The UPS Problem

Cristina G. Fernandes[1,⋆] and Till Nierhoff[2,⋆⋆]

[1] Departamento de Ciência da Computação
Instituto de Matemática e Estatística
Universidade de São Paulo - Brazil
cris@ime.usp.br

[2] Institut für Informatik
Humboldt-Universität zu Berlin
nierhoff@informatik.hu-berlin.de

Abstract. The UPS Problem consists of the following: given a vertex set V, vertex probabilities $(p_v)_{v \in V}$, and distances $l : V^2 \to R^+$ that satisfy the triangle inequality, find a Hamilton cycle such that the expected length of the shortcut that skips each vertex v with probability $1 - p_v$ (independently of the others) is minimum. This problem appears in the following context. Drivers of delivery companies visit customers daily to deliver packages. For the company, the shorter the distance traversed, the better. For a driver, routes that change dramatically from one day to the other are inconvenient; it is better if one only has to shortcut a fixed route. The UPS problem, whose objective captures these two points of view, is at least as hard to approximate as the Metric TSP. Given that one of the vertices has probability one, we show that the performance ratio of a TSP tour for the UPS problem is $1/p_{\min}$, where $p_{\min} := \min_{v \in V} p_v$. We also show that this is tight. Consequently, Christofides' algorithm for the TSP has a performance ratio of $3/(2p_{\min})$ for the UPS problem and the approximation threshold for the UPS problem is at most $1/p_{\min}$ times the one for the TSP.

1 Introduction

1.1 Motivation

Package delivery companies, like the *United Parcel Service (UPS)*, have to deliver packages daily to several of their customers. The order of delivery is chosen so that to minimize the distance traversed by the drivers. Each delivery concerns only a subset of the customers. Therefore each delivery could be optimized individually. It is, however, easier for a driver

⋆ Research partially done while at Humboldt-Universität zu Berlin, supported in part by CAPES/DAAD Proc. 089/99, CNPq Proc. 301174/97-0, FAPESP Proc. 96/04505-2 and ProNEx 107/97 – MCT/FINEP (Brazil).

⋆⋆ Research supported in part by Deutsche Forschungsgemeinschaft, Pr296/6-1.

to shortcut a fixed route than to travel each time a completely different route. In this paper we study a variation of the *Traveling Salesman Problem (TSP)* which captures the issue described above.

The delivery company has the information on how often each customer receives a package. From this information one can estimate the probability that a customer receives a package per day. Roughly speaking, the here called *UPS Problem* consists of the following: find an ordering of all customers that minimizes the expected length of the route that starts and ends at a company location and visits in this ordering a randomly chosen (according to the customer's estimated probabilities) subset of the customers. The problem, as well as the described application, was proposed by [8]. Also according to [8], even the special case where customers are divided into two clusters—the customers who receive packages often and the ones who receive packages not so often—is of interest. The setup is conceivable in other delivery systems as well. We therefore expect the study of this problem to have several applications.

1.2 Notation and Problem Statement

Let $G = (V, E)$ be the complete graph on n vertices. A *path* is a sequence $\langle v_0, v_1, \ldots, v_k \rangle$ of distinct vertices of G. A *cycle* is a sequence $\langle v_0, v_1, \ldots, v_k \rangle$ of vertices of G, where $v_0, v_1, \ldots, v_{k-1}$ are distinct and $v_k = v_0$. For a path (cycle) $P = \langle v_0, v_1, \ldots, v_k \rangle$, we denote by $V(P)$ the set $\{v_0, v_1, \ldots, v_k\}$ and we say that P is a path (cycle) on $V(P)$. A *tour* (or a *Hamilton cycle*) is a cycle on V.

We denote the set $\{(v_0, v_1), (v_1, v_2), \ldots, (v_{k-1}, v_k)\}$ by $E(P)$. The *length of P with respect to a function $l : V^2 \to R^+$* is denoted by $l(P)$, and is given by

$$l(P) := \sum_{e \in E(P)} l(e).$$

The *Traveling Salesman Problem (TSP)* is the following: given a complete graph $G = (V, E)$ and a function $l : V^2 \to R^+$, find a tour of minimum length. We refer to such a tour as a *TSP tour*.

Unless specified otherwise, we consider in the following only functions $l : V^2 \to R^+$ that satisfy the *triangle inequality*: for any x, y, z in V, $l(x, y) \leq l(x, z) + l(z, y)$. Under this condition, TSP is called *Metric TSP*.

