \frac{1}{2}$.

Proof: We will use a reduction from the restricted Hamiltonian cycle problem (RHC). The objective in RHC is, given an unweighted, undirected graph G and a Hamiltonian path P in G which cannot be trivially extended to a Hamiltonian cycle by joining its endpoints, to decide whether a Hamiltonian cycle in G exists. This problem is well-known to be NP-complete (see, for example, [15]).

The reduction uses an idea analogous to the standard reduction from the Hamiltonian cycle problem to TSP: Let (G, P) be an instance of RHC where $G = (V, E)$, $V = \{v_1, \dots, v_n\}$, and $P = (v_1, \dots, v_n)$. From this, we construct an instance $(G_O, G_N, \overline{C})$ of LM-Δ_β-TSP as follows: Let $G_O = (V, \tilde{E}, c_O)$ and $G_N = (V, \tilde{E}, c_N)$ where $(V, \tilde{E})$ is a complete graph, $c_O(e) = 1$ for all $e \in E \cup \{\{v_n, v_1\}\}$ and $c_O(e) = 2\beta$ otherwise, and $c_N(\{v_n, v_1\}) = 2\beta$. Let $\overline{C} = (v_1, v_2, \dots, v_n, v_1)$. Clearly, this reduction can be done in polynomial time, and it is easy to see that there is a Hamiltonian cycle in G iff there is a Hamiltonian cycle of cost n in G_N.

3. The Metric Case

In what follows, we will show that LM-Δ-TSP admits a $\frac{7}{5}$-approximation, which beats the naïve approach of using Christofides' algorithm (which would yield a $\frac{3}{2}$-approximation), whereby the input cycle $(V, \overline{C})$ would be ignored altogether.

Theorem 3 *There is a* 1.4*-approximation algorithm for* LM-Δ-TSP.

In order to prove Theorem 3, we will need the following few lemmas. Our crucial observation is that in a metric graph, all of the neighboring edges of short edges can only be modified by small amounts.

Lemma 4 *Let* $G_1 = (V, E, c_1)$ *and* $G_2 = (V, E, c_2)$ *be metric graphs such that* c_1 *and* c_2 *coincide, except for one edge* $e \in E$. *Then, every edge adjacent to* e *has a cost of at least* $\frac{1}{2}|c_1(e) - c_2(e)|$.

Proof: We set $\{a, a'\} := \{c_1(e), c_2(e)\}$ such that $a' > a$ and $\delta := a' - a$. Let $f \in E$ be any edge adjacent to e, and for any such f, let $f' \in E$ be the one edge that is adjacent to both e and f. Then, by the triangle inequality, we have:

$$a' \leq c(f) + c(f') \qquad c(f') \leq c(f) + a$$

and hence $a' - a \leq 2c(f)$.

We will have to distinguish two cases. Either, an edge becomes more expensive, or it becomes less expensive. In either case, our strategy is to compare the input solution (to the old problem instance) with an approximate solution (to the new problem instance).

Let us start with the latter case.

Lemma 5 *Let* $(G_O, G_N, \overline{C})$ *be an admissible input for* LM-Δ-TSP *such that* $\delta := c_O(e) - c_N(e) > 0$ *for the edge* e. *If* $\frac{\delta}{OT_{G_N}} \leq \frac{2}{5}$, *it is a* $\frac{7}{5}$*-approximation to output the feasible solution* $C := \overline{C}$ *for* LM-Δ-TSP.

Proof:

$$\frac{c_N(\overline{C})}{OT_{G_N}} \leq \frac{c_O(\overline{C})}{OT_{G_N}} = \frac{OT_{G_O}}{OT_{G_N}} \leq \frac{OT_{G_N} + \delta}{OT_{G_N}} = 1 + \frac{\delta}{OT_{G_N}} \leq 1 + \frac{2}{5} = \frac{7}{5}$$

Lemma 6 *Let* $(G_O, G_N, \overline{C})$ *be an admissible input for* LM-Δ-TSP *such that* $\delta := c_O(e) - c_N(e) > 0$ *for the edge* e. *If* $\frac{\delta}{OT_{G_N}} \geq \frac{2}{5}$, *there is a* $\frac{7}{5}$*-approximation for* LM-Δ-TSP.

Proof: We may assume that optimal TSP tours in G_N use the edge e. For if they did not, $\overline{C}$ would already constitute an optimal solution. Fix one such optimal tour C_{OPT} in G_N. In C_{OPT}, e is adjacent to two edges f and f'. Let v be the vertex incident with f, but not with e, and let v' be the vertex incident with f', but not with e. By P, denote the path from v to v' in C_{OPT} that does *not* involve e.

Consider the following algorithm: For every pair $\tilde{f}$, $\tilde{f}'$ of disjoint edges, both of which are adjacent to e, compute an approximate solution to the TSP path problem on the subgraph of G_N induced by the vertex set $V \setminus e$ (*i.e.*, without *two* vertices) with start vertex $\tilde{v}$ and end vertex $\tilde{v}'$ where $\{\tilde{v}\} = \tilde{f} \setminus e$ and $\{\tilde{v}'\} = \tilde{f}' \setminus e$. It is known [11,13] that this can be done with an approximation guarantee of $\frac{5}{3}$. Each of these paths is augmented by $\tilde{f}$, e, and $\tilde{f}'$ so as to yield a TSP tour. The algorithm concludes by outputting the least expensive of all of these tours.

