European Journal of Operational Research 65 (1993) 68-95
North-Holland

Theory and Methodology

Further results on the probabilistic traveling salesman problem

Dimitris Bertsimas

Sloan School of Management, Massachusetts Institute of Technology, Cambridge, USA

Louis H. Howell

Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

Received June 1990; revised March 1991

Abstract: In 1985, Jaillet introduced the probabilistic traveling salesman problem (PTSP), a variant of the classical TSP in which only a subset of the nodes may be present in any given instance of the problem. The goal is to find an a priori tour of minimal expected length, with the strategy of visiting the present nodes in a particular instance in the same order as they appear in the a priori tour. In this paper we reexamine the PTSP using a variety of theoretical and computational approaches. We sharpen the best known bounds for the PTSP, derive several asymptotic relations, and compare from various viewpoints the PTSP with the re-optimization strategy, i.e., finding an optimal tour in every problem instance. When a Euclidean metric is used and the nodes are uniformly distributed in the unit square, a heuristic for the PTSP is shown to be very close to the re-optimization strategy. We examine some PTSP heuristics with provable worst-case performance, and address the question of finding constant-guarantee heuristics. Implementations of various heuristics, some based on sorting and some on local optimality, permit us to discuss the qualitative and quantitative properties of computational problems with up to 5000 nodes.

Keywords: Probabilistic traveling salesman problem; Probabilistic analysis; Combinatorial optimization; Heuristics

1. Introduction

The traveling salesman problem (TSP) is probably the central problem in combinatorial optimization. Many new ideas in combinatorial optimization have been tested on the TSP, including asymptotic analysis [3], dynamic programming [13], cutting planes [11], Lagrangian relaxation [14,15], and recently polyhedral combinatorics [12]. A very natural and fundamentally different variation of the TSP was introduced in Jaillet's Ph.D. thesis [16], called the probabilistic traveling salesman problem (PTSP).

Correspondence to: D. Bertsimas, Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

0377-2217/93/$06.00 © 1993 – Elsevier Science Publishers B.V. All rights reserved

Consider a problem of routing through a set of n known points. On any given instance of the problem only a subset consisting of k out of n points ($0 \leqslant k \leqslant n$) must be visited, the active subset being determined by a known probability distribution. Ideally we would like to re-optimize, i.e., find an optimal TSP tour for every instance. In some cases, however, we may not have the resources to do this, or even if we have them it may be very time consuming or otherwise inefficient to do so. Instead, we wish to find a priori a tour through all n points. On any given instance of the problem, the k points present will then be visited in the same order as they appear in the a priori tour. The problem of finding such an a priori tour which is of minimum length in the expected value sense is defined as the PTSP.

The expectation is computed over all possible instances of the problem, i.e., over all subsets of the vertex set $V = \{1, 2, \ldots, n\}$. That is, given an a priori tour τ, if problem instance $S(\subseteq V)$ occurs with probability $p(S)$ and requires covering a total distance $L_\tau(S)$ to visit the subset S of customers, that problem instance will receive a weight of $p(S)L_\tau(S)$ in the computation of the expected length. If we denote the length of the tour τ by L_τ (a random variable), then our problem is to find an a priori tour τ_p through all n potential customers, which minimizes the quantity

$$E[L_\tau] = \sum_{S \subseteq V} p(S) L_\tau(S), \tag{1}$$

with the summation being over all subsets of V.

In general, PTSP's arise in practice whenever a company, on any given day, is faced with the problem of collections (deliveries) from (to) a random subset of its (known) global set of customers in an area. The reason for not re-optimizing the tour on every problem instance could be that the system's operator may have other priorities that could best be attained by following a vehicle routing-like strategy, such as regularity and personalization of service by having the same vehicle and driver visit a particular customer every day. Examples in this category that have been described in the literature include a 'hot meals' delivery system [2], routing of forklifts in a cargo terminal or in a warehouse and, interestingly, the daily delivery of mail to homes and business by Post Office mail carriers everywhere. It is also interesting that this generic problem was first formulated by Odoni [25] in an attempt to deal with this particular application.

In a routing context other areas of application include strategic planning and transportation models. In a non-routing context, PTSP models can also be of interest in many situations in which a sequence of entities has to be found and that sequence has to be preserved even when some of the entities may be absent, for example in the area of job-shop scheduling. For a further discussion of areas of application of PTSP models see [17].

In the following section we review briefly the results of Jaillet's thesis. In Section 3 we examine some combinatorial and functional properties of the problem by sharpening the best bounds for the PTSP, examining some solvable special cases, and exploring the relation between the PTSP and the the probabilistic minimum spanning tree (PMST) problem defined in [5]. In Section 4 we compare the PTSP with the re-optimization strategy, i.e., finding an optimal tour in every problem instance. In Section 5 we examine some heuristics with provable worst-case performance and address the question of finding constant-guarantee heuristics for the PTSP.

In Section 6 we prove that the tour produced by the space-filling curve heuristic of Bartholdi and Platzman [1] is approximately within 25% of the re-optimization strategy if the points are randomly distributed in the unit square. Finally, we perform a probabilistic analysis of the nearest neighbor heuristic and show that on average it behaves very poorly in the PTSP case. In Section 7 we discuss and explain the implementation of various heuristics for the PTSP, which are of two different types: heuristics which are based on ideas of local optimality (2-popt, 2-opt, 3-opt, 1-shift), and heuristics which are based on sorting (angular sorting, space-filling curve). We report computational results for problems of up to 5000 nodes and we discuss some qualitative properties of the PTSP which are apparent in the computed solutions. The final section contains some concluding remarks.

2. Review of previous results

We briefly review here the main results of Jaillet's thesis in an attempt to make the present paper self-contained.

Computation of the expected length of a given PTSP tour

Suppose we are given an a priori PTSP tour τ through n points. Point i has probability p_i $(1 \leqslant i \leqslant n)$ of requiring a visit on any given instance of the problem independently of all other points. Is there an efficient way of computing $E[L_\tau]$, the expected length of the tour τ?

In his dissertation Jaillet [16] has derived the following efficient closed-form expressions for computing $E[L_\tau]$.

Result 1. If the distance between points i, j is $d(i, j)$ and we assume, without loss of generality, that the a priori tour is $\tau = (1, 2, \dots, n, 1)$, then

$$E[L_\tau] = \sum_{i=1}^{n} \sum_{j=i+1}^{n} d(i, j) p_i p_j \prod_{k=i+1}^{j-1} (1 - p_k)$$
$$+ \sum_{j=1}^{n} \sum_{i=1}^{j-1} d(j, i) p_i p_j \prod_{k=j+1}^{n} (1 - p_k) \prod_{k=1}^{i-1} (1 - p_k). \tag{2}$$

Expression (2) is derived by looking at the probability of every link being present. This expression takes, in general, $O(n^2)$ time to compute. In special cases (e.g., when all points i have $p_i = p$, a constant for all i) further simplifications, which lead to combinatorial interpretations of the expressions, can be achieved. Specifically, if the a priori tour is $\tau = (1, 2, \dots, n, 1)$, then

$$E[L_\tau] = p^2 \sum_{r=0}^{n-2} (1 - p)^r L_\tau^{(r)} \tag{3}$$

where $L_\tau^{(r)} \triangleq \sum_{j=1}^{n} d(j, (j + 1 + r) \bmod n)$. The $L_\tau^{(r)}$'s have the combinatorial interpretation of being the lengths of $\gcd(n, r + 1)$ subtours.

Similar results are derived for the case where $m < n$ of the points are deterministically present and $n - m$ are present with probability p. Thus, given a PTSP tour τ, we can now efficiently compute its expected length $E[L_\tau]$ over all instances of the problem.

Relationship between PTSP and TSP

It is natural to inquire about the links between solutions to the PTSP and to the TSP. In other words, how well would a TSP tour through n probabilistic points do as a solution to the true problem, i.e., the PTSP? The answer is: potentially very poorly. Under some very special conditions the solutions to the PTSP and TSP are identical (for example, when the n points lie at the corner points of a convex n-gon). In general, however, no assurance can be offered that this will be the case, except for trivially small problems. For example, when the n points are present independently of each other, $p_i = p$ for all $i = 1, 2, \dots, n$, and distances between points satisfy the triangle inequality, the following is true:

Result 2. The optimum TSP tour is guaranteed to solve the PTSP optimally for problems with only 5 or fewer points (and with only 3 or fewer points when the matrix of distances between points is not symmetric).

The errors that can result from using the optimal TSP tour as a solution to the corresponding PTSP can be very large as suggested by the following: for the n-point problem let us denote by τ_p the optimal PTSP tour and by L_{TSP} the length of the optimal traveling salesman tour (TST). Then we get

Result 3. Under the triangle inequality, if the number of deterministically present points is m, then for $m \geqslant 1$,

$$E\left[L_{\tau_p}\right] \geqslant pL_{\text{TSP}}, \qquad \frac{E[L_{\text{TSP}}] - E\left[L_{\tau_p}\right]}{E\left[L_{\tau_p}\right]} \leqslant \frac{1-p}{p}, \tag{4}$$

and for $m = 0$ and n prime,

$$E\left[L_{\tau_p}\right] \geqslant pL_{\text{TSP}}\left(1 - (1-p)^{n-1}\right). \tag{5}$$

Jaillet conjectures that (4) and (5) hold even if $m = 0$ and n not necessarily prime. The left-hand side of (4) represents the error (as a fraction of $E[L_{\tau_p}]$) that would result if the optimum TSP tour were used to solve the PTSP instead of the optimal solution τ_p. As p gets smaller, Result 3 suggests that the optimum TSP tour can be an arbitrarily poor substitute for τ_p, and indeed Jaillet [16] has shown actual examples where this is true.

Two additional (and counter-intuitive) results

The following two results underscore the point that the probabilistic aspects of the PTSP induce some characteristics which are distinctly different from those of the TSP.

Result 4. For a PTSP in the Euclidean plane, the optimal PTSP tour may intersect itself.

Result 4, of course, is in sharp contrast with one of the first known properties of optimal TSP tours, namely that in a Euclidean metric the optimal tour does not interset itself. As this property is often exploited by heuristic algorithms for solving the TSP in the plane (e.g., [21]), one must be careful when using modified versions of such algorithms with PTSP's.

Result 5. The extension to the PTSP of the well-known dynamic programming solution approach to the TSP [13] fails to solve the problem. Both Results 4 and 5 were also noted in Dror et al. [9].

The dynamic programming approach to the TSP is probably the most widely-known one and seemingly so fundamentally simple that one might reasonably expect it to solve in principle PTSP's, as well. It turns out that this is not the case, due to the fact that PTSP's cannot be decomposed in a straightforward way into stages.

Bounds and asymptotic results

The derivation of upper and lower bounds for the optimal solutions of combinatorial problems and of asymptotic results (that apply in the limit as problem size increases) has become in recent years a standard and rewarding area of research in combinatorial optimization. Such results are very useful in a number of situations as, for instance: obtaining approximations for large-scale instances of routing problems; analyzing the properties of some heuristic algorithms; or exploring the boundaries between good and poor heuristics in the probabilistic (as opposed the worst-case) sense. The fundamental concepts, in this respect, have been pioneered in the seminal papers of Beardwood et al. [3] and of Karp [19] – see also the recent paper by Karp and Steele [20].

The simplest asymptotic result, in the case where n points are uniformly and independently distributed within the unit square and each point has probability p of being present, independently of the others, states that as $n \to \infty$, with probability 1,

$$\lim_{n \to \infty} \frac{E\left[L_{\tau_p}\right]}{\sqrt{n}} = c(p), \tag{6}$$

where $c(p)$ is independent of n.

Moreover, it can be shown that, for the Euclidean metric,

$$\beta_{\text{TSP}}\sqrt{p} \leqslant c(p) \leqslant \min\left(\beta_{\text{TSP}}, 0.9204\sqrt{p}\right), \quad 0 \leqslant p \leqslant 1$$

where β_{TSP} is the TSP constant in the theorem of Beardwood et al. [3], who showed that $0.625 \leqslant \beta_{TSP} \leqslant 0.9204$.

More general limit theorems can be derived, including cases simultaneously involving arbitrary numbers of probabilistic and deterministic (always present) points and extensions to any bounded Lebesque measurable set of d-dimensional Euclidean space as well as other metrics.