Note that l may be given partially. Then the length of an edge is considered to be the infimum of the lengths of all paths between its end vertices. This *closure* satisfies the triangle inequality.

Given a path P and a subset S of $V(P)$, the *shortcut of P induced by S*, denoted by $sc_S(P)$, is the path on S given by the subsequence of P containing exactly the vertices of S. Similarly we can define the *shortcut of a cycle C induced by a subset S of $V(C)$* and denote it by $sc_S(C)$.

Assume each vertex v in V has an associated probability p_v and let $\mathbf{p} := (p_v)_{v \in V}$. These probabilities induce a probability distribution on the subsets of V: for each $S \subseteq V$,

$$\Pr[S] := \prod_{v \in S} p_v \prod_{v \notin S} (1 - p_v).$$

Let $\mathbf{p}$ and $\mathbf{q}$ be two sets of vertex probabilities. We say that $\mathbf{p}$ dominates $\mathbf{q}$ if $p_v \geq q_v$ for all $v \in V$.

Given a cycle C, denote by $lsc_{\mathbf{p}}(C)$ the expected value of the length of the shortcut of C induced by a vertex subset which is randomly chosen according to $\mathbf{p}$:

$$lsc_{\mathbf{p}}(C) := \sum_{S \subseteq V} \Pr[S] \, l(sc_S(C)).$$

Now we are ready to state the UPS problem:

Definition 1. *Given the complete graph $G = (V, E)$, a function $l : V^2 \to R^+$ satisfying the triangle inequality, and probabilities $\mathbf{p} = (p_v)_{v \in V}$, the* UPS problem *asks for a tour C that minimizes $lsc_{\mathbf{p}}(C)$.*

The performance ratio of a tour C for the UPS problem is the ratio $lsc_{\mathbf{p}}(C)/\text{opt}$, where opt denotes the optimal value of the UPS problem, that is, $\text{opt} = lsc_{\mathbf{p}}(C^{\text{UPS}})$ for some optimal tour C^{UPS} of the UPS problem.

Throughout the paper, we consider the UPS problem under the additional assumption that $p_u = 1$ for at least one vertex u (representing, say, a UPS location).

1.3 Results

We start studying a restricted class of vertex probabilities for the UPS problem. Let G, l, and $\mathbf{p}^*$ be the input of the UPS problem, where we assume that, for some $0 \leq p \leq 1$, $p^*_v \in \{p, 1\}$ for all v. Our first result is a lower bound on the objective function for this particular case in terms of the TSP optimum.

Theorem 1. *Let C be a tour and C^{TSP} be a TSP tour. Then $lsc_{\mathbf{p}^*}(C) \geq p \cdot l(C^{TSP})$.*

The assumption that the vertex probabilities only attain the values p and 1 can be removed with the help of the following proposition.

Proposition 1. *Let* $\mathbf{p}$ *and* $\mathbf{q}$ *be two sets of vertex probabilities, where* $\mathbf{p}$ *dominates* $\mathbf{q}$*. Then* $lsc_{\mathbf{q}}(C) \leq lsc_{\mathbf{p}}(C)$ *for any tour* C*.*

The performance ratio of C^{TSP} as a solution for the general UPS problem follows from Theorem 1, using Proposition 1:

Corollary 1. *Let* C^{TSP} *be a TSP tour, let* C^{UPS} *be an optimal solution of the UPS problem, and let* $p_{\min} := \min_{v \in V} p_v$*. Denote by* opt *the optimal value* $lsc_{\mathbf{p}}(C^{UPS})$*. Then*

$$\frac{lsc_{\mathbf{p}}(C^{TSP})}{\mathrm{opt}} \leq \frac{1}{p_{\min}}.$$

Our second result is that this bound is tight.

Theorem 2. *For every* $\epsilon > 0$ *there is an instance* I_ϵ *to the UPS problem such that there are two TSP tours* C_1 *and* C_2 *with* $lsc_{\mathbf{p}}(C_2) \leq (p_{\min} + \epsilon) lsc_{\mathbf{p}}(C_1)$*.*

The tightness then follows from $\mathrm{opt} \leq lsc_{\mathbf{p}}(C_2)$.

When studying approximations for a computational problem, it is certainly necessary to explore the complexity theoretical limitations of that approach. We prove the following hardness of approximation result.

Theorem 3. *The approximation threshold of the UPS problem with any constantly bounded probability set is not less than the approximation threshold of the Metric TSP.*

The proof of Theorem 1 will be given in Section 2. In Section 3 we will sketch the proof of Proposition 1 and give more details on Corollary 1. Theorems 2 and 3 will be proved in Sections 4 and 5, respectively.