Note that since *all* pairs $\tilde{f}$, $\tilde{f}'$ are taken into account, one of the considered tours uses exactly those edges $\tilde{f} = f$, $\tilde{f}' = f'$ that C_{OPT} uses. This is why the algorithm outputs a tour of cost at most

$$\begin{aligned} c(f) + c(f') + c_N(e) + \frac{5}{3}c(P) \\ &= \big(OT_{G_N} - c(P)\big) + \frac{5}{3}c(P) \\ &= OT_{G_N} + \frac{2}{3}c(P) \end{aligned}$$

(where c is short-hand notation for c_N wherever c_O and c_N coincide) and thus achieves an approximation guarantee of

$$1 + \frac{2}{3} \cdot \frac{c(P)}{OT_{G_N}} \quad .$$

Since by Lemma 4, $\min\{c(f), c(f')\} \geq \frac{\delta}{2}$ for $i \in \{1, 2\}$, we have $OT_{G_N} - c(P) \geq \delta$ and hence:

$$\frac{c(P)}{OT_{G_N}} \leq 1 - \frac{\delta}{OT_{G_N}} \leq \frac{3}{5} \quad .$$

So, we obtain an overall approximation guarantee of $1 + \frac{2}{5} = \frac{7}{5}$.

Corollary 7 *There is a $\frac{7}{5}$-approximation algorithm for the subproblem of* LM-Δ-TSP *where edges may only become less expensive.*

Proof: Compute, as laid out in Lemma 6, an approximate solution to LM-Δ-TSP and compare it with the input solution $\overline{C}$. Output the less expensive of the two solutions. Depending on whether the value of $\frac{\delta}{OT_{G_N}}$ (where $\delta := c_O(e) - c_N(e) > 0$) is less or greater than $\frac{2}{5}$ (which we cannot necessarily tell), one of the considered two feasible solutions is a $\frac{7}{5}$-approximation.

We will now turn to the case where an edge becomes more expensive. We can state a lemma akin to Lemma 5, but notice that by reusing a formerly optimal solution, we incur a certain extra cost.

Lemma 8 *Let $(G_O, G_N, \overline{C})$ be an admissible input for* LM-Δ-TSP *such that $\delta := c_N(e) - c_O(e) > 0$ for the edge e. If $\frac{\delta}{OT_{G_N}} \leq \frac{2}{5}$, it is a $\frac{7}{5}$-approximation to output the feasible solution $C := \overline{C}$ for* LM-Δ-TSP.

Proof:

$$\begin{aligned} \frac{c_N(\overline{C})}{OT_{G_N}} &\leq \frac{c_O(\overline{C}) + \delta}{OT_{G_N}} = \frac{OT_{G_O} + \delta}{OT_{G_N}} \leq \frac{OT_{G_N} + \delta}{OT_{G_N}} \\ &= 1 + \frac{\delta}{OT_{G_N}} \leq 1 + \frac{2}{5} = \frac{7}{5} \end{aligned}$$

When computing an approximate solution, things become slightly different from what they used to be like in Lemma 6: We may assume that e *used* to be a part of $\overline{C}$ and that a new solution should no longer use it. Instead, it will use two edges f and f' such that f and f' are non-disjoint and both incident with the same vertex of e. This pair may be chosen at either end-point of e, a choice which is completely *arbitrary.*

We conjecture that, if an improvement of the approximation guarantee is possible, this is precisely the point where to start at.

Lemma 9 *Let $(G_O, G_N, \overline{C})$ be an admissible input for* LM-Δ-TSP *such that $\delta := c_N(e) - c_O(e) > 0$ for the edge e. If $\frac{\delta}{OT_{G_N}} \geq \frac{2}{5}$, there is a $\frac{7}{5}$-approximation for* LM-Δ-TSP.

Proof: We may assume that optimal TSP tours in G_N do not use the edge e. For if they did, $\overline{C}$ would already constitute an optimal solution. Fix one such optimal tour C_{OPT}, and fix one vertex w incident with e. In C_{OPT}, w is incident with two edges f and f'. Let v be the vertex incident with f, but not with e, and let v' be the vertex incident with f', but not with e. By P, denote the path from v to v' in C_{OPT} that does *not* involve w.

Consider the following algorithm: For every pair $\tilde{f}$, $\tilde{f}'$ of edges incident with w, compute an approximate solution to the TSP path problem on the subgraph of G_2 induced by the vertex set $V \setminus \{w\}$ with start vertex

$\tilde{v}$ and end vertex $\tilde{v}'$ where $\{\tilde{v}\} = \tilde{f} \setminus e$ and $\{\tilde{v}'\} = \tilde{f}' \setminus e$. It is known [11,13] that this can be done with an approximation guarantee of $\frac{5}{3}$. Each of these paths is augmented by $\tilde{f}$ and $\tilde{f}'$ so as to yield a TSP tour. The algorithm concludes by outputting the least expensive of all of these tours.