3. Combinatorial and functional properties of the PTSP

In this section we examine some further properties of the PTSP. Specifically, we first improve the best bounds known for the problem under various probabilistic assumptions and examine some solvable special cases. We further examine some functional properties of the problem and explore the relation between the PTSP and the PMST problem.

3.1. Improved bounds for the PTSP

In this subsection we improve the best upper and lower bounds for the PTSP in three cases of probabilistic assumptions: 1) $p_i = p$, 2) $p_i \neq p_j$ and 3) $p_1 = 1$, $p_i \neq p_j$.

1) The case $p_i = p$. In the bounds proved by Jaillet there is a strange duality in the simplest possible case. If $p_i = p$ for all nodes i, then he proves that (5) holds, if the number of nodes n is a prime number, while if the number of nodes is not prime his proof technique does not work. Applying the closed-form expression (3) for the optimal PTSP tour τ_p we obtain:

$$E\left[L_{\tau_p}\right] = p^2 \sum_{r=0}^{n-2} (1-p)^r L_{\tau_p}^{(r)} \tag{7}$$

where $L_{\tau_p}^{(r)}$ is the length of a collection of gcd(n, $r+1$) subtours $\tau_p^{(r)}$. For n prime, each one of the $n-1$ $\tau_p^{(r)}$'s is a tour and thus $L_{\tau_p}^{(r)} \geqslant L_{TSP}$, since L_{TSP} is the length of the minimum tour. Using this bound, (5) follows. The above analysis is independent of the triangle inequality, but clearly does not hold for n not prime, because then there exists an r such that $\tau_p^{(r)}$ has gcd(n, $r+1) \neq 1$ subtours. Faced with this unnatural distinction Jaillet conjectures that (5) holds even if n is not prime.

In proving a strengthened version of his conjecture for $n = 2k+1$, we give the following definition: Given a tour τ, we say that it has cycle structure $C_\tau \triangleq (\tau^{(0)}, \tau^{(1)}, \ldots, \tau^{(n-2)})$, where $\tau^{(r)}$ is the set of gcd(n, $r+1$) subtours which is formed by skipping r points. For an example see Figure 1.

Then we can prove the following strengthened version of Jaillet's conjecture:

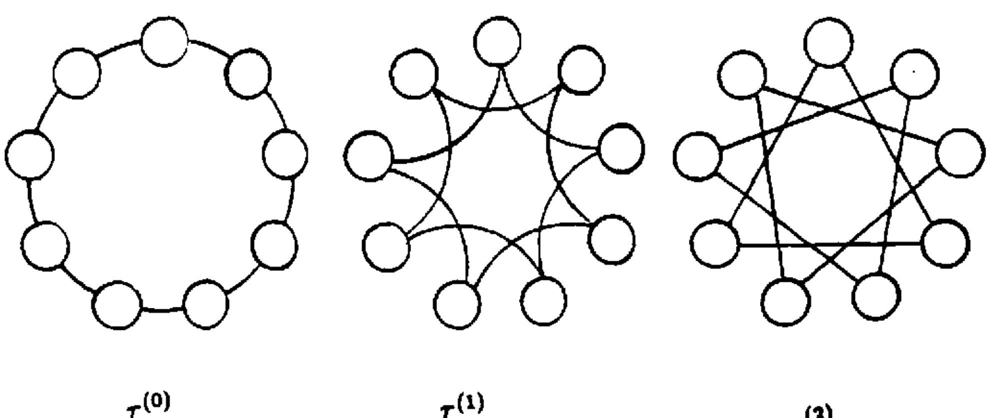

Figure 1. The cycle structure C_τ of a tour τ

Proposition 1. *If τ_p is the optimal probabilistic traveling salesman tour* (PTST), *then for* $n = 2k + 1$,

$$E\left[L_{\tau_p}\right] \geqslant pL_{\tau_p}\left\{1 + (1-p)^{2k-1} - (1-p)^k(2-p)\right\}. \tag{8}$$

Proof. Let the number of points be an odd number $n = 2k + 1$. Assuming that the cost matrix is symmetric we obtain that for every tour τ, $L_\tau^{(r)} = L_\tau^{(n-2-r)}$, since

$$\begin{aligned} L_\tau^{(n-2-r)} &= \sum_{j=1}^{n} d(j, (j+n-r-1) \bmod n) \\ &= \sum_{j=1}^{n} d((j+1+r) \bmod n, [(j+1+r) \bmod n + n - r - 1] \bmod n) \\ &= \sum_{j=1}^{n} d((j+1+r) \bmod n, j) = L_\tau^{(r)} \end{aligned}$$

where we used the symmetry in the last equality.

Suppose then that the optimal tour has cycle structure

$$\left(\tau_p = \tau_p^{(0)}, \tau_p^{(1)}, \ldots, \tau_p^{(k-1)}, \tau_p^{(k-1)}, \ldots, \tau_p^{(1)}, \tau_p^{(0)}\right).$$

Since $\gcd(n, 1+1) = 1$, $\tau_1 \triangleq \tau_p^{(1)}$ is a tour. Consider now the cycle structure of the tour τ_1. The key observation is that the cycle structure of τ_1 is:

$$\left(\tau_1 = \tau_p^{(1)}, \tau_p^{(2)}, \ldots, \tau_p^{(k-1)}, \tau_p^{(0)}, \tau_p^{(0)}, \tau_p^{(k-1)}, \ldots, \tau_p^{(2)}, \tau_p^{(1)}\right).$$

As a result, the expectations of the tours τ_p, τ_1 are:

$$\begin{aligned} E\left[L_{\tau_p}\right] &= p^2 L_{\tau_p}\left[1 + (1-p)^{n-2}\right] + p^2(1-p)\sum_{r=0}^{k-1}(1-p)^{r-1}L_{\tau_p}^{(r)} + p^2\sum_{r=1}^{k-1}(1-p)^{n-r-2}L_{\tau_p}^{(r)}, \\ E\left[L_{\tau_1}\right] &= p^2 L_{\tau_p}\left[(1-p)^{k-1} + (1-p)^k\right] + p^2\sum_{r=0}^{k-1}(1-p)^{r-1}L_{\tau_p}^{(r)} \\ &\quad + p^2(1-p)\sum_{r=1}^{k-1}(1-p)^{n-r-2}L_{\tau_p}^{(r)}. \end{aligned}$$

With the definitions

$$a \triangleq p^2\sum_{r=0}^{k-1}(1-p)^{r-1}L_{\tau_p}^{(r)} \text{ and } b \triangleq p^2\sum_{r=1}^{k-1}(1-p)^{n-r-2}L_{\tau_p}^{(r)},$$

these expectations become

$$\begin{aligned} E\left[L_{\tau_p}\right] &= p^2 L_{\tau_p}\left[1 + (1-p)^{n-2}\right] + (1-p)a + b, \\ E\left[L_{\tau_1}\right] &= p^2 L_{\tau_p}(1-p)^{k-2}(2-p) + a + (1-p)b. \end{aligned}$$

As a result,

$$(1-p)a + b = E\left[L_{\tau_p}\right] - p^2 L_{\tau_p}\left[1 + (1-p)^{n-2}\right], \tag{9}$$

$$a + (1-p)b = E\left[L_{\tau_1}\right] - p^2 L_{\tau_p}(1-p)^{k-2}(2-p). \tag{10}$$

But, since

$$(1-p)a+b=(1-p)\left(a+\frac{b}{1-p}\right)\geqslant(1-p)(a+(1-p)b), \tag{11}$$

if we substitute (9), (10) to (11), we obtain

$$E\left[L_{\tau_p}\right]-p^2L_{\tau_p}\left[1+(1-p)^{n-2}\right]\geqslant(1-p)\Big(E\left[L_{\tau_1}\right]-p^2L_{\tau_p}(1-p)^{k-1}(2-p)\Big).$$

Since τ_1 is a tour, $E[L_{\tau_1}]\geqslant E[L_{\tau_p}]$, since $E[L_{\tau_p}]$ is the PTST. As a result, we find that

$$E\left[L_{\tau_p}\right]-p^2L_{\tau_p}\left[1+(1-p)^{n-2}\right]\geqslant(1-p)\Big(E\left[L_{\tau_p}\right]-p^2L_{\tau_p}(1-p)^{k-1}(2-p)\Big),$$

from which (8) follows. □

A proof along the same lines can be given in the case of an asymmetric cost matrix. Note that the result does not assume the triangle inequality and is stronger than Jaillet's conjecture in the sense that it involves L_{τ_p} rather than L_{TSP}. In addition, the fact that n was odd was used extensively, since τ_1 is a tour. For n even the above proof does not work.

2) The case $p_i\neq p_j$. In the case of equal probabilities the combinatorial interpretation of expression (3) for the expected length enables us to prove some lower bounds. In the case of unequal probabilities, we obtain lower bounds for the expected length of the PTST by using a mathematical programming rather than a combinatorial approach. In this case, we assume that node i is present with probability p_i. We then use an idea suggested in [4].

Proposition 2. *If τ_p is the optimal* PTST, *then*

$$E\left[L_{\tau_p}\right]\geqslant z^* \tag{12}$$

where z^* is the optimal solution to the transportation problem

$$\begin{aligned} z^*=\min\ &\sum_{i,j}x_{i,j}d(i,j),\\ \text{s.t.}\quad &\sum_i x_{i,j}=p_j\Big(1-\prod_{k\neq i}(1-p_k)\Big),\\ &\sum_j x_{i,j}=p_i\Big(1-\prod_{k\neq i}(1-p_k)\Big),\\ &x_{i,j}\geqslant 0. \end{aligned} \tag{13}$$

Proof. Let $d_1(i,j)\triangleq d(i,j)-u_i-v_j$ and we will choose u_i, v_i to obtain the best possible bounds. By appropriately renumbering the nodes, let $\tau_p=(1,2,\ldots,n,1)$ be the optimum PTST with respect to the original distances $d(i,j)$. Let $E[L_{\tau_p}]$, $E[L^1_{\tau_p}]$ be the expected lengths of the tour τ_p with respect to distances $d(i,j)$, $d_1(i,j)$, respectively. Then

$$\begin{aligned} E\left[L^1_{\tau_p}\right]=E\left[L_{\tau_p}\right]&-\sum_{i=1}^{n}\sum_{j=i+1}^{n}(u_i+v_j)p_ip_j\prod_{k=i+1}^{j-1}(1-p_k)\\ &-\sum_{j=1}^{n}\sum_{i-1}^{j-1}(u_j+v_i)p_ip_j\prod_{k=j+1}^{n}(1-p_k)\prod_{k=1}^{i-1}(1-p_k). \end{aligned}$$

After some algebraic manipulations we find that

$$E\left[L^1_{\tau_p}\right]=E\left[L_{\tau_p}\right]-\sum_{i=1}^{n}(u_i+v_i)p_i\left(1-\prod_{k\neq i}(1-p_k)\right).$$

If we demand that $d_1(i, j)=d(i, j)-u_i-v_j\geqslant 0$, we have the trivial bound $E[L^1_{\tau_p}]\geqslant 0$. As a result, we obtain

$$E\left[L_{\tau_p}\right]\geqslant\sum_{i=1}^{n}(u_i+v_i)p_i\left(1-\prod_{k\neq i}(1-p_k)\right).$$

Since we want the best possible bound, we wish to choose u_i, v_i satisfying:

$$\begin{aligned}\max\quad&\sum_{i=1}^{n}(u_i+v_i)p_i\left(1-\prod_{k\neq i}(1-p_k)\right),\\ \text{s.t.}\quad&d(i, j)-u_i-v_j\geqslant 0.\end{aligned}\tag{14}$$

The dual of the linear program (14) is (13) and hence (12) follows from strong duality. □

Corollary. *If we let $p_i=p$ in* (13), *we obtain that*

$$E\left[L_{\tau_p}\right]\geqslant p\left(1-(1-p)^{n-1}\right)L_a\tag{15}$$

where L_a is the length of the minimum assignment problem.

Note that this bound is weaker than the bound (8). A combinatorial proof of (15) can be easily obtained by noticing that from (3) each of the terms $L^{(r)}_{\tau_p}\geqslant L_a$ since $\tau^{(r)}_p$ is a solution to the assignment problem.