1.4 Conclusions and Open Problems

The UPS problem extends the TSP in that not only the length of the tour, but also the lengths of the subtours determine its objective value. As one might expect, the tradeoff depends on the vertex probabilities. We give matching upper and lower bounds on the rate in Theorems 1 and 2.

This result, being interesting in its own right, has several consequences for the approximation properties of the UPS problem. The currently best known approximation algorithm for the Metric TSP, Christofides' algorithm [3], has a performance ratio of $\frac{3}{2}$. As a consequence of Corollary 1, the same algorithm has a performance ratio of $\frac{3}{2p_{\min}}$ for the UPS problem.

Similarly, every other approximation algorithm for the Metric TSP can be applied to the UPS problem, while the performance ratio is multiplied by a factor of $\frac{1}{p_{\min}}$. Thus, the approximation threshold of the UPS problem is at most $\frac{\theta}{p_{\min}}$, where θ is the approximation threshold for the Metric TSP (for the definition of the approximation threshold and related notions see, e.g., Chapter 13 in [5]). This fact is complemented by Theorem 3, which states that it is at least θ.

One of the first questions that one might ask in this context concerns the influence of different probabilities. The factor of $\frac{1}{p_{\min}}$ might seem too pessimistic, if there were only few vertices with probability $p_{\min}$ and lots of vertices with much larger probabilities. However, the tight examples given in the proof of Theorem 2 can be modified so as to show that the bound given in Theorem 1 is very accurate.

The situation is less clear in the case of approximation algorithms. Here it is conceivable that an algorithm takes into account the distances given by l and combines them with the individual vertex probabilities in a clever way. The non-approximability result in Theorem 3 does not set any limit for that, however it is not less conceivable that the hardness result could be improved to show that the approximation threshold is actually $\frac{\theta}{p_{\min}}$.

The same consideration applies to the important special case of Euclidean instances. An instance of the TSP is called Euclidean if there is a point in the plane for every vertex such that the distance given by l is the Euclidean distance between the points. For this special case, there exist polynomial-time approximation schemes (PTAS) [2, 4] for TSP. (A PTAS consists of a polynomial-time algorithm for the problem with a performance ratio of at most $1+\epsilon$, for each $\epsilon > 0$.) Thus, for each $\epsilon > 0$, by Theorem 1 there exists a polynomial time algorithm for the UPS problem with Euclidean instances with a performance ratio of at most $\frac{1+\epsilon}{p_{\min}}$. On the other hand, it is conceivable both that there is a PTAS and that the approximation threshold is up to $\frac{1}{p_{\min}}$.

2 The UPS Problem with Probabilities 1 or p

The aim of this section is to prove Theorem 1. Recall that there is a $u \in V$ with $p^*_u = 1$ and that, for that theorem, the vertex probabilities are restricted to values of p and 1.

Let C^{TSP} be a TSP tour and C be a tour. To bound $lsc_{\mathbf{p}^*}(C)$ in terms of $l(C^{\mathrm{TSP}})$ we first need another formulation for the corresponding expectation. To this end we introduce some more notation.

Let U be the set of vertices of probability 1, $\bar{U} := V \setminus U$, and let t be the number of vertices in $\bar{U}$. If $t = 0$ then $U = V$ and $lsc_{\mathbf{p}^*}(C) = l(C)$. Since $l(C) \geq l(C^{\text{TSP}})$, the theorem clearly holds in this case. So we may assume $t \geq 1$.

For every $u, v \in \bar{U}$, let P_{uv} be the subsequence of C beginning at u and ending at v (circularly). Let P^1_{uv} be the shortcut of P_{uv} induced by $\{u, v\} \cup U$. Note that P^1_{vv} denotes a cycle—the shortcut of C induced by $\{v\} \cup U$.