Note that since *all* pairs $\tilde{f}$, $\tilde{f}'$ are taken into account, one of the considered tours uses exactly those edges $\tilde{f} = f$, $\tilde{f}' = f'$ that C_{OPT} uses. This is why the algorithm outputs a tour of cost at most

$$c(f) + c(f') + \frac{5}{3}c(P) = \big(OT_{G_N} - c(P)\big) + \frac{5}{3}c(P) = OT_{G_N} + \frac{2}{3}c(P),$$

just as in the proof of Lemma 6.

Using the same arguments as in the proof of Corollary 7, the preceding lemma yields the following corollary.

Corollary 10 *There is a $\frac{7}{5}$-approximation algorithm for the subproblem of* LM-Δ-TSP *where edges may only become more expensive.*

4. The Near-Metric Case

The algorithm outlined in Lemma 6 can be generalized to graphs which are not necessarily metric, but only near-metric, *i.e.*, where the metricity constraint is relaxed by a factor of $\beta > 1$. Since it will be useful later, let us pay extra attention to the fact that input instances for all the problems from Definition 2 contain two distinct graphs, potentially obeying relaxed triangle inequalities according to different values of β.

Notice that the parameter β need not be greater for the graph with the costlier edge. Under some circumstances, it might even decrease when we modify the cost of a single edge. In the following generalization of Lemma 4, the convention is therefore that c_1 is the cost function of the less expensive graph, c_2 that of the more expensive one, and both c_i obey the Δ_{β_i}-inequality, $i \in \{1,2\}$.

Lemma 11 *Let $G_1 = (V, E, c_1)$ and $G_2 = (V, E, c_2)$ be graphs such that c_i obeys the Δ_{β_i}-inequality for $i \in \{1,2\}$ and some values $\beta_1, \beta_2 \geq 1$ and such that c_1 and c_2 coincide, except for one edge $e \in E$. By convention, let $c_1(e) \leq c_2(e)$. Then, every edge adjacent to e has a cost of at least $\frac{c_2(e) - \beta_1\beta_2 c_1(e)}{\beta_1\beta_2 + \beta_2}$.*

Proof: We set $a := c_1(e)$ and $a' := c_2(e)$. Let $f \in E$ be any edge adjacent to e, and for any such f, let $f' \in E$ be the one edge that is adjacent to both e and f. Then, by the relaxed triangle inequality, we have:

$$a' \leq \beta_2 \cdot (c(f) + c(f')) \quad c(f') \leq \beta_1 \cdot (a + c(f))$$

and hence

$$c(f) \geq \frac{a' - \beta_1\beta_2 a}{\beta_2 + \beta_1\beta_2}.$$

Note that for relatively small changes, the value $c_2(e) - \beta_1\beta_2 c_1(e)$ may well be non-positive, rendering Lemma 11 trivial in such a case.

The algorithm from Lemmas 6 and 8 should be adjusted to accommodate for the relaxation of the triangle inequality. More precisely, in order to find a Hamiltonian path between a given pair of vertices in a β-metric graph, we will employ the algorithm by Forlizzi et al. [9], a variation of the path-matching Christofides algorithm (PMCA, see [6]) for the path version of near-metric TSP, which yields an approximation guarantee of $\frac{5}{3}\beta^2$. This gives us Algorithm 1.

Algorithm 1

Input: *An instance $(G_O, G_N, \overline{C})$ of* LM-Δ_β-TSP *where $\beta > 1$, $G_O = (V, E, c_O)$ and $G_N = (V, E, c_N)$.*

(1) Let $e \in E$ be the edge where $c_O(e) \neq c_N(e)$. Let $\mathcal{E}$ be the set of all unordered pairs $\{f, f'\} \subseteq E$ where $f \neq f'$ are edges adjacent to e such that if $c_O(e) < c_N(e)$: $f \cap f' \cap e$ is a singleton; and if $c_O(e) > c_N(e)$: $f \cap f' = \emptyset$.

(2) For all $\{f, f'\} \in \mathcal{E}$, compute a Hamiltonian path between the two vertices from $(f \cup f') \setminus e$ on the graph $G \setminus (e \cap (f \cup f'))$, using the PMCA path variant by Forlizzi et al. [9]. Augment this path by edges f, f', and, if $c_O(e) > c_N(e)$, edge e to obtain the cycle $C_{\{f,f'\}}$.

(3) Let C be the least expensive of the cycles in the set $\{\overline{C}\} \cup \{C_{\{f,f'\}} \mid \{f, f'\} \in \mathcal{E}\}$.