For the lower bound we have not assumed that the triangle inequality holds, but in order to find a useful upper bound we will need the triangle inequality, since otherwise $E[L_{\tau_p}]$ can be arbitrarily large while the optimum *TST* is small.

Proposition 3. *Under the triangle inequality,*

$$E\left[L_{\tau_p}\right]\leqslant L_{\mathrm{TSP}}.\tag{16}$$

Proof. For any tour τ, $E[L_\tau]=\Sigma_{S\subseteq V}p(S)L_\tau(S)$. From the triangle inequality, clearly $L_\tau(S)\leqslant L_\tau$. Then $E[L_{\tau_p}]\leqslant E[L_{TSP}]\leqslant L_{\mathrm{TSP}}$. □

3) The case $p_1=1$, $p_i\neq p_j$. In this case we are trying to improve the bound (16) if the triangle inequality holds. We prove the following:

Proposition 4. *Under the triangle inequality, any tour τ satisfies*

$$E[L_\tau]\leqslant\sum_{i=2}^{n}p_i[d(i, 1)+d(1, i)].\tag{17}$$

Proof. Consider a tour $\tau=(1, 2,\dots, n, 1)$. In the case with $p_1=1$ expression (2) becomes:

$$\begin{aligned}E[L_\tau]=&\sum_{i=2}^{n}d(1, i)p_i\prod_{k=2}^{i-1}(1-p_k)+\sum_{i=2}^{n}d(i, 1)p_i\prod_{k=i+1}^{n}(1-p_k)\\&+\sum_{i=2}^{n}\sum_{j=i+1}^{n}d(i, j)p_ip_j\prod_{k=i+1}^{j-1}(1-p_k).\end{aligned}$$

From the triangle inequality, $d(i, j) \leqslant d(i, 1) + d(1, j)$. As a result,

$$\sum_{i=2}^{n} \sum_{j=i+1}^{n} d(i, j) p_i p_j \prod_{k=i+1}^{j-1} (1-p_k)$$

$$\leqslant \sum_{i=2}^{n} d(i, 1) p_i \sum_{j=i+1}^{n} p_j \prod_{k=i+1}^{j-1} (1-p_k) + \sum_{j=2}^{n} d(1, j) p_j \sum_{i=1}^{j-1} p_i \prod_{k=i+1}^{j-1} (1-p_k)$$

$$= \sum_{i=2}^{n} d(i, 1) p_i \left(1 - \prod_{k=i+1}^{n} (1-p_k)\right) + \sum_{j=2}^{n} d(1, j) p_j \left(1 - \prod_{k=2}^{j-1} (1-p_k)\right),$$

from which (17) follows. □

Note that the result is independent of the tour used. For the special case of $p_i = p$ and symmetric cost function we find that any tour τ satisfies $E[L_\tau] \leqslant 2pL_{T_*}$ where T_* is the length of the star tree rooted at node 1. Therefore,

$$p\left(1 - O(1-p)^n\right) L_{\mathrm{TSP}} \leqslant E\left[L_{\tau_p}\right] \leqslant 2pL_{T_*},$$

which means that if we fix the cost matrix and vary p, $E[L_{\tau_p}] = \Theta(p)$ where the constant depends on properties of the cost matrix. Following this path of analysis, in the following section we investigate the dependence of the PTSP on the coverage probability p.

3.2. Functional properties of the PTSP

In this subsection we examine the case with equal probabilities $p_i = p$. In this case we are trying to find a tour τ_p that minimizes the expression

$$\frac{E\left[L_{\tau_p}\right]}{p^2} = f(p) \triangleq \min_\tau f_\tau(p) \triangleq \min_\tau \left\{ \sum_{r=0}^{n-2} (1-p)^r L_\tau^{(r)} \right\}. \tag{18}$$

Intuitively, $f(p)$ examines the dependence of the optimal PTSP length as a function of the coverage probability p. Expression (18) is clearly a function of the coverage probability p. For different values of p the corresponding optimal probabilistic tours which minimize (18) are different. We first investigate the properties of the function $f(p)$. Some initial observations are stated in the following proposition.

Proposition 5. *If the cost matrix is positive, the function $f(p)$ is continuous, decreasing, piecewise differentiable and piecewise convex.*

Proof. Each function $f_\tau(p)$ is analytic, so $f(p)$ must be continuous and piecewise analytic, since it is the minimum of a finite number of analytic functions. For each τ it is easily established from (18) that if the costs are positive, then

$$\frac{\mathrm{d}}{\mathrm{d}p} f_\tau(p) < 0, \qquad \frac{\mathrm{d}^2}{\mathrm{d}p^2} f_\tau(p) > 0,$$

so each $f_\tau(p)$ is decreasing and convex. Since $f(p)$ is continuous, it must be globally decreasing, but since $f(p)$ is only piecewise differentiable, it need only be piecewise convex. □

As $p \to 1$, the optimal PTSP solution is clearly the optimum solution to TSP. But what is the limit of the solution as $p \to 0$? Since $f_\tau(0) = \sum_{r=0}^{n-2} L_\tau^{(r)} = \sum_{i,j} d(i, j) = K_0$, which is a constant independent of the tour, all tours are optimal at $p = 0$. As $p \to 0$,

$$f_\tau(p) = K_0 - p \sum_{r=0}^{n-2} r L_\tau^{(r)} + p^2 \sum_{r=0}^{n-2} \binom{r}{2} L_r^{(r)} + O(p^3).$$

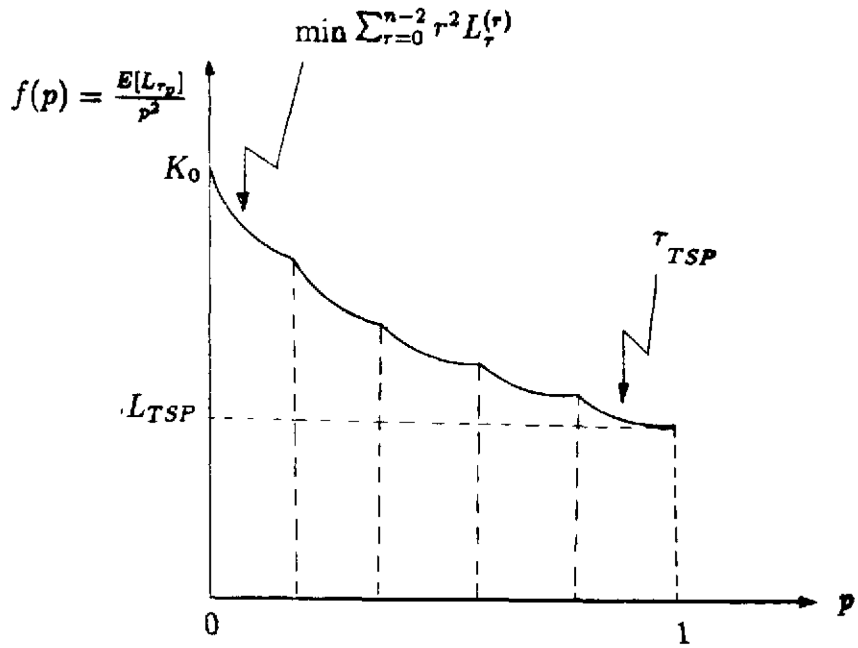

Figure 2. The PTSP as a function of the coverage probability p

If the cost matrix is asymmetric, then the optimal PTSP tour as $p \to 0$ is the one maximizing $\sum_{r=0}^{n-2} rL_\tau^{(r)}$. If the cost matrix is symmetric, however, then $L_\tau^{(r)} = L_\tau^{(n-2-r)}$. Therefore, if $n = 2k+1$, then

$$\sum_{r=0}^{n-2} rL_\tau^{(r)} = \sum_{r=0}^{k-1} rL_\tau^{(r)} + \sum_{r=k}^{n-2} rL_\tau^{(r)} = (n-2)\sum_{r=0}^{k-1} L_\tau^{(r)} = \tfrac{1}{2}(n-2)K_0,$$

which is independent of the tour. Similarly, if $n = 2k$, $\sum_{r=0}^{n-2} rL_\tau^{(r)} = \frac{1}{2}(n-2)K_0$.

Therefore, as $p \to 0$, under the symmetry assumption the PTSP tour is the one minimizing $\sum_{r=0}^{n-2} \binom{r}{2} L_\tau^{(r)}$, which is equivalent, because of symmetry, to the one minimizing $\sum_{r=0}^{n-2} r^2 L_\tau^{(r)}$.

We can now combine Proposition 5 and the above observations to sketch a possible graph of the function $f(p) = E[L_{\tau_p}]/p^2$ in Figure 2.

3.3. Relation between the PTSP and the PMST

The minimum spanning tree plays an important role in the best known heuristic [8] for the TSP with triangle inequality. In this subsection we explore the relationship between the PTSP and the probabilistic

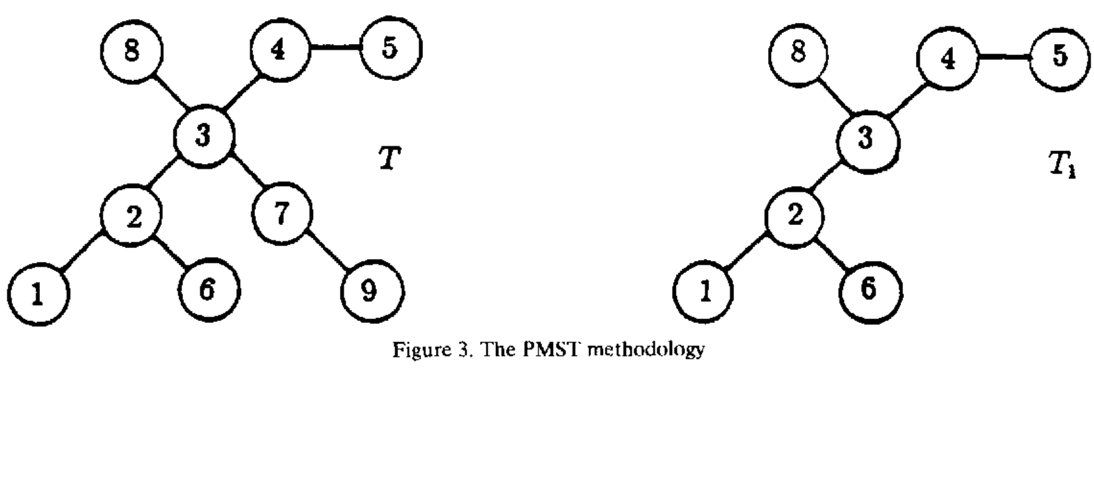

Figure 3. The PMST methodology

minimum spanning tree (PMST) studied in [6], which is defined as follows: Given a weighted graph $G=(V, E)$ and a probability of presence $p(S)$ for each subset S of V, we want to construct an *a priori* spanning tree of minimum expected length in the following sense: on any given instance of the problem, delete the vertices and their adjacent edges among the set of absent vertices provided that the tree remains connected. The problem of finding an a priori spanning tree of minimum expected length is the probabilistic minimum spanning tree (PMST) problem. In order to clarify the definition of the PMST problem, consider the example in Figure 3. If the a priori tree is T and nodes 2, 7, 9 are the only ones not present, the tree becomes T_1.

Consider a tree T. We say that a tour τ is induced by a tree T if and only if for all subtrees of T the nodes in the subtree are visited in τ before another node outside the subtree is visited in τ. Let $I \triangleq \{\tau \mid \tau$ is induced by $T\}$. For example, in Figure 3 the tour 832167945 is induced by the tree while the tour 832791645 is not, since node 7 is visited before the subtree emanating from node 2 is visited. Clearly, when the triangle inequality holds, $L_\tau \leqslant 2L_T$ for all $\tau \in I$. In the following proposition we prove that the same inequality holds in the probabilistic case.