Denote by $v_0, v_1, \ldots, v_{t-1}$ the vertices of $\bar{U}$ in the order given by C. For $i = 0, \ldots, t-1$, set $\mathcal{C}_i := \{P^1_{v_j v_{j+i+1}} : 0 \leq j < t\}$, where indices are taken modulo t, and $l(\mathcal{C}_i) := \sum_{P \in \mathcal{C}_i} l(P)$. Each $\mathcal{C}_i$ is a collection of paths (cycles if $i = t-1$) in G whose concatenation results in an Eulerian subgraph of G. Because $U \neq \emptyset$ and each vertex in $\bar{U}$ appears in some path (cycle if $i = t-1$) in $\mathcal{C}_i$, each of these Eulerian subgraphs is connected and spanning. Therefore, for each i,

$$l(\mathcal{C}_i) \geq l(C^{\text{TSP}}). \tag{1}$$

Using the notation above we can give the following characterization of $lsc_{\mathbf{p}^*}(C)$:

Lemma 1. $lsc_{\mathbf{p}^*}(C) = (1-p)^t l(sc_U(C)) + p(1-p)^{t-1} l(\mathcal{C}_{t-1}) + \sum_{i=0}^{t-2} p^2 (1-p)^i l(\mathcal{C}_i)$.

Proof. By definition, $lsc_{\mathbf{p}^*}(C) = \sum_{U \subseteq S \subseteq V} \Pr[S]\, l(sc_S(C))$. Note that the summands where $|S \setminus U| \leq 1$ contribute with $(1-p)^t l(sc_U(C)) + p(1-p)^{t-1} l(\mathcal{C}_{t-1})$ to $lsc_{\mathbf{p}^*}(C)$.

Let $\mathcal{S} := \{S : U \subseteq S \subseteq V, |S \setminus U| \geq 2\}$ be the collection of the other vertex subsets. For any $S \in \mathcal{S}$, let $E_S := E(sc_{S \setminus U}(C))$. Then, adding indices modulo t,

$$\begin{aligned}
\sum_{S \in \mathcal{S}} \Pr[S]\, l(sc_S(C)) &= \sum_{S \in \mathcal{S}} \Pr[S] \sum_{(v_i, v_j) \in E_S} l(P^1_{v_i v_j}) \\
&= \sum_{i=0}^{t-1} \sum_{j=0, j \neq i}^{t-1} l(P^1_{v_i v_j}) \sum_{\substack{S \in \mathcal{S} \\ (v_i, v_j) \in E_S}} \Pr[S] \\
&= \sum_{i=0}^{t-2} \sum_{j=0}^{t-1} l(P^1_{v_j v_{j+i+1}}) \Pr[(v_j, v_{j+i+1}) \in E_S] \\
&= \sum_{i=0}^{t-2} l(\mathcal{C}_i) p^2 (1-p)^i
\end{aligned}$$

and the lemma holds. $\square$

Putting (1) and Lemma 1 together, we have that

$$\begin{aligned} lsc_{\mathbf{p}^*}(C) &\geq p(1-p)^{t-1} l(C^{\text{TSP}}) + \sum_{i=0}^{t-2} p^2(1-p)^i l(C^{\text{TSP}}) \\ &= l(C^{\text{TSP}}) \left(p(1-p)^{t-1} + \sum_{i=0}^{t-2} p^2(1-p)^i \right). \end{aligned}$$

Straightforward calculation shows that the right hand side is equal to $l(C^{\text{TSP}})p$, concluding the proof of Theorem 1. $\square$

3 Arbitrary Probabilities

Our result on the UPS problem with arbitrary probabilities, i.e. Corollary 1, is a consequence of Proposition 1. The proof of Proposition 1 is based on the FKG-Inequality [1, p. 75] and we only sketch it here.

Let $\mathbf{p}$ and $\mathbf{q}$ be two sets of vertex probabilities and assume that $\mathbf{p}$ dominates $\mathbf{q}$. Let C be any tour. For $S \subseteq V$, let $f(S) := l(sc_S(C))$ and $g(S) := \prod_{v \in S}(p_v/q_v) \cdot \prod_{v \notin S}(1-p_v)/(1-q_v)$. Observe that $S \mapsto \Pr_{\mathbf{q}}[S]$ is log-supermodular, that f is increasing because of the triangle inequality, and that g is increasing because $\mathbf{p}$ dominates $\mathbf{q}$. Thus, the requirements of the FKG inequality are met and we have

$$\sum_{S \subseteq V} f(S) \Pr_{\mathbf{q}}[S] \cdot \sum_{S \subseteq V} g(S) \Pr_{\mathbf{q}}[S] \leq \sum_{S \subseteq V} f(S) g(S) \Pr_{\mathbf{q}}[S] \cdot \sum_{S \subseteq V} \Pr_{\mathbf{q}}[S].$$

Since $\Pr_{\mathbf{p}}[S] = \Pr_{\mathbf{q}}[S] \cdot g(S)$, this implies Proposition 1. $\square$