Output: *The Hamiltonian cycle C.*

Lemma 12 *Algorithm 1 achieves an approximation guarantee of*

$$\beta_{\mathsf{L}}\beta_{\mathsf{H}} \cdot \frac{15\beta_{\mathsf{L}}^2 + 5\beta_{\mathsf{L}} - 6}{10\beta_{\mathsf{L}}^2 + 3\beta_{\mathsf{L}}\beta_{\mathsf{H}} + 3\beta_{\mathsf{H}} - 6} \tag{2}$$

for input graph pairs (G_O, G_N) such that G_O obeys the Δ_{β_O}-inequality and G_N obeys the Δ_{β_N}-inequality and where $\beta_{\mathsf{L}} := \min\{\beta_O, \beta_N\} \geq 1$ and $\beta_{\mathsf{H}} := \max\{\beta_O, \beta_N\} > 1$.

Proof: Adhering to the convention of Lemma 11, set $\{c_1, c_2\} = \{c_O, c_N\}$ such that $c_1(e) \leq c_2(e)$ for all edges $e \in E$. In other words, we have $c_2 = c_N$ if an edge becomes more expensive and $c_1 = c_N$ otherwise.

We may assume that optimal TSP tours in $G_N = (V, E, c_N)$ use the edge e iff $c_N = c_1$; otherwise, $\overline{C}$ is an optimal solution, and we are done. Fix one such optimal tour C_{OPT} in G_N, and let $\{f, f'\} \in \mathcal{E}$ be such that C_{OPT} uses both f and f'. By P, denote the path that results from C_{OPT} by removing edges f, f', and, potentially, e. Set

$$\alpha := \frac{C(P)}{OT_{G_N}} \quad \text{and let, for brevity,}$$
$$\vartheta := \beta_{\mathsf{L}}\beta_{\mathsf{H}} \cdot \frac{15\beta_{\mathsf{L}}^2 + 5\beta_{\mathsf{L}} - 6}{10\beta_{\mathsf{L}}^2 + 3\beta_{\mathsf{L}}\beta_{\mathsf{H}} + 3\beta_{\mathsf{H}} - 6}$$

denote the approximation guarantee claimed in (2). In terms of α, Algorithm 1 always achieves an approximation guarantee of

$$1 - \alpha + \frac{5}{3}\beta_L 2\alpha,$$

even if we did not have $\overline{C}$ at our disposal. Here, the term $1 - \alpha$ corresponds to the edges f, f' and (potentially) e, which are chosen optimally, and the term $\frac{5}{3}\beta_L 2\alpha$ corresponds to the approximation of the path P.

(Note that the strategy to approximate P may rely on the $\Delta_{\beta_{\mathsf{L}}}$ inequality, *i.e.*, the less relaxed one of the two because this strategy removes the edge e from the graph.) Hence, unless

$$\alpha > \frac{\vartheta - 1}{\frac{5}{3}\beta_{\mathsf{L}}^2 - 1} \quad , \tag{3}$$

we are done. Let us therefore assume that (3) holds. By Lemma 11, we have

$$\begin{aligned}\min\{c(f), c(f')\} &\geq \frac{c_2(e) - \beta_1\beta_2 c_1(e)}{\beta_1\beta_2 + \beta_2}\\ &\geq \frac{c_2(e) - \beta_{\mathsf{L}}\beta_{\mathsf{H}} c_1(e)}{\beta_{\mathsf{L}}\beta_{\mathsf{H}} + \beta_{\mathsf{H}}}\end{aligned}$$

and hence

$$1 - \alpha \geq \frac{2 \cdot (c_2(e) - \beta_{\mathsf{L}}\beta_{\mathsf{H}} c_1(e))}{OT_{G_N} \cdot (\beta_{\mathsf{L}}\beta_{\mathsf{H}} + \beta_{\mathsf{H}})}.$$

Putting this together with (3), we know that

$$\frac{\vartheta - 1}{\frac{5}{3}\beta_{\mathsf{L}}^2 - 1} \leq 1 - \frac{2 \cdot (c_2(e) - \beta_{\mathsf{L}}\beta_{\mathsf{H}} c_1(e))}{OT_{G_N} \cdot (\beta_{\mathsf{L}}\beta_{\mathsf{H}} + \beta_{\mathsf{H}})},$$

which yields

$$\frac{c_2(e) - \beta_{\mathsf{L}}\beta_{\mathsf{H}} c_1(e)}{OT_{G_N}} \leq \frac{\beta_{\mathsf{L}}\beta_{\mathsf{H}} + \beta_{\mathsf{H}}}{2} - \frac{(\vartheta - 1) \cdot (\beta_{\mathsf{L}}\beta_{\mathsf{H}} + \beta_{\mathsf{H}})}{\frac{10}{3}\beta_{\mathsf{L}}^2 - 2}.$$