Proposition 6. Under the triangle inequality, if τ is induced by a tree T, then

$$E[L_\tau] \leqslant 2E[L_T] \tag{19}$$

where the first expectation means in the PTSP *sense and the second in the* PMST *sense.*

Proof. Let $L_\tau(S)$ and $L_T(S)$ denote the lengths of the tour $\tau(S)$ in the PTSP sense and the tree $T(S)$ in the PMST sense, respectively, when only vertices from the set S are present. We claim that $L_\tau(S) \leqslant 2L_T(S)$.
The PTSP strategy maintains the order of visiting the points which are present. Furthermore, the PMST strategy changes the structure of the a prior tree in a way that disconnects only absent subtrees. The structure of the present subtrees is maintained. Therefore, the property that for all subtrees the nodes in the subtree are visited in the tour before another node outside the subtree is visited in τ is maintained. Then, if S_1 is the set S along with the nodes needed by the PMST strategy so that the tree does not become disconnected if only nodes from S are present, then $\tau(S_1)$ is induced by $T(S_1)$ and thus $L_\tau(S_1) \leqslant 2L_T(S_1)$. But $L_T(S) = L_T(S_1)$ and $L_\tau(S) \leqslant L_\tau(S_1)$ since $\tau(S)$ is a subsequence of $\tau(S_1)$. Hence $L_\tau(S) \leqslant 2L_T(S)$. Multiplying by $p(S)$ and adding for all sets S, (19) follows. □

Note that the above result is quite general since it does not make any assumptions on the probabilities of presence of points. It holds even if the probabilities of presence of sets of points are correlated. A very interesting consequence of the above proposition is that we can compare the PTSP and the PMST strategy if the triangle inequality holds.

Theorem 7. *If the cost matrix is symmetric and satisfies the triangle inequality, then*

$$E[\Sigma_{\text{MST}}] \leqslant E\left[L_{\tau_p}\right] \leqslant 2E\left[L_{T_p}\right] \tag{20}$$

where τ_p and T_p are the optimal PTST *and* PMST, *respectively and $E[\Sigma_{\text{MST}}]$ is the expectation of the* MST *re-optimization strategy.*

Proof. Consider a tour τ induced by T_p. From Proposition 6, $E[L_\tau] \leqslant 2E[L_{T_p}]$. Since τ_p is the optimum probabilistic tour, $E[L_{\tau_p}] \leqslant E[L_\tau]$, from which the right inequality follows. Furthermore, since $L_{\text{MST}}(S) \leqslant L_{\text{TSP}}(S) \leqslant L_{\tau_p}(S)$, the left inequality follows. □

An interesting observation is that Theorem 7 provides a way of comparing the MST re-optimization strategy Σ_{MST} and the PMST strategy when the triangle inequality holds. It is quite surprising that this analysis used the PTSP.

In the deterministic case, $L_{\mathrm{MST}} \leqslant L_{\mathrm{TSP}} \leqslant 2L_{\mathrm{MST}}$, and these inequalities offer a heuristic of worst-case guarantee of 2 for the TSP. In the probabilistic case it is not in general true that $E[L_{T_p}] \leqslant L_{\tau_p}$ and furthermore the PMST is an NP-*hard* problem. But when can the PMST problem help us in providing some heuristics for the PTSP?

In addressing this question we consider the MST T_*. If $T_*(S)$ and $T_{\mathrm{MST}}(S)$ are the trees produced by the PMST strategy when the a priori tree is T_* and under the re-optimization strategy, respectively, then the following lemma is useful in the analysis.

Lemma 8. *If an edge $e=(i, j) \in T_*$ and $i, j \in S$, then there exists an* MST *in the instance S with $e \in T_{\mathrm{MST}}(S)$. Furthermore, if the distances are distinct, $T_{\mathrm{MST}}(S)$ is unique.*

Proof. Assume that for all the MST's on instance S, $\in \notin T_{\mathrm{MST}}(S)$. Then $e \cup T_{\mathrm{MST}}(S)$ creates a cycle C, in which $d(e) > d(a)$, $\forall a \in C$. Consider now the MST T_* on the original instance. If we delete e from T_*, then T_* will be disconnected into two trees T_1, T_2. Since the cycle C joins nodes $i \in T_1$ and $j \in T_2$ there exists an edge $a=(u, v) \in C$ such that if we add a in $T_* - e$, a new tree will be created. Since $d(e) > d(a)$ the new tree will have less cost than T_*, which contradicts the optimality of T_*. □

Based on this lemma, we want to compare the lengths of $T_*(S)$ and $T_{\mathrm{MTT}}(S)$.

Proposition 9. *Suppose that the cost matrix has distinct costs and satisfies the triangle inequality. If T_* is the* MST, *then*

$$\frac{E[L_{T_*}]}{E[\Sigma_{\mathrm{MST}}]} < d \tag{21}$$

where d is the diameter of T_*, i.e., the maximum number of edges joining two nodes in T_*.

Proof. We assume that the costs are distinct so that $T_{\mathrm{MST}}(S)$ is unique. This is not a truly restrictive assumption, because a perturbation ε will result in distinct costs. From Lemma 8, $T_{\mathrm{MST}}(S)$ contains all edges $(i, j) \in T_*$ if $(i, j) \in S$. The set S induces a partition in T_*, namely the connected components of T_*, when only nodes from S are present. For example, in Figure 4b there are three connected

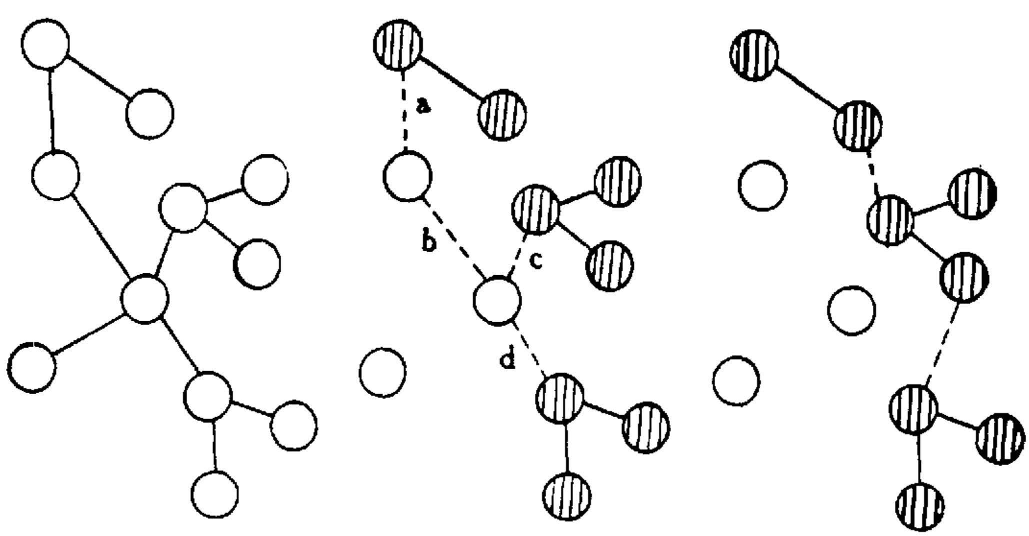

(a) T_* (b) $T_*(S)$ (c) $T_{MST}(S)$

Figure 4. The trees $T_*(S)$ and $T_{\mathrm{MST}}(S)$

components. In $T_*(S)$ these components are connected by paths from T_* and in $T_{\text{MST}}(S)$ by edges (see Figure 4b,c).

Let $k(S)$ be the number of components of T_* induced by S (in Figure 4b, $k(S)=3$), $\theta(S)$ be the length of these components and e_i, $i=1,\dots,k(S)-1$, be the set of edges in $T_{\text{MST}}(S)$ joining the $k(S)$ components (see Figure 4c). Finally, let $B \triangleq \{e \in T_*$ joining the $k(S)$ components of T_* in $T_*(S)\}$. In the example of Figure 3b, $B=\{a, b, c, d\}$. With these definitions the corresponding lengths become:

$$L_{T_*}(S) = \sum_{e \in B} d(e) + \theta(S), \qquad L_{\text{MST}}(S) = \sum_{i=1}^{k(S)-1} d(e_i) + \theta(S).$$

But for all $e_i \in T_{\text{MST}}(S) - T_*(S)$, $e_i \notin T_*$. Therefore, since T_* is the MST, $d(e_i) > d(e)$ for all e belonging to the cycle in T_* generated by adding edge e_i to T_*. Consider now two components joined by edge e_i in $T_{\text{MST}}(S)$ and by the path $f_i=(a_1,\dots,a_r)$ in $T_*(S)$. As a result, $rd(e_i) > d(f_i)$, from which we find that $r_{\max} d(e_i) > d(f_i)$, where $r_{\max}$ is the maximum number of edges in a path joining two components in T_*. Adding with respect to $i=1,\dots,k(S)-1$, we find

$$r_{\max} \sum_{i=1}^{k(S)-1} d(e_i) > \sum_{i=1}^{k(S)-1} d(f_i) \geqslant \sum_{e \in B} d(e),$$

since each edge in $e \in B$ belongs to at least one path. As a result, $r_{\max} L_{\text{MST}}(S) > L_{T_*}(S)$. Thus, if we define $d \triangleq$ the largest number of edges joining two nodes in the MST (the diameter of T_*), then $r_{\max} \leqslant d$, which results in $dL_{\text{MST}}(S) > L_{T_*}(S)$. Summing over all sets S, (21) follows. □

Based on the above analysis we propose the following heuristic for the PTSP, for symmetric cost matrices satisfying the triangle inequality:

1. Find the MST T_*.
2. T_* induces a tour τ_h, which is the proposed solution.

The performance of the heuristic can be described as follows:

Theorem 10.

$$\frac{E\left[L_{\tau_h}\right]}{E\left[L_{\tau_p}\right]} < 2d.$$

Proof. From (19)–(21),

$$\frac{E\left[L_{\tau_h}\right]}{E\left[L_{\tau_p}\right]} \leqslant 2\frac{E\left[L_{T_*}\right]}{E\left[\Sigma_{\text{MST}}\right]} < 2d. \qquad \square$$

The *MST* heuristic behaves well (constant-guarantee) for d small. For example, if the MST is a star tree, then $d=2$. However, d can in general be $O(n)$.

The above discussion tried to investigate whether the PMST problem can be helpful in addressing the PTSP. Surprisingly, we identified some cases where it can indeed be useful.

3.4. A special case

In this subsection we exploit the closed-form expressions (2) to find a special case in which we can solve the PTSP in polynomial time. This is the 'constant PTSP'.

Let the probabilities of presence be p_i for every node i. Let the cost matrix be of the form $d(i, j) = u_i + v_j$. In [22] this special case is defined as the constant TSP, since all tours have the same length. We prove that the same property holds for the PTSP case:

Proposition 11. *If node i is present with probability p_i, any tour τ has the same expected length*:

$$E[L_\tau] = \sum_{i=1}^{n} (u_i + v_i) p_i \Big(1 - \prod_{k \neq i} (1 - p_k)\Big). \tag{22}$$

Proof. By renumbering the nodes, consider any tour $\tau = (1, 2, \ldots, n, 1)$. From (2) we have:

$$E[L_\tau] = \sum_{i=1}^{n} \sum_{j=i+1}^{n} (u_i + v_j) p_i p_j \prod_{k=i+1}^{j-1} (1 - p_k)$$

$$+ \sum_{j=1}^{n} \sum_{i=1}^{j-1} (u_j + v_i) p_i p_j \prod_{k=j+1}^{n} (1 - p_k) \sum_{k=1}^{i-1} (1 - p_k).$$

After some simple algebraic manipulations we find that

$$E[L_\tau] = \sum_{i=1}^{n} (u_i + v_i) p_i \Big(1 - \prod_{k \neq i} (1 - p_k)\Big).$$

Note that (22) does not depend on the tour used, i.e., all tours have the same expected length. □

4. Relationship between the PTSP and the TSP re-optimization strategy

We have already observed that the PTSP provides an efficient strategy to update traveling salesman tours when problem instances are modified probabilistically because of the absence of certain nodes from the graph. If τ_p is the optimal PTST, then in the instance S, i.e., when only nodes in the set S are present, the strategy produces a tour $\tau_p(S)$ with length $L_{\tau_p}(S)$, which is the length of the tour which visits nodes from the set S of present nodes in the same order as they appear in the a priori tour τ_p. In the context of the following discussion the letter Σ denotes the strategy used.

An alternative strategy is clearly the re-optimization strategy Σ_{TSP}, in which we find the minimum TST of the set of present nodes in every instance. Let $L_{TSP}(S)$ be the length of the TST of the nodes in the set S. We then define the expectation of this re-optimization strategy as follows:

$$E[\Sigma_{TSP}] \triangleq \sum_{S \subseteq V} p(S) L_{TSP}(S) \tag{23}$$

where $p(S)$ was defined earlier to be the probability that only nodes in S are present.