Corollary 1 follows by sandwiching $\mathbf{p}$ between two appropriate sets of vertex probabilities. More precisely, let

$$p_v^* := \begin{cases} p_{\min} & \text{if } p_v < 1, \\ 1 & \text{otherwise.} \end{cases}$$

Then $\mathbf{p}^*$ is dominated by $\mathbf{p}$, which in turn is dominated by the all-ones probability set. The corollary follows from

$$\text{opt} = lsc_{\mathbf{p}}(C^{\text{UPS}}) \geq lsc_{\mathbf{p}^*}(C^{\text{UPS}}) \geq p_{\min}\, l(C^{\text{TSP}}) \geq p_{\min}\, lsc_{\mathbf{p}}(C^{\text{TSP}}),$$

where the first and the last inequality are implied by Proposition 1 and the second inequality by Theorem 1, applied to $C = C^{\text{UPS}}$.

4 Tight Examples

Assume that $p < 1$ and let $\epsilon > 0$. In this section we give the construction of an instance I_ϵ of the UPS problem with probabilities p and 1. It has two TSP tours C_1 and C_2 and (3) states that their UPS values differ by a factor of at least $\frac{1}{p+\epsilon}$. This implies Theorem 2.

We assume w.l.o.g. that $\epsilon < 1 - p$. Let k be a positive integer, large enough so that $k + \log_{1-p}(8k^2) \geq \log_{1-p} \epsilon$. Let n be a prime such that $2k^2 < n < 4k^2$. Then

$$2n(1-p)^k \leq 8k^2(1-p)^k \leq \epsilon. \tag{2}$$

Let $V := \{0, \ldots, n-1\}$ and let $H := (V, E)$, where

$$E := \{ij : j - i \pmod{n} \leq k\}.$$

That is, $H = C_n^k$ is a cycle on n vertices plus all chords of length at most k. Let $l(e) = 1$ for all $e \in E$. Then two TSP tours for H and l are (indices are taken modulo n, as usual)

$$C_1 := \langle 0, k, \ldots, ik, \ldots, nk \rangle \text{ and}$$
$$C_2 := \langle 0, 1, \ldots, i, \ldots, n \rangle.$$

Note that C_1 is a tour because of the primality of n.

Let $p_0 := 1$ and $p_i := p$ for $i \geq 1$. Then I_ϵ consists of (the closure of) H, l, and $\mathbf{p}$. In the rest of this section we shall prove that

$$lsc_{\mathbf{p}}(C_2) \leq (p + \epsilon) lsc_{\mathbf{p}}(C_1). \tag{3}$$

Let S be a randomly chosen subset of V. Call S *dense for* C_1 if S intersects any set of k consecutive vertices of C_1. The probability of that event is at most $n(1-p)^k$, where $n(1-p)^k \leq \epsilon/2$ by (2). The event that S is *dense for* C_2 is defined analogously, and its probability is the same.

Assume that S is dense for C_1 and let ik and $jk > ik$ be two subsequent vertices of $sc_S(C_1)$. Then $j - i \leq k$ because S is dense for C_1. But then $jk - ik \leq k^2 < n/2$, and by the choice of H the distance between ik and jk is $j - i$. Assume that $sc_S(C_1) = \langle i_1 k, \ldots, i_{|S|} k \rangle$ and let $i_0 := i_{|S|}$. Then

$$n = \sum_{j=1}^{|S|} (i_j - i_{j-1} \pmod{n}) = \sum_{j=1}^{|S|} l(i_j k, i_{j-1} k) = l(sc_S(C_1)),$$

and therefore

$$lsc_{\mathbf{p}}(C_1) \geq n\Pr\left[S \text{ is dense for } C_1\right] \geq (1-\epsilon/2)n. \tag{4}$$

If S is dense for C_2, then there is a chord between any two subsequent vertices of $sc_S(C_2)$ and thus $l(sc_S(C_2)) = |S|$. This implies that

$$lsc_{\mathbf{p}}(C_2) \leq n\Pr\left[S \text{ is not dense for } C_2\right] + \sum_{S \subseteq V} |S|\Pr\left[S\right] \leq (p+\epsilon/2)n. \tag{5}$$

As a consequence of (4) and (5),

$$\frac{lsc_{\mathbf{p}}(C_2)}{lsc_{\mathbf{p}}(C_1)} \leq \frac{p+\epsilon/2}{1-\epsilon/2},$$

which implies (3), using $\epsilon < 1-p$. □

5 Hardness of Approximation

The Metric TSP is APX-complete [7] and the currently best lower bound on the approximation threshold is $\frac{41}{40}$ in the asymmetric case and $\frac{129}{128}$ in the symmetric case [6]. It is trivial that the UPS problem has the same lower bounds, because the objective functions coincide for the all-ones probability set $p_v \equiv 1$. It might, however, be interesting to verify that the same holds for the probability set $p_v \equiv p$, where $0 < p < 1$. This, together with Proposition 1, proves Theorem 3.