By adding $(\beta_{\mathsf{L}}\beta_{\mathsf{H}} - 1)\frac{c_1(e)}{OT_{G_N}}$ to both sides, we are given:

$$\frac{c_2(e) - c_1(e)}{OT_{G_N}} \leq \frac{\beta_{\mathsf{L}}\beta_{\mathsf{H}} + \beta_{\mathsf{H}}}{2} - \frac{(\vartheta - 1) \cdot (\beta_{\mathsf{L}}\beta_{\mathsf{H}} + \beta_{\mathsf{H}})}{\frac{10}{3}\beta_{\mathsf{L}}^2 - 2} + (\beta_{\mathsf{L}}\beta_{\mathsf{H}} - 1) \cdot \underbrace{\frac{c_1(e)}{OT_{G_N}}}_{\leq 1}$$

and thus, substituting the value (2) for ϑ,

$$\begin{aligned}\frac{c_2(e) - c_1(e)}{OT_{G_N}} &\leq \frac{3}{2}\beta_{\mathsf{L}}\beta_{\mathsf{H}} + \frac{1}{2}\beta_{\mathsf{H}} - 1 - \frac{(\vartheta - 1) \cdot (\beta_{\mathsf{L}}\beta_{\mathsf{H}} + \beta_{\mathsf{H}})}{\frac{10}{3}\beta_{\mathsf{L}}^2 - 2}\\ &= \frac{3}{2}\beta_{\mathsf{L}}\beta_{\mathsf{H}} + \frac{1}{2}\beta_{\mathsf{H}} - 1 - \frac{(\beta_{\mathsf{L}}\beta_{\mathsf{H}} \cdot \frac{15\beta_{\mathsf{L}}^2 + 5\beta_{\mathsf{L}} - 6}{10\beta_{\mathsf{L}}^2 + 3\beta_{\mathsf{L}}\beta_{\mathsf{H}} + 3\beta_{\mathsf{H}} - 6} - 1)(\beta_{\mathsf{L}}\beta_{\mathsf{H}} + \beta_{\mathsf{H}})}{\frac{10}{3}\beta_{\mathsf{L}}^2 - 2}\end{aligned}$$

(tedious calculations) $= \cdots$

$$= \beta_{\mathsf{L}}\beta_{\mathsf{H}} \cdot \frac{15\beta_{\mathsf{L}}^2 + 5\beta_{\mathsf{L}} - 6}{10\beta_{\mathsf{L}}^2 + 3\beta_{\mathsf{L}}\beta_{\mathsf{H}} + 3\beta_{\mathsf{H}} - 6} - 1 = \vartheta - 1.$$

Since, by the same reasoning as that of Lemmas 5 and 8, reusing the input optimal solution $\overline{C}$ inflicts a deviation from the new optimum by at most $c_2(e) - c_1(e) \leq (\vartheta - 1) \cdot OT_{G_N}$, Algorithm 1 is a ϑ-approximation algorithm.

Hence, whenever the β values of G_O and G_N coincide, we have Theorem 13.

Theorem 13 *There is a (polynomial-time)* $\beta^2 \cdot \dfrac{15\beta^2 + 5\beta - 6}{13\beta^2 + 3\beta - 6}$*-approximation algorithm for* LM-Δ_β-TSP *for* $\beta > 1$.

Interestingly, Algorithm 1 achieves a better approximation guarantee not just than PMCA [6], but also than Bender's and Chekuri's 4β-approximation algorithm [3] for the most practically relevant values of β. The turning point is about at $\beta^* \approx 3.34899$. More to the point,

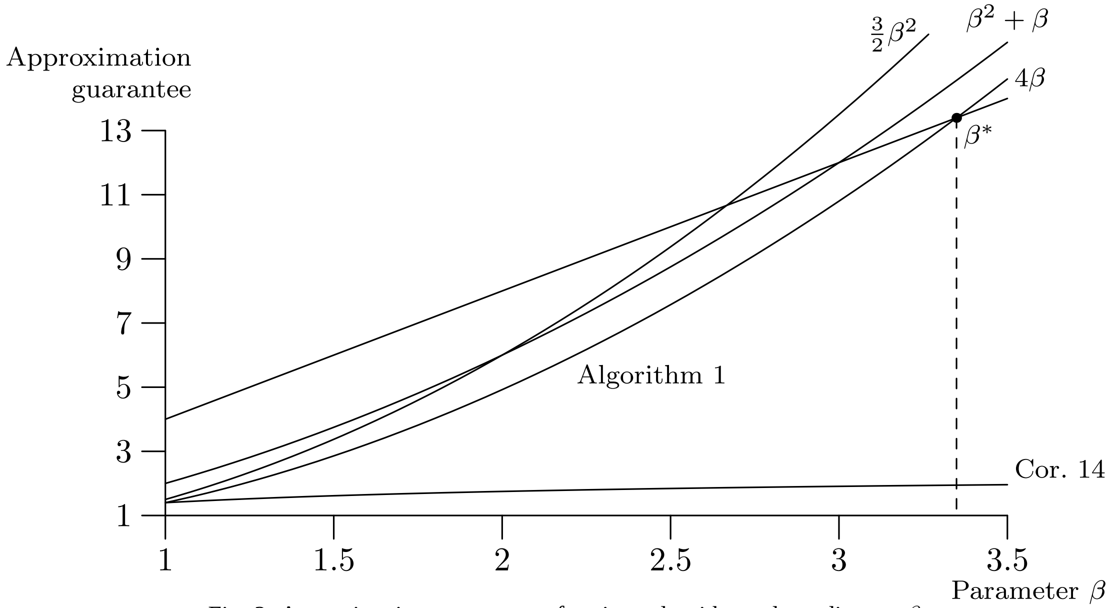

Fig. 2. Approximation guarantees of various algorithms, depending on β

Andreae's $(\beta^2+\beta)$-approximation [1], which performs better than 4β only when $\beta < 3$, always performs worse than Algorithm 1 in the interval $\beta \in (1, \beta^*)$. These observations are illustrated in Figure 2.