In the introductory section we have discussed some of the possible reasons that one might use the PTSP strategy rather than the re-optimization strategy Σ_{TSP}. In general it is extremely difficult to find an efficient way to compute $E[\Sigma_{TSP}]$, since we have to compute a sum of $O(2^n)$ terms, each of which requires the solution of an NP-hard problem. Therefore, it is apparent that the evaluation of $E[\Sigma_{TSP}]$ is computationally infeasible. Our hope then lies in proving some useful bounds. The following proposition relates the strategy of re-optimization and the seemingly irrelevant nearest neighbor heuristic.

Proposition 12. *If $p_1 = 1$, all the other points have the same coverage probability p and the triangle inequality holds, then*

$$E[\Sigma_{TSP}] \geqslant \frac{2p}{\log(np + 1 - 2p) + 2} L_{NN}, \tag{24}$$

where L_{NN} is the tour length produced by the nearest-neighbor heuristic.

Proof. From definition (23),

$$E[\Sigma_{\text{TSP}}] \triangleq \sum_{S\subseteq V} p(S) L_{\text{TSP}}(S) = \sum_{k=1}^{n-1} p^k(1-p)^{n-1-k} \sum_{S\subseteq V-\{1\},|S|=k} L_{\text{TSP}}(S). \tag{25}$$

Following [28], Lemma 1: If for each node i there exists a number l_i such that $d(i, j) \geq \min(l_i, l_j)$, $l_1 \leqslant \frac{1}{2}L_{\text{TSP}}(S)$, then

$$(\lceil \log|S| \rceil + 1) L_{\text{TSP}}(S) \geqslant 2 \sum_{i\in S} l_i.$$

Summing over all $S \subseteq V-\{1\}$, $|S| = k$, we obtain

$$(\lceil \log k \rceil + 1) \sum_{S\subseteq V-\{1\},|S|=k} L_{\text{TSP}}(S) \geqslant 2 \sum_{S\subseteq V-\{1\},|S|=k} \sum_{i\in S} l_i \geqslant 2\binom{n-2}{k-1} \sum_{i=1}^{n} l_i,$$

since each node $i = 2,\dots, n$ is present in $\binom{n-2}{k-1}$ sets and node 1 is present in $\binom{n-1}{k} \geqslant \binom{n-2}{k-1}$ sets. From (25) we then obtain

$$E[\Sigma_{\text{TSP}}] \geqslant 2\left(\sum_{l=1}^{n} l_i\right) \sum_{k=1}^{n-1} \binom{n-2}{k-1} p^k(1-p)^{n-1-k} \frac{1}{\lceil \log k \rceil + 1}.$$

Hence,

$$E[\Sigma_{\text{TSP}}] \geqslant 2p\left(\sum_{l=1}^{n} l_i\right) \sum_{k=0}^{n-2} \binom{n-2}{k} p^k(1-p)^{n-k-2} \frac{1}{\log(k+1)+2}$$
$$= 2p\left(\sum_{l=1}^{n} l_i\right) E\left\{\frac{1}{\log(X+1)+2}\right\},$$

where the random variable X is binomial with parameters $n-2$ and p. Applying Jensen's inequality to the convex function $f(x) = 1/(\log(x+1)+2)$ we obtain

$$E[\Sigma_{\text{TSP}}] \geqslant 2p\left(\sum_{l=1}^{n} l_i\right) E[f(X)] \geqslant 2p\left(\sum_{l=1}^{n} l_i\right) f(E[X])$$
$$= 2p\left(\sum_{l=1}^{n} l_i\right)\left[\frac{1}{\log[(n-2)p+1]+2}\right].$$

The question now is to find good bounds on $\sum_{l=1}^{n} l_i$. Consider the application of the nearest-neighbor heuristic to the original instance. If we let l_i be the length of the edge leaving node i and going to the node selected as the nearest neighbor of i, then clearly $d(i, j) \geqslant l_i$ if node i was selected by the algorithm before node j, or $d(i, j) \geqslant l_j$ if node j was selected before node i. Since one of the nodes was selected first, we have $d(i, j) \geqslant \min(l_i, l_j)$. Also for any problem instance S, $l_1 \leqslant l_1(S)$, the label of the 'black' node 1 if the nearest-neighbor heuristic was applied to instance S, because at this step whatever choices for the nearest neighbor of node 1 are available at instance S were also available in the initial instance. As a result, $2l_1 \leqslant L_{\text{TSP}}(S)$ and thus the hypotheses of Lemma 1 of [28] are satisfied. But in this case $\sum_{l=1}^{n} l_i = L_{\text{NN}}$, which proves (24). □

From a worst-case perspective, if we do not have the triangle inequality it is trivial to construct examples with $E[L_{\tau_p}]/E[\Sigma_{\text{TSP}}] \to \infty$. In the next section we prove that in the Euclidean case, $E[L_{\tau_p}]/E[\Sigma_{\text{TSP}}] = O(\log n)$, even in the case where we have dependencies among the probabilities of sets of points being present.

From a probabilistic perspective, we prove in Section 6 that if the points are randomly distributed in the unit square, then these two strategies have the same expectation almost surely. As a result of the previous discussion, one can observe that the PTSP strategy is a natural one and produces tours which under the triangle inequality are close to the optimal ones.

5. Worst-case analysis for PTSP heuristics

In trying to analyze the worst-case behavior of heuristics for the PTSP, we should not hope to find a constant-guarantee polynomial time heuristic for the PTSP on a general cost matrix, since the corresponding problem for the TSP is NP-hard. But, what if the triangle inequality holds? We examine this question next.

5.1. TSP heuristics applied to PTSP

In Theorem 10 we have showed that the MST heuristic for the TSP works well for the PTSP if the diameter of the MST is small. Following this path of analysis, how does the best known heuristic for the TSP, i.e., the Christofides heuristic, work for the PTSP? Or more generally, how would a constant-guarantee heuristic for the TSP behave for the PTSP?

Proposition 13. *For a symmetric cost matrix satisfying the triangle inequality, if τ_H is the tour produced by any heuristic H which is within a constant c of the* TST, *then for the case $p_i = p$, $i = 1, 2, \ldots, n$,*

$$\frac{E[L_{\tau_H}]}{E[L_{\tau_p}]} \leqslant \frac{c}{p(1 - O(1-p)^n)}. \tag{26}$$

Proof. From (16), $E[L_{\tau_H}] \leqslant L_{\tau_H} \leqslant cL_{\text{TSP}}$, and from (8), $E[L_{\tau_p}] \geqslant p(1 - O(1-p)^n)L_{\text{TSP}}$. Combining these bounds, (26) follows. □

The result simply confirms our intuition that when p is large (close to 1), then the PTSP is close to the TSP and therefore the approximation of the PTSP with any constant-guarantee heuristic should work well. For $p \to 0$ the bound is not informative and for this reason we prove the following:

Proposition 14. *Under the triangle inequality every tour τ satisfies*

$$\frac{E[L_\tau]}{E[L_{\tau_p}]} \leqslant \tfrac{1}{2}n. \tag{27}$$

Proof. Consider a tour τ, the expectation of which is given by $E[L_\tau] = \sum_{S \subseteq V} p(S) L_\tau(S)$. Since for every $(i, j) \in \tau(S)$ there exist paths $(i, i_2, \ldots, i_l, j)$, $(j, i_{l+2}, \ldots, i_{|S|}, i)$ in $\tau_p(S)$, then $L_\tau(S) \leqslant \frac{1}{2}|S|\, L_{\tau_p}(S)$. Therefore,

$$\frac{E[L_\tau]}{E[L_{\tau_p}]} = \frac{\sum_{S \subseteq V} p(S) L_\tau(S)}{\sum_{S \subseteq V} p(S) L_{\tau_p}(S)} \leqslant \frac{\sum_{S \subseteq V} p(S) \frac{1}{2}|S|\, L_{\tau_p}(S)}{\sum_{S \subseteq V} p(S) L_{\tau_p}(S)} \leqslant \tfrac{1}{2}n \qquad \square$$

For the case that $p_i = p$, which might depend on n ($p = p(n)$), and $c(i, j) \leqslant B$, then as $np \to \infty$,

$$\frac{E[L_\tau]}{E[L_{\tau_p}]} \leqslant \frac{\sum_{S \subseteq V} p(S) \frac{1}{2}|s|\, L_{\tau_p}(S)}{\sum_{S \subseteq V} p(S) L_{\tau_p}(S)} \sim \frac{np \sum_{S \subseteq V, |S| = np} L_{\tau_p}(S)}{2 \sum_{S \subseteq V, |S| = np} L_{\tau_p}(S)} = \tfrac{1}{2}np, \tag{28}$$

since only the contribution from the terms $|S| = np$ is asymptotically important. Using this limiting behavior we can give an asymptotic characterization of the behavior of heuristics for the PTSP which under triangle inequality are within a constant of the optimal TST.

Theorem 15. *For the case $p_i = p$, as $np \to \infty$, the expected tour length obtained through every heuristic which is of constant-guarantee for the* TSP *is within* $O(\sqrt{n})$ *of the optimal* PTSP *tour.*

Proof. Let heuristic H produce a tour τ_H with the property $L_{\tau_H}/L_{\text{TSP}} \leqslant c$. Combining bounds (26) and (28), we find that as $n \to \infty$,

$$\frac{E\left[L_{\tau_H}\right]}{E\left[L_{\tau_p}\right]} \leqslant \min\left\{\tfrac{1}{2}np,\ \frac{c}{p\left(1-O(1-p)^n\right)}\right\}.$$

The worst-case is when both terms are of the same order, which happens at $p = O(1/\sqrt{n})$ and leads to

$$\frac{E\left[L_{\tau_H}\right]}{E\left[L_{\tau_p}\right]} \leqslant O(\sqrt{n}). \qquad \square \tag{29}$$

The critical question then is: Can we do better? Is there a constant-guarantee heuristic for the PTSP with the triangle inequality? In trying to address these questions, we next consider the space-filling curve heuristic proposed in [1] for the TSP.

5.2. *The space-filling curve heuristic*

The space-filling curve heuristic for the Euclidean TSP is an $O(n \log n)$ method based on sorting which works as follows:

1. Given the n coordinates (x_i, y_i) of the points in the plane, compute the number $f(x_i, y_i)$ for each point. The function $f: \mathbb{R}^2 \to R$ is called the Sierpinski curve (for details on the computation of $f(x, y)$ see [1]).
2. Sort the numbers $f(x_i, y_i)$ and visit the initial points in the same order, producing a tour τ_{SF}.

The key property of the space-filling curve heuristic that makes its analysis for the PTSP possible is the following: Consider an instance S of the problem. Suppose the space-filling curve heuristic produces a tour $\tau_{\text{SF}}(S)$ if we run the heuristic on the instance S. Consider now the tour τ_{SF} produced by the heuristic on the original instance. What is the tour that the PTSP strategy would produce in instance S if the a priori tour is τ_{SF}?

The answer is precisely $\tau_{\text{SF}}(S)$, because sorting has the property of preserving the order, which is exactly the property of the PTSP strategy as well. Based on this critical observation we can then analyze the space-filling curve heuristic.

Theorem 16. *For the Euclidean* PTSP, *the space-filling curve heuristic produces a tour τ_{SF} with the property*

$$\frac{E\left[L_{\tau_{\text{SF}}}\right]}{E\left[L_{\tau_p}\right]} \leqslant \frac{E\left[L_{\tau_{\text{SF}}}\right]}{E\left[\Sigma_{\text{TSP}}\right]} = O(\log n). \tag{30}$$

Proof. In Platzman and Bartholdi [26] it is proven that the length of the space-filling curve heuristic satisfies $L_{\tau_{\text{SF}}}/L_{\text{TSP}} = O(\log n)$. Consider an instance S of the problem. If the space-filling curve heuristic is applied to the instance S, it will similarly produce a tour $\tau_{\text{SF}}(S)$ with length $L_{\tau_{\text{SF}(S)}}/L_{\text{TSP}}(S) = O(\log |S|) = O(\log n)$. But since $\tau_{\text{SF}}(S)$ is the tour produced by the PTSP strategy at instance S, we have

$$\frac{E\left[L_{\tau_{\text{SF}}}\right]}{E\left[\Sigma_{\text{TSP}}\right]} = \frac{\sum_{S \subseteq V} p(S) L_{\tau_{\text{SF}}}(S)}{\sum_{S \subseteq V} p(S) L_{\text{TSP}}(S)} \leqslant \frac{\sum_{S \subseteq V} p(S) O(\log n) L_{\text{TSP}}(S)}{\sum_{S \subseteq V} p(S) L_{\text{TSP}}(S)} = O(\log n). \qquad \square$$

Note that the above result does not depend on the probabilities of points being present. It holds even if there are dependencies on the presence of the points. Observe also that the heuristic ignores the probabilistic nature of the problem but surprisingly produces a tour which is globally (in every instance) close to the optimal.