Next we present a reduction from Metric TSP to the UPS problem in instances with probability set $p_v \equiv p$ for each $\epsilon > 0$. The reduction preserves the approximation ratio up to a factor of $1-\epsilon$.

Let V be a vertex set and let $l : V^2 \to R^+$ be an instance of the Metric TSP on V. The corresponding instance of the UPS problem consists of the following. Add $c_{p,\epsilon}(n)$ copies of every vertex to get V' and let l' be the extension of l to V' such that all copies of the same vertex have distance 0 to each other and copies of different vertices have the same distance as the original vertices. Here $c_{p,\epsilon}(n) = O(\log_{\frac{1}{1-p}} n)$ is chosen large enough that $(1-(1-p)^{c_{p,\epsilon}(n)})^n \geq 1-\epsilon$. Note that this can be done in polynomial time.

Any tour on V' whose performance ratio (for the UPS instance) is at most η can be converted into a tour on V whose performance ratio (for the original TSP instance) is at most $\eta/(1-\epsilon)$. Indeed, let C' be a tour on V' whose performance ratio is at most η for the UPS problem. We may assume w.l.o.g. that all copies of each original vertex occur subsequently in C'. Then $l(sc_S(C')) = l(C')$ as long as S contains at least one copy of each original vertex. Since the probability for that event is at least $1-\epsilon$, we have that $lsc_{\mathbf{p}}(C') \geq (1-\epsilon)l'(C') = (1-\epsilon)l(C)$, where C is the ordering of V induced by C'. Now let C^{TSP} be a TSP tour on V. If we extend it to a tour on V' by visiting all copies of every vertex subsequently, we know that $lsc_{\mathbf{p}}(C^{\mathrm{TSP}}) \leq l'(C^{\mathrm{TSP}}) = l(C^{\mathrm{TSP}})$. Therefore the optimal UPS solution has length at most $l(C^{\mathrm{TSP}})$. Thus $(1-\epsilon)l(C) \leq lsc_{\mathbf{p}}(C) \leq \eta lsc_{\mathbf{p}}(C^{\mathrm{UPS}}) \leq \eta l(C^{\mathrm{TSP}})$. This completes the analysis of the reduction, and, together with Proposition 1, the proof of Theorem 3. $\square$

6 Acknowledgment

We thank Martin Savelsbergh for suggesting the problem and Deryk Osthus and Hanno Lefmann for fruitful discussions. Carlos Eduardo Rodrigues Alves gave a simple proof of a weaker version of Theorem 1 (with p^2 instead of p), which inspired our proof. An anonymous referee gave several useful hints on the presentation of the paper.

References

1. N. Alon, J. Spencer, P. Erdős, *The Probabilistic Method*, Wiley, New York, 1992.
2. S. Arora, *Polynomial Time Approximation Schemes for Euclidean TSP and other Geometric Problems*, Proceedings of the 37th Symposium on the Foundations of Computer Science, 2–11 (1996).
3. N. Christofides, *Worst-case Analysis of a new Heuristic for the Travelling Salesman Problem*, Technical Report (Graduate School of Industrial Administration, Carnegie-Mellon University, Pittsburgh, PA), 1976.
4. J. Mitchell, *Guillotine Subdivisions Approximate Polygonal Subdivisions: A Simple Polynomial-Time Approximation Scheme for Geometric TSP, k-MST, and Related Problems*, SIAM Journal on Computing 28, 1298–1309 (1999).
5. C.H. Papadimitriou, *Computational Complexity*, Addison-Wesley, 1994.
6. C.H. Papadimitriou and S. Vempala, *On the Approximability of the Traveling Salesman Problem*, Proceedings of the 32nd Annual ACM Symposium on Theory of Computing, 126–133 (2000).
7. C.H. Papadimitriou and M. Yannakakis, *The Traveling Salesman Problem with Distances One and Two*, Mathematics of Operations Research 18, 1–11 (1993).
8. M. Savelsbergh, personal communication (1999).