Another practical special case is that where $\beta_{\mathrm{L}} = 1$, *i.e.*, where we start with a metric graph, but changing the cost of an edge will violate the Δ-inequality.

Corollary 14 LM-Δ_β-TSP *for* $\beta > 1$, *restricted to those inputs where* G_O *is metric, admits a* $\frac{7\beta}{2+3\beta}$-*approximation.*

5. The Super-Metric Case

We will now deal with the case of *super-metric* graphs, i. e. with graphs satisfying the Δ_β-inequality for some $\beta < 1$. Please note that $\beta \geq \frac{1}{2}$ holds in any case where $\beta = \frac{1}{2}$ corresponds to the trivial case where all edge costs are equal. Thus, we will assume $\frac{1}{2} < \beta < 1$ for the remainder of this section. As it turns out, LM-Δ_β-TSP is fairly easy for super-metric graphs. In this section, we will show that even the conceivably most naïve algorithm for LM-Δ_β-TSP on super-metric graphs is a PTAS.

First of all, we note that, for super-metric graphs, there is a bound on the ratio of the maximal and minimal edge costs.

Lemma 15 ([5]) *Let G be a graph which obeys the Δ_β-inequality for some $\beta < 1$. Let $c_{\max}$ and $c_{\min}$ denote the cost of its most and least expensive edge, respectively. Then,*

$$\frac{c_{\max}}{c_{\min}} \leq \frac{2\beta^2}{1-\beta} \quad . \tag{4}$$

Proof: To be found in [5], Lemma 2 (b).

Moreover, neighboring edges in super-metric graphs never differ by a factor of more than $\frac{1}{1-\beta}$ [5]. Therefore, the maximal edge costs in the two graphs of a LM-Δ_β-TSP input instance are similarly related.

Lemma 16 *Let G_O and G_N be two weighted graphs such that G_O obeys the Δ_{β_O}-inequality and G_N obeys the Δ_{β_N}-inequality where $\max\{\beta_O, \beta_N\} < 1$. For $i \in \{O, N\}$, let $c_{\max,i}$ and $c_{\min,i}$ denote the maximal and minimal cost of an edge in G_i, respectively. Let the edge costs in G_O and G_N agree except for one edge. Then,*

$$\begin{aligned} c_{\max,N} &\leq \frac{1}{1-\beta_N} c_{\max,O}, \\ c_{\min,N} &\leq \frac{1}{1-\beta_O} c_{\min,O}, \\ c_{\max,O} &\leq \frac{1}{1-\beta_O} c_{\max,N}, \\ \text{and} \quad c_{\min,O} &\leq \frac{1}{1-\beta_N} c_{\min,N}. \end{aligned} \tag{5}$$

Proof: Let e be the edge such that $c_O(e) \neq c_N(e)$. Since all the neighbors of e have the same cost in G_O as in G_N and since these costs are bounded by $c_{\max,O}$, we have $c_N(e) \le \frac{1}{1-\beta_N} c_{\max,O}$. Hence,

$$c_{\max,N} \le \max\{c_N(e), c_{\max,O}\} \le \frac{1}{1-\beta_N} c_{\max,O}.$$

Likewise, $c_O(e) \ge (1-\beta_O)c_{\min,N}$, and hence,

$$c_{\min,O} \ge \min\{c_O(e), c_{\min,N}\} \ge (1-\beta_O)c_{\min,N}.$$

The two remaining inequalities are symmetric.

Theorem 17 *Let $(G_O, G_N, \overline{C})$ be an input instance of* LM-Δ_β-TSP *such that G_O obeys the Δ_{β_O}-inequality and G_N obeys the Δ_{β_N}-inequality and where $\beta_{\mathsf{L}} := \min\{\beta_O, \beta_N\}$ and $\beta_{\mathsf{H}} := \max\{\beta_O, \beta_N\} < 1$. Then, it is a $(1 + \frac{2\beta_{\mathsf{L}}^2 - (1-\beta_{\mathsf{L}})(1-\beta_{\mathsf{H}})}{(1-\beta_{\mathsf{L}})(1-\beta_{\mathsf{H}})|V|})$-approximation to simply output $\overline{C}$.*

Proof: It is straightforward to verify that

$$c_{\min,i} \le \frac{OT_{G_i}}{|V|}. \tag{6}$$

Suppose that e is the edge whose cost is altered and suppose that $c_N(e) > c_O(e)$. W.l.o.g., assume that e is a part of $\overline{C}$. (If it is not, $\overline{C}$ is already an optimal tour in G_N.) Then,

$$OT_{G_O} \le OT_{G_N} \tag{7}$$

and

$$\begin{aligned} c_N(\overline{C}) &= OT_{G_O} + c_N(e) - c_O(e) \\ &\le OT_{G_O} + c_{\max,N} - c_{\min,O} \qquad (8) \\ &\overset{(4)}{\le} OT_{G_O} + \frac{1}{1-\beta_N} c_{\max,O} - c_{\min,O} \\ &\overset{(3)}{\le} OT_{G_O} + \frac{1}{1-\beta_N} \cdot \frac{2\beta_O^2}{1-\beta_O} \\ &\quad \cdot c_{\min,O} - c_{\min,O}. \qquad (9) \end{aligned}$$