A further corollary of the space-filling curve heuristic is that we can compare the PTSP and the re-optimization strategies from a worst-case perspective. For the Euclidean PTSP, $E[L_{\tau_p}]/E[\Sigma_{\mathrm{TSP}}] = O(\log n)$. Platzman and Bartholdi [26] conjecture that the space-filling curve heuristic is a constant-guarantee heuristic. If the above conjecture is correct, then it would follow that $E[L_{\tau_{\mathrm{SF}}}]/E[\Sigma_{\mathrm{TSP}}] = O(1)$. Unfortunately, Bertsimas and Grigni [7] refute the conjecture by exhibiting an example which achieves the $\log n$ bound.

6. Probabilistic analysis

In this section we depart from the examination of the combinatorial (deterministic) properties of the problem to examine its behavior probabilistically. In trying to understand the behavior of the TSP heuristics for the PTSP from many viewpoints, we perform a probabilistic analysis of the space-filling curve heuristic and of one of the most common heuristics for the TSP, the nearest-neighbor heuristic, variations of which were proposed in [16] and [18] as possible solutions to the PTSP. We prove that, even on average, the nearest-neighbor heuristic produces poor solutions to the PTSP.

6.1. The space-filling curve heuristic in the random Euclidean model

In Theorem 16 we have shown that the space-filling curve heuristic is $O(\log n)$ from the strategy of re-optimization in the worst case. Our goal in this subsection is to find how the heuristic behaves asymptotically if the n points $X^{(n)} = (X_1, \ldots, X_n)$ are uniformly and independently distributed in the unit squared. Platzman and Bartholdi [26] proved that with probability 1,

$$\limsup_{n\to\infty} \frac{L^n_{\tau_{\mathrm{SF}}}(X^{(n)})}{\sqrt{n}} = \beta_1, \qquad \liminf_{n\to\infty} \frac{L^n_{\tau_{\mathrm{SF}}}(X^{(n)})}{\sqrt{n}} = \beta_2 \tag{31}$$

where β_1, β_2 are constants such that $\beta_1 - \beta_2 \leqslant 0.0001$. From simulation experiments it is found that $\beta_i \sim 0.956$, $i = 1, 2$. By a simple modification of the method in [26] we can prove the following theorem:

Theorem 17. *Let $X^{(n)}$ be a sequence of points distributed independently and uniformly in the unit square and p the coverage probability of each point. Then with probability* 1,

$$\limsup_{n\to\infty} \frac{E\left[L^n_{\tau_{\mathrm{SF}}}(X^{(n)})\right]}{\sqrt{n}} = \beta_1\sqrt{p}, \qquad \liminf_{n\to\infty} \frac{E\left[L^n_{\tau_{\mathrm{SF}}}(X^{(n)})\right]}{\sqrt{n}} = \beta_2\sqrt{p} \tag{32}$$

where $E[L^n_{\tau_{\mathrm{SF}}}(X^{(n)})]$ is the expectation in the PTSP sense of the tour produced by the space-filling curve heuristic on the points $X^{(n)}$.

As a result, the space-filling heuristic produces solutions which are close asymptotically to the re-optimization strategy in terms of performance.

6.2. Probabilistic analysis of the nearest neighbor heuristic applied to PTSP

In this subsection we investigate the possibility that TSP heuristics might work well for the PTSP on average, although they can be bad in the worst case. A natural way to investigate this type of question is to perform a probabilistic analysis under some reasonable probabilistic model. We therefore consider the random length model, in which we assume that the costs $d(i, j)$ are uniformly distributed in (0, 1) and that $p_1 = 1$ and $p_i = p$. Let the tour found by the nearest-neighbor heuristic be $\tau_{NN} = (1, 2, \ldots n, 1)$. Then if E_u means the expectation with respect to the random cost lengths and, as usual, and E means the expectation in the PTSP sense, then we can prove the following theorem:

Theorem 18. *Under the random-length model,*

$$\lim_{n\to\infty} \frac{E_u E(L_{\tau_{NN}})}{n} = \tfrac{1}{2}p(1-p). \tag{33}$$

Proof. From (2),

$$E[L_{\tau_{NN}}] = pd(1, 2) + p\sum_{r=2}^{n-1}(1-p)^{r-1}d(1, r+1) + p^2\sum_{r=2}^{n-1}d(r, r+1)$$
$$+p^2\sum_{i=2}^{n-2}\sum_{r=1}^{n-i-1}(1-p)^r d(i, i+r+1) + p\sum_{r=2}^{n}(1-p)^{n-r}d(r, 1). \tag{34}$$

We will need some ideas from order statistics in order to compute $E_u E[L_{\tau_{NN}}]$. Let X_i^n, $i = 1, \ldots, n$, be a sequence of n independent random variables and $X_{(i)}^n$ be the i-th-order statistic. We need the following results:

$$E_u(X_{(1)}^n) = \frac{1}{n+1}, \qquad E_u(X_i^n \mid X_i^n \neq X_{(1)}^n) = \frac{n+2}{2(n+1)}.$$

As a result, since $d(i, i+1) = \min[d(i, i+1), \ldots, d(i, n)]$,

$$E_u[d(i, i+1)] = E_u[X_{(1)}^{n-i}] = \frac{1}{n-i+1},$$
$$E_u[d(i, i+r+1)] = E_u[X_i^{n-i} \mid X_i^{n-i} \neq X_{(1)}^{n-i}] = \frac{n-i+2}{2(n-i+1)},$$
$$E_u[d(i, 1)] = \tfrac{1}{2}.$$

With these equalities, (34) becomes:

$$E_u E[L_{\tau_{NN}}] = p\frac{1}{n} + p\sum_{r=2}^{n-1}(1-p)^{r-1}\frac{n+1}{2n} + p^2\sum_{r=2}^{n-1}\frac{1}{n-r+1}$$
$$+p^2\sum_{i=2}^{n-2}\sum_{r=1}^{n-i-1}(1-p)^r\frac{n-i+2}{2(n-i+1)} + p\sum_{r=2}^{n}(1-p)^{n-r}\tfrac{1}{2}.$$

If we fix p and let $n \to \infty$, then there will only be a nonzero contribution in the limit from the fourth term. As a result, (33) follows. □

In order to understand the behavior of this heuristic we need some estimate of $E[\Sigma_{TSP}]$. Karp and Steele [20] prove that $E_u(L_{TSP}^n) \leqslant E_u(L_a^n) + O(1/\sqrt{n})$, where L_a^n is the cost of the minimum assignment problem. Since $E_u(L_a^n) \leqslant 2$, it follows that as $n \to \infty$, $E_u(L_{TSP}^n) \leqslant 2$, provided that the limit exists. As a result, we can easily prove that as $n \to \infty$, $E_u(\Sigma_{TSP}^n) \leqslant 2$, i.e., the expectation of the strategy of re-optimization is bounded, while the nearest-[neighbor heuristic is $\Theta(n)$.

7. Practical optimization

In our numerical experiments we have obtained near-optimal solutions to Euclidean PTSP's by means of two different types of heuristics. The first, the space-filling curve heuristic, we have already described

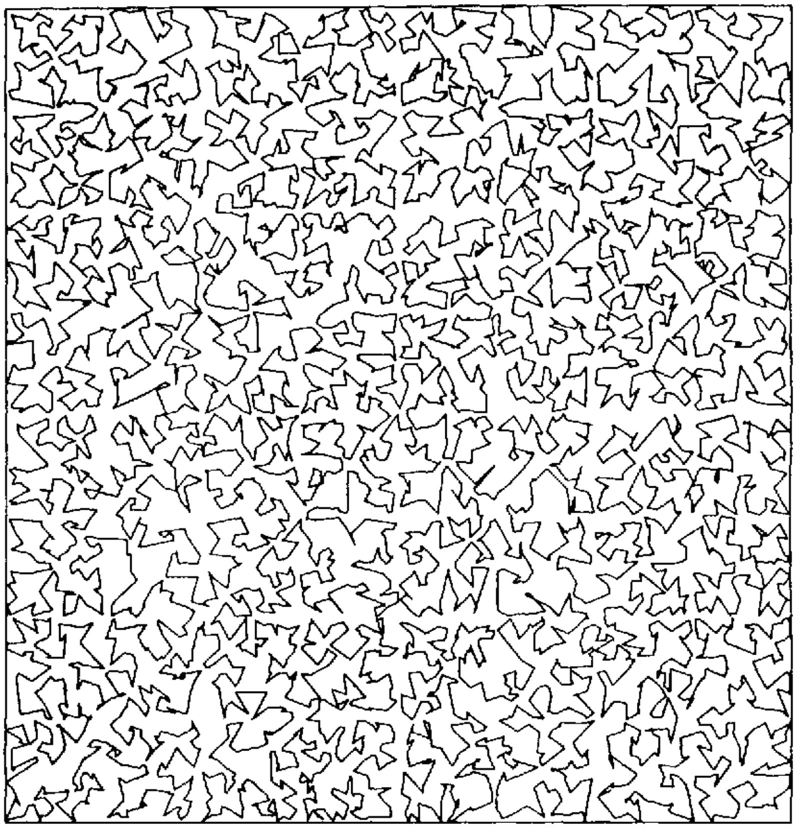

Figure 5. The space-filling curve heuristic applied to a problem with 5000 uniformly distributed nodes

at length in Section 5.2. Our implementation uses heapsort for the sorting part of the procedure, and thus requires only $O(n \log n)$ time to find a nearly optimal tour τ_{SF}. Interestingly enough, this is even faster than the computation of the expected length of that tour, $E[L_{\tau_{SF}}]$, which requires $O(n^2)$ time. A sample solution involving 5000 nodes is shown in Figure 5. Since the computed tour τ_{SF} is independent of the probabilities p_i, the space-filling curve heuristic can be used when these probabilities are not all the same, or even when they are not accurately known.

For problems involving equal probabilities $p_i = p$ and not more than a few hundred nodes, we have had considerable success with two separate iterative improvement algorithms based on the idea of local optimality. Given a tour τ and a set $S(\tau)$ of tours which are minor modifications of τ, the tour τ is said to be locally optimal if

$$E[L_\tau] \leqslant \min_{\tau' \in S(\tau)} E[L_{\tau'}]. \tag{35}$$

The iterative improvement algorithm works by choosing an initial tour τ_0, then testing to see if τ_0 is locally optimal. If a better tour τ_1 is found, it then replaces τ_0 and is itself tested. Since there are only a finite number (although potentially a very large number) of possible tours, this procedure must eventually converge to a locally optimal tour τ_* – which it is hoped will be a nearly optimal solution to the problem. In our experiments we have found that varying the initial guess τ_0 does not usually affect the expected length of the solution τ_* in any predictable fashion – except, of course, for the fact that $E[L_{\tau_*}]$ cannot be greater than $E[L_{\tau_0}]$. However, choosing a fairly good τ_0, such as the space-filling curve solution, is often worthwhile since it may reduce the number of steps required for convergence.