But we also have

$$\begin{aligned} c_N(\overline{C}) &\overset{(7)}{\le} OT_{G_O} + c_{\max,N} - c_{\min,O} \\ &\overset{(3)}{\le} OT_{G_O} + \frac{2\beta_N^2}{1-\beta_N} c_{\min,N} - c_{\min,O} \\ &\overset{(4)}{\le} OT_{G_O} + \frac{2\beta_N^2}{1-\beta_N} \cdot \frac{1}{1-\beta_O} \\ &\quad \cdot c_{\min,O} - c_{\min,O}. \end{aligned} \tag{10}$$

The combination of (9) and (10) yields

$$\begin{aligned} c_N(\overline{C}) &\le OT_{G_O} + \frac{\min\{2\beta_O^2, 2\beta_N^2\}}{(1-\beta_N)\cdot(1-\beta_O)} \\ &\quad \cdot c_{\min,O} - c_{\min,O} \\ &= OT_{G_O} + \Big(\frac{2\beta_L^2}{(1-\beta_L)\cdot(1-\beta_H)} - 1\Big) \\ &\quad \cdot c_{min,O} \\ &\overset{(5)}{\le} OT_{G_O} + \Big(\frac{2\beta_L^2}{(1-\beta_L)(1-\beta_H)} - 1\Big) \\ &\quad \cdot \frac{OT_{G_O}}{|V|} \\ &\overset{(6)}{\le} OT_{G_N} + \Big(\frac{2\beta_L^2 - (1-\beta_L)(1-\beta_H)}{(1-\beta_L)(1-\beta_H)}\Big) \\ &\quad \cdot \frac{OT_{G_N}}{|V|} \\ &= \Big(1 + \frac{2\beta_L^2 - (1-\beta_L)(1-\beta_H)}{(1-\beta_L)(1-\beta_H)|V|}\Big) \cdot OT_{G_N}. \end{aligned}$$

Now, suppose that $c_N(e) < c_O(e)$. W.l.o.g., assume that e is *not* a part of $\overline{C}$. (If it is, $\overline{C}$ is already an optimal tour in G_N.) Then,

$$\begin{aligned} c_N(\overline{C}) &= OT_{G_O} \le OT_{G_N} + c_O(e) - c_N(e) \\ &\le OT_{G_N} + c_{\max,O} - c_{\min,N} \qquad (11) \\ &\overset{(4)}{\le} OT_{G_N} + \frac{1}{1-\beta_O} c_{\max,N} - c_{\min,N} \\ &\overset{(3)}{\le} OT_{G_N} + \frac{1}{1-\beta_O} \cdot \frac{2\beta_N^2}{1-\beta_N} \\ &\quad \cdot c_{\min,N} - c_{\min,N}. \qquad (12) \end{aligned}$$

But we also have

$$\begin{aligned} c_N(\overline{C}) &\overset{(10)}{\le} OT_{G_N} + c_{\max,O} - c_{\min,N} \\ &\overset{(3)}{\le} OT_{G_N} + \frac{2\beta_O^2}{1-\beta_O} c_{\min,O} - c_{\min,N} \\ &\overset{(4)}{\le} OT_{G_N} + \frac{2\beta_O^2}{1-\beta_O} \cdot \frac{1}{1-\beta_N} \\ &\quad \cdot c_{\min,N} - c_{\min,N}. \end{aligned} \tag{13}$$

The combination of (12) and (13) yields

$$
\begin{aligned}
c_N(\overline{C}) &\leq OT_{G_N} + \frac{\min\{2\beta_N^2, 2\beta_O^2\}}{(1-\beta_O)\cdot(1-\beta_N)} \\
&\quad \cdot c_{\min,N} - c_{\min,N} \\
&= OT_{G_N} + \Big(\frac{2\beta_L^2}{(1-\beta_L)\cdot(1-\beta_H)} - 1\Big) \\
&\quad \cdot c_{min,N} \\
&\overset{(5)}{\leq} OT_{G_N} + \Big(\frac{2\beta_L^2}{(1-\beta_L)(1-\beta_H)} - 1\Big) \\
&\quad \cdot \frac{OT_{G_N}}{|V|} \\
&= \Big(1 + \frac{2\beta_L^2 - (1-\beta_L)(1-\beta_H)}{(1-\beta_L)(1-\beta_H)|V|}\Big) \cdot OT_{G_N}.
\end{aligned}
$$

So, as $|V| \to \infty$, the approximation guarantee of an algorithm which simply outputs $\overline{C}$ approaches 1. Since we can also estimate this guarantee, the following simple strategy results in a PTAS for LM-Δ_β-TSP on supermetric graphs: Count the number of vertices in the input graphs, compute the approximation guarantee (given by the formula of Theorem 17) and decide, whether it is good enough. If so, output $\overline{C}$. If not, perform exhaustive search for an optimal solution. Since this happens for finitely many inputs only, this algorithm is a PTAS.

6. Conclusion

In this work, we have introduced and successfully applied the concept of reusing optimal solutions when input instances are locally modified. In the case of metric TSP, we are able to improve on the previously-known upper bound of 1.5, as achieved by Christofides' algorithm (applied to the new instance, ignoring the given optimal solution). In the case of near-metric TSP, we have shown how to non-trivially extend our approach to the most practical values of β.

As an open problem, we state the question whether the NP-hard LM-Δ-TSP is also APX-hard.

Acknowledgements

This work was partially supported by SNF grant `200021-109252/1`, by the research project GRID.IT, funded by the Italian Ministry of Education, University and Research, and by SBF grant `C06.0108` as part of the COST 293 (GRAAL) project funded by the European Union. A part of the results was presented at the World Computer Congress 2006 [4].

References

[1] T. Andreae. On the traveling salesman problem restricted to inputs satisfying a relaxed triangle inequality. *Networks*, 38:59–67, 2001.

[2] T. Andreae and H.-J. Bandelt. Performance guarantees for approximation algorithms depending on parameterized triangle inequalities. *SIAM Journal on Discrete Mathematics*, 8:1–16, 1995.

[3] M. Bender and C. Chekuri. Performance guarantees for TSP with a parametrized triangle inequality. *Information Processing Letters*, 73:17–21, 2000.

[4] H.-J. Böckenhauer, L. Forlizzi, J. Hromkovič, J. Kneis, J. Kupke, G. Proietti, and P. Widmayer. Reusing optimal TSP solutions for locally modified input instances (extended abstract). In *Proc. of the 4th IFIP International Conference on Theoretical Computer Science*, pages 251–270, 2006.

[5] H.-J. Böckenhauer, J. Hromkovič, R. Klasing, S. Seibert, and W. Unger. Approximation algorithms for TSP with sharpened triangle inequality. *Information Processing Letters*, 75:133–138, 2000.

[6] H.-J. Böckenhauer, J. Hromkovič, R. Klasing, S. Seibert, and W. Unger. Towards the notion of stability of approximation for hard optimization tasks and the traveling salesman problem. *Theoretical Computer Science*, 285:3–24, 2002.

[7] H.-J. Böckenhauer and S. Seibert. Improved lower bounds on the approximability of the traveling salesman problem. *RAIRO Theoretical Informatics and Applications*, 34:213–255, 2000.

[8] N. Christofides. Worst-case analysis of a new heuristic for the travelling salesman problem. Technical Report 388, Graduate School of Industrial Administration, Carnegie-Mellon University, Pittsburgh, 1976.

[9] L. Forlizzi, J. Hromkovič, G. Proietti, and S. Seibert. On the stability of approximation for Hamiltonian path problems. *Algorithmic Operations Research*, 1(1): 31–45, 2006.

[10] H. Greenberg. An annotated bibliography for post-solution analysis in mixed integer and combinatorial optimization. *Advances in Computational and Stochastic Optimization, Logic Programming, and Heuristic Search*, pages 97–147, 1998. ISBN 0-7923-8078-9.

[11] N. Guttmann-Beck, R. Hassin, S. Khuller, and B. Raghavachari. Approximation algorithms with bounded performance guarantees for the clustered traveling salesman problem. *Algorithmica*, 28:422–437, 2000.

[12] S. van Hoesel and A. Wagelmans. On the complexity of postoptimality analysis of 0/1 programs. *Discrete Applied Mathematics*, 91:251–263, 1999.

[13] J. A. Hoogeveen. Analysis of Christofides' heuristic: Some paths are more difficult than cycles. *Operations Research Letters*, 10:178–193, 1978.

[14] J. Hromkovič. Stability of approximation algorithms for hard optimization problems. In *Proc. of the 26th SOFSEM*, volume 1725 of *Lecture Notes in Computer Science*, pages 29–47. Springer-Verlag, 1999.

[15] J. Hromkovič. *Algorithmics for Hard Problems. Introduction to Combinatorial Optimization, Randomization, Approximation, and Heuristics.* Springer-Verlag, 2003.

[16] M. Libura. Sensitivity analysis for minimum Hamiltonian path and traveling salesman problems. *Discrete Applied Mathematics*, 30:197–211, 1991.

[17] M. Libura, E. S. van der Poort, G. Sierksma, and J. A. A. van der Veen. Stability aspects of the traveling salesman problem based on k-best solutions. *Discrete Applied Mathematics*, 87:159–185, 1998.

[18] Ch. Papadimitriou and K. Steiglitz. Some examples of difficult traveling salesman problems. *Operations Research*, 26:434–443, 1978.

[19] Y. N. Sotskov, V. K. Leontev, and E. N. Gordeev. Some concepts of stability analysis in combinatorial optimization. *Discrete Applied Mathematics*, 58:169–190, 1995.

Received December 4, 2006; revised 17 April 2007; accepted 5 June 2007