Lin [24] used an iterative improvement algorithm for the TSP based on what he called the λ-opt local neighborhood. For a given tour τ consisting of n links between nodes, the neighborhood $S_\lambda(\tau)$ consists of those tours which differ from τ by no more than λ links. For $\lambda = 2$ this is the set of tours which can be obtained by reversing a section of τ; for $\lambda = 3$ it is the set of tours obtainable by removing a section of τ and inserting it, with or without a reversal, at another place in the tour. We have implemented both the 2-opt and 3-opt TSP algorithms, since when p is greater than about 0.5 the TSP solutions provide useful starting points for our more general PTSP routines. The testing procedure for 2-optimality consists of simply running through all of the tours in $S_2(\tau)$, and thus requires $O(n^2)$ time. Our 3-opt routine is based on the one given in [24], but we have reduced the number of necessary operations from $O(n^3)$ to nearly $O(n^2)$ by eliminating inner loops whenever possible. Specifically, let us number the nodes from 0 to $n-1$ in order of their appearance in τ, and let these indices be treated modulo n. Let $d(i, j)$ be the distance from node i to node j, let $d^{\tau}_{\max}$ be the length of the longest link in the tour, and let $d_{\min}$ be the smallest distance between any two nodes, regardless of whether or not the corresponding link is part of τ. Each tour in $S_3(\tau)$ can be obtained by removing the section $(i, \ldots, j)$ from τ and re-inserting it between two other nodes k and $k+1$. For each i and j the inner loop consists of examining every possible value of k. Regardless of k, however, we know that the links leaving the tour will have lengths $d(i-1$, $d(j, j+1)$, and some other length not greater than $d^{\tau}_{\max}$. Of the links entering the tour, one will have length $d(i-1, j+1)$ and the other two cannot be shorter than $d_{\min}$. If

$$d(i-1, j+1) + 2d_{\min} - d(i-1, i) - d(j, j+1) - d^{\tau}_{\max} \geqslant 0,$$

we need not execute the inner loop at all since none of the tours involved could be shorter than τ. Our own program is somewhat more elaborate than this, in that it uses more detailed bounding information than we have described here. The details are unimportant, however, since an even more sophisticated TSP algorithm using higher values of λ was described by Lin and Kernighan [23].

Unlike in the TSP case, the expected length $E[L_{\tau}]$ in the PTSP sense depends on all $\frac{1}{2}(n^2-n)$ independent elements of the distance matrix. We cannot, therefore, speak of some links leaving and others entering the tour, rather it is just the weight given to each of the $d(i, j)$ by (3) which changes. We can still use Lin's λ-opt neighborhoods, but the computation of the changes in expected length becomes considerably more complicated. It takes $O(n^2)$ time to calculate the change in expected length from τ to an arbitrary tour in $S_2(\tau)$, so it would seem at first that testing for even 2-p-optimality would take $O(n^4)$ time. We can, however, reduce this to $O(n^2)$ if we examine the tours in the proper sequence and maintain certain auxiliary arrays of information as the computation proceeds.

Let us denote by $\Delta E_{i,j}$ the change in the expected length $E[L_{\tau}]$ caused by reversing a section $(i, \ldots, j)$ of the tour. The first phase in our test for 2-p-optimality is computing $\Delta E_{i,i+1}$ for every value of i – these n calculations require $O(n)$ time apiece, or $O(n^2)$ time in all. At the same time we accumulate two matrices of partial results which will be useful later,

$$A_{i,k} = \sum_{r=k}^{n-1} (1-p)^{r-1} d(i, i+r) \quad \text{and} \quad B_{i,k} = \sum_{r=k}^{n-1} (1-p)^{r-1} d(i-r, i), \quad 1 \leqslant k \leqslant n-1 \tag{36}$$

where as before the nodal indices are taken modulo n. Each time we encounter a negative ΔE we immediately switch the two nodes involved. Note that since

$$\Delta E_{i,i+1} = p^3 \left[A_{i,2}/(1-p) - (B_{i,1} - B_{i,n-1}) - (A_{i+1,1} - A_{i+1,n-1}) + B_{i+1,2}/(1-p) \right], \tag{37}$$

only two rows of A and B are involved in each comparison. Thus we can immediately proceed to the next pair of nodes without recomputing each entire matrix. At the end of this phase we will have reached a tour for which every $\Delta E_{i,i+1}$ is positive, and the matrices A and B will be complete and correct for that tour.

In the second phase we compute the remaining $\Delta E_{i,j}$ recursively by means of the formula

$$\begin{aligned} \Delta E_{i,j} = \Delta E_{i+1,j} + p^2 \big[& (q^{-k}-1) A_{i,k+1} + (q^k - 1)(B_{i,1} - B_{i,n-k}) \\ & + (q^k - 1)(A_{j,1} - A_{j,n-k}) + (q^{-k}-1) B_{j,k+1} \\ & + (q^{n-k}-1)(A_{i,1} - A_{i,k}) + (q^{k-n}-1) B_{i,n-k+1} \\ & + (q^{k-n}-1) A_{j,n-k+1} + (q^{n-k}-1)(B_{j,1} - B_{j,k}) \big] \end{aligned} \tag{38}$$

where $k = j - i$ and $q = 1 - p$. The first four terms inside the square brackets represent the interaction between nodes i and j and the nodes outside the reversed section; the last four terms represent the interaction between these two nodes and those inside the section. Some of the terms in this equation can be combined for more efficient computation; we have written it in this form merely to provide greater clarity. Since each $\Delta E_{i,j}$ in phase two can be found in O(1) time, this phase, and thus the entire 2-popt checking sequence, can be performed in only $O(n^2)$ time.

Once any phase-two inversion is actually performed, the computation must immediately return to phase one since the matrices A and B must be recomputed. The structure of formula (42) requires that the computation proceed in the direction of increasing k, so it is fortunate that in practice most negative values of $\Delta E_{i,j}$ are found when k is relatively small. It is easy to see why this should be the case: Suppose we already have a fairly good tour – which is not surprising since we can precondition with the

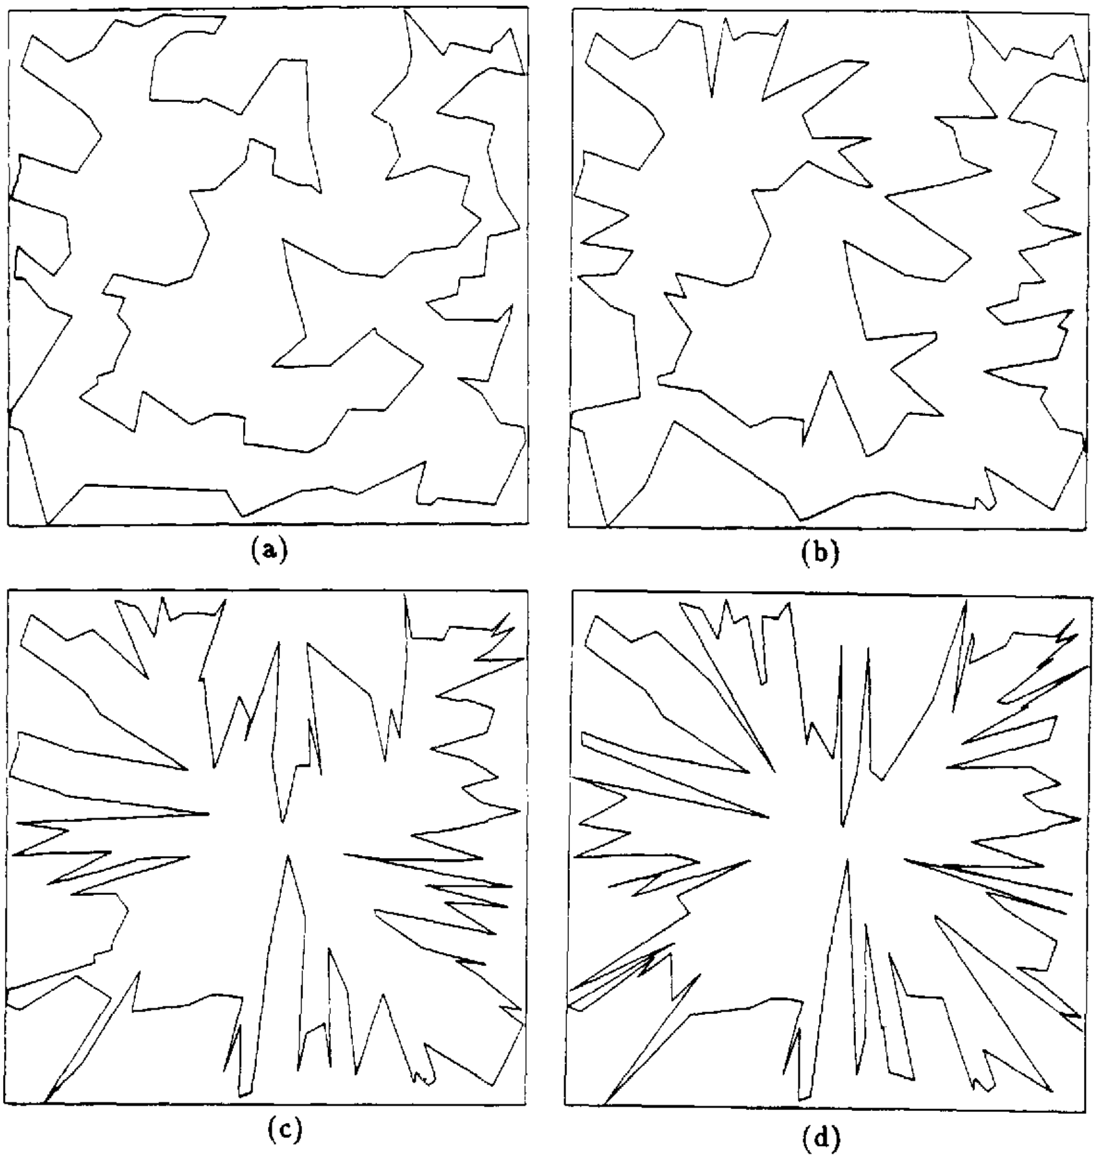

Figure 6. Near-optimal tours for a 150-node problem at four different probabilities: (a) $p = 1$, (b) $p = 0.343$, (c) $p = 0.064$, (d) $p = 0.001$. The tour shown in part (d) remains optimal as p approaches 0

space-filling curve heuristic, and since phase one has already been completed. In the TSP there is no global coupling between nodes, so inverting any section of the tour alters only two links. In the PTSP, however, global coupling becomes increasingly important as p decreases. Inverting a long section of a tour is therefore a very major perturbation, and if the tour is already fairly good, then any major perturbation is likely to make it worse instead of better. This coupling effect also tends to make phase one more efficient for low values of p, since individual nodes can 'feel' their proper places in the tour from a greater distance.

Though global coupling tends to decrease the time necessary to compute 2-p-optimal tours, it raises the question of whether the 2-opt neighborhood is a good choice when p is small. The idea behind iterative improvement is to examine minor modifications of a given state to see if a better state can be found. If most of the modifications in the 'local' neighborhood are not minor at all, then perhaps a different neighborhood should be used. For the PTSP a minor alteration is one which leaves the ordering of the majority of nodes unchanged; this is easily achieved by simply moving a single node to another point in the tour, rather than reversing an entire section. The corresponding neighborhood, which we call the 1-shift neighborhood, has roughly twice as many members as S_2, is a subset of S_3, and yields much better results than S_2 in our experiments. The algorithm follows the same lines as the 2-p-opt algorithm: All phase-one computations, including the accumulation of the matrices A and B, proceed in the same way. The phase-two computations even use a similar recursive formula,

$$\begin{aligned}\Delta E'_{i,j} = \Delta E'_{i,j-1} + p^2\big[&(q^{-k} - q^{-(k-1)})A_{i,k+1} + (q^k - q^{k-1})(B_{i,1} - B_{i,n-k}) \\ &+ (q-1)(A_{j,1} - A_{j,n-k}) + (q^{-1} - 1)B_{j,k+1} \\ &+ (q^{n-k} - q^{n-(k-1)})(A_{i,1} - A_{i,k}) \\ &+ (q^{k-n} - q^{(k-1)-n})B_{i,n-k+1} \\ &+ (1 - q^{-1})A_{j,n-k+1} + (1-q)(B_{j,1} - B_{j,k})\big]\end{aligned} \tag{39}$$

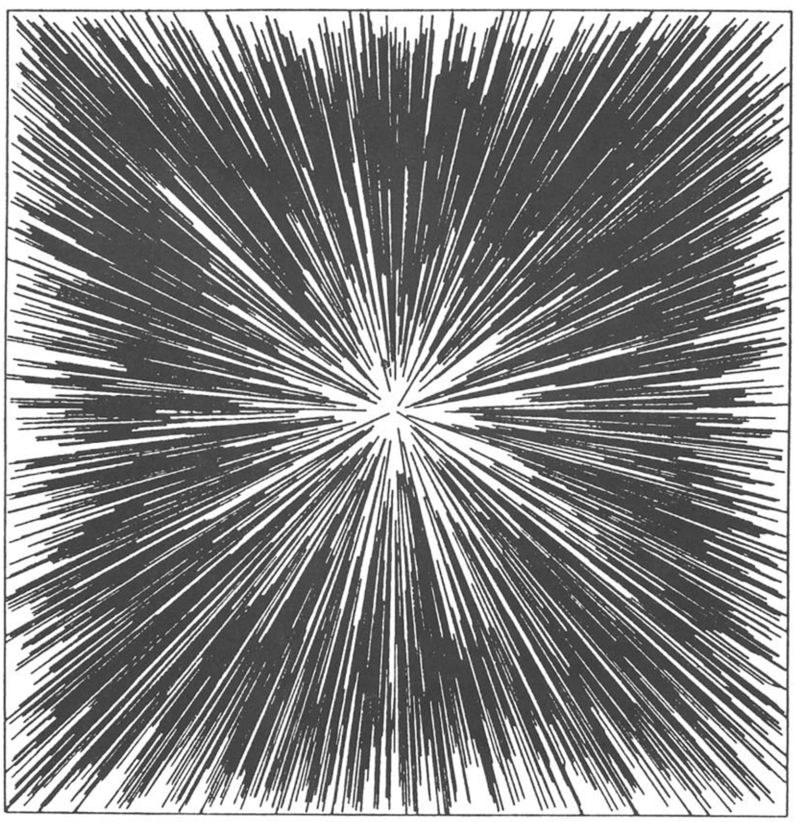

Figure 7. Angular sorting applied to the 5000-node problem of Figure 5

where $\Delta E'_{i,j}$ denotes the change in $E[L_\tau]$ which results when node i is moved to position j, with the intervening nodes being shifted backwards one space accordingly. As before, this equation can be simplified for more efficient computation – we have chosen this form because it makes the significance of each term more readily apparent.

Despite the similarities, there is one major difference between the implementations of our 2-popt and 1-shift algorithms. With 2-popt significant improvements were most likely to be found at small k, so whenever a negative ΔE appeared the change was immediately accepted, sending the algorithm back into phase one. With the 1-shift algorithm, however, the best improvement may appear at any k, so we have found it advantageous to run through the entire phase-two checking sequence before deciding which change to accept. Since the program chooses the *best*, rather than the *first*, tour found, it typically requires fewer cycles through phase one to reach a locally optimal solution.

Figure 6 shows the best near-optimal tours found for the same set of 150 nodes but different values of p. To obtain these solutions we first ran the space-filling curve heuristic, then the 3-opt TSP algorithm, and finally the 1-shift PTSP routine for each p.

The appearance of the solutions for small p suggests an alternative heuristic which is even simpler to implement than the space-filling curve algorithm: Simply sort the nodes by their angular positions with respect to some central point. Figure 7 shows the effect of this heuristic, which we call angular sorting, on the 5000 node problem of Figure 5. (The 'center of mass' of the nodes was used as the central point – it is quite possible that some other choice would be better.) As p approaches 0, this solution is actually better than the one produced by the space-filling curve heuristic. The range where this is true is rather narrow, though, so we doubt that this algorithm will be very useful in practice. It is interesting that even

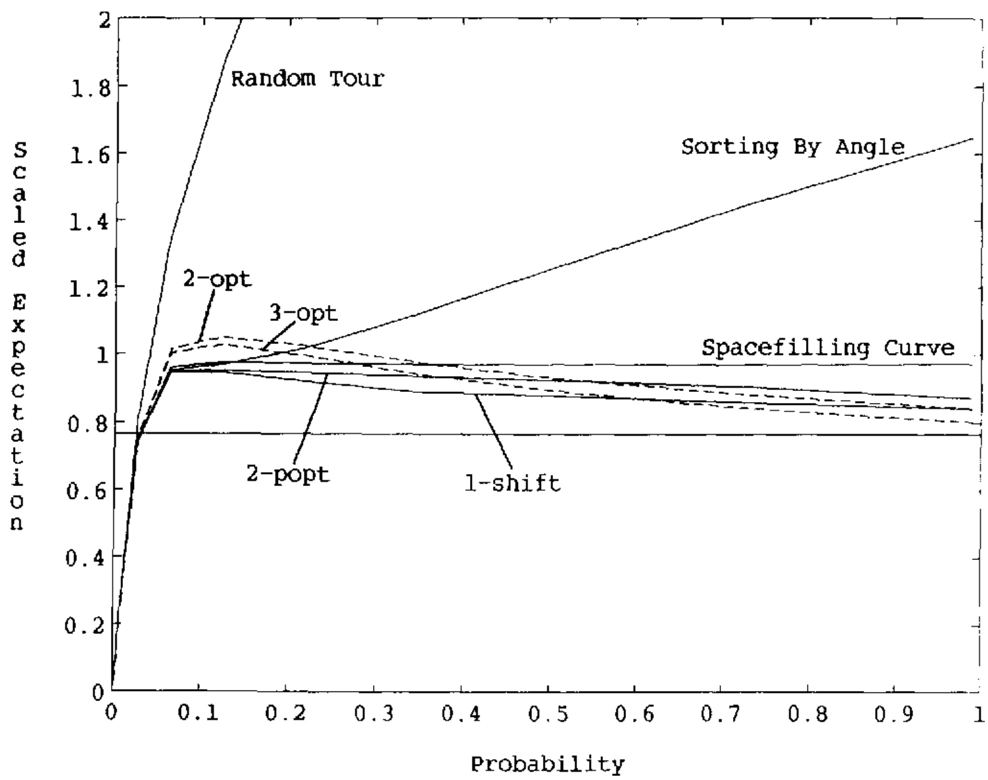

Figure 8. A summary of results for several PTSP heuristics on 100-node problems scaled by $\sqrt{np}$. Solutions obtained via the 2-opt and 3-opt TSP algorithms (dashed lines) are shown for comparison. The horizontal line shows the value of $\beta_{TSP} = 0.765$. The heuristics are 1) random tour, 2) angular sorting, 3) space-filling curve, 4) 2-popt and 5) 1-shift

as p approaches 0, the optimal tour is only approximately, not exactly, starlike. (A starlike tour would have a 'center' such that sorting by angle about this point would yield the optimal tour.) If this behavior were better understood, it might be possible to construct the optimal solution for $p \to 0$ in polynomial time.

A summary of the behavior of each of the heuristics we have used is shown in Figure 8. The space-filling curve solutions were used as starting positions for the 2-popt and 1-shift algorithms; this greatly reduces the amount of work required and does not affect the results for small p. When p is large, however, the effect on the 2-popt results is somewhat detrimental. The 2-opt and 3-opt TSP algorithms were started from random positions-note that near $p = 1$, 2-opt gives significantly better results than 2-popt because of the different starting positions. The more powerful 3-opt and 1-shift algorithms do not seem to suffer from this effect: 3-opt gives excellent results for large p regardless of the starting position, and for small p the 1-shift solutions are usually optimal. (This conclusion is based on the fact that the algorithm always converges to the same tour regardless of the starting position.) The best general approach seems to be to first use the space-filling curve algorithm, followed by 3-opt if p is fairly large, and then finish by applying 1-shift. Adding 2-popt to this procedure helps only rarely, and does not seem to be worth the trouble. The turnover point below which 3-opt ceases to be helpful is uncertain and probably depends strongly on the details of the problem. For problems with more than a few hundred nodes both the running time and the memory required for the distance and the auxiliary matrices A and B begin to become excessive. At this point we were forced to switch to heuristics like the space-filling curve algorithm which do not require $O(n^2)$ memory.

In the calculations for Figure 8 (also 9 and 10), data from 10 separate 100-node problems were averaged in order to minimize the effects of statistical fluctuations. The nodes for each problem were

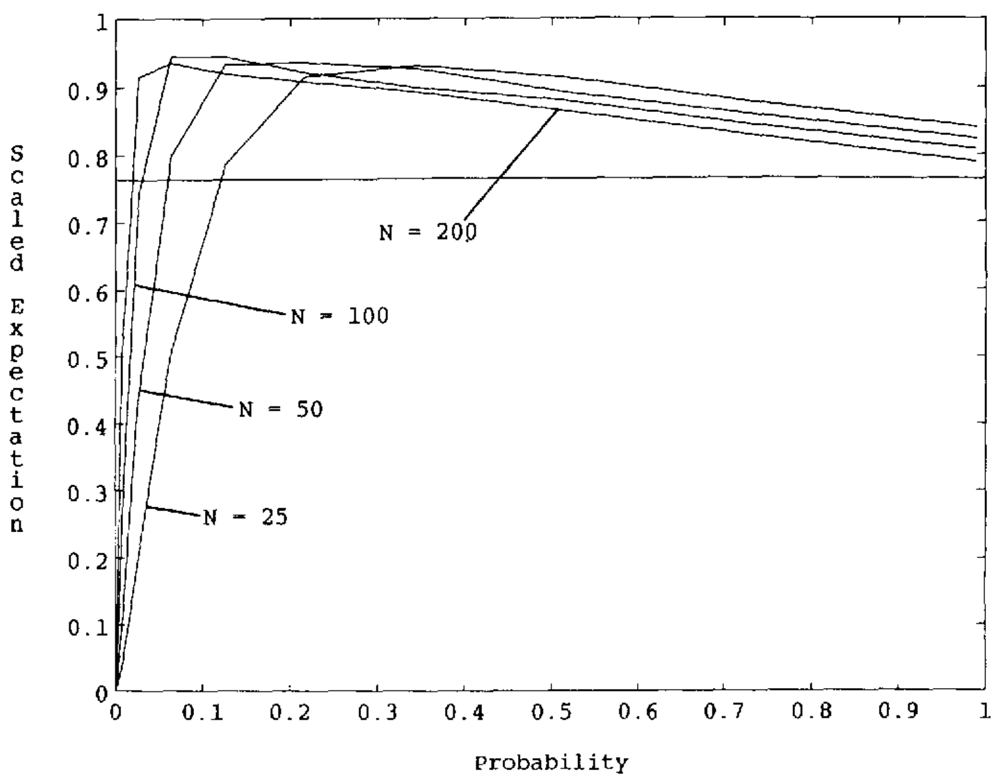

Figure 9. Near-optimal tours for various values of n. These tours were obtained by applying first the space-filling curve heuristic, then the 3-opt TSP algorithm, then the 1-shift PTSP algorithm. The values of n shown are $n = 25, 50, 100, 200$

chosen from a uniform distribution in the unit square, and we have scaled the expected lengths $E[L_\tau]$ by $\sqrt{np}$. Our heuristics seem to suggest that data from optimal tours would follow a horizontal line on the plot. Figure 9 shows the behavior of near-optimal solutions for different values of n – the curves do seem to be converging towards the horizontal line $E[L_\tau]/\sqrt{np} = \beta_{TSP}$, but the convergence is non-uniform. A look at (3) explains the behavior at low p: Once p is small enough for $(1-p)^n$ to be a significant fraction of 1, it is unlikely that event two nodes will be present in any given instance of the problem. The expected length for *any* tour in this range is then $O(n^2p^2)$, rather than $O(\sqrt{np})$. As n becomes large, however, the range in p for which this behavior is important becomes negligible.

Even though they do not match the performance of the PTSP heuristics, the 2-opt and 3-opt results in Figure 8 seem surprisingly good over the whole range of probabilities. If these results are indicative of the behavior of good TSP tours on PTSP problems, then Jaillet's bound (4) would seem to be a false alarm, not applicable except in contrived examples. There would then be little point in working with the PTSP at all, since TSP tours would give very similar results. If we experiment with larger problems, though, it becomes apparent that the expected poor behavior is simply masked by the transition into the low-p regime described above. Figure 10 shows the performance of 2-opt TSP tours for $n = 100$ and $n = 500$, and we can see that the 2-opt tours, and thus presumably the optimal TSP tours, do in fact behave increasingly badly as n becomes large. Since practical PTSP problems would typically involve large numbers of nodes, the PTSP approach should therefore prove useful in dealing more efficiently with such systems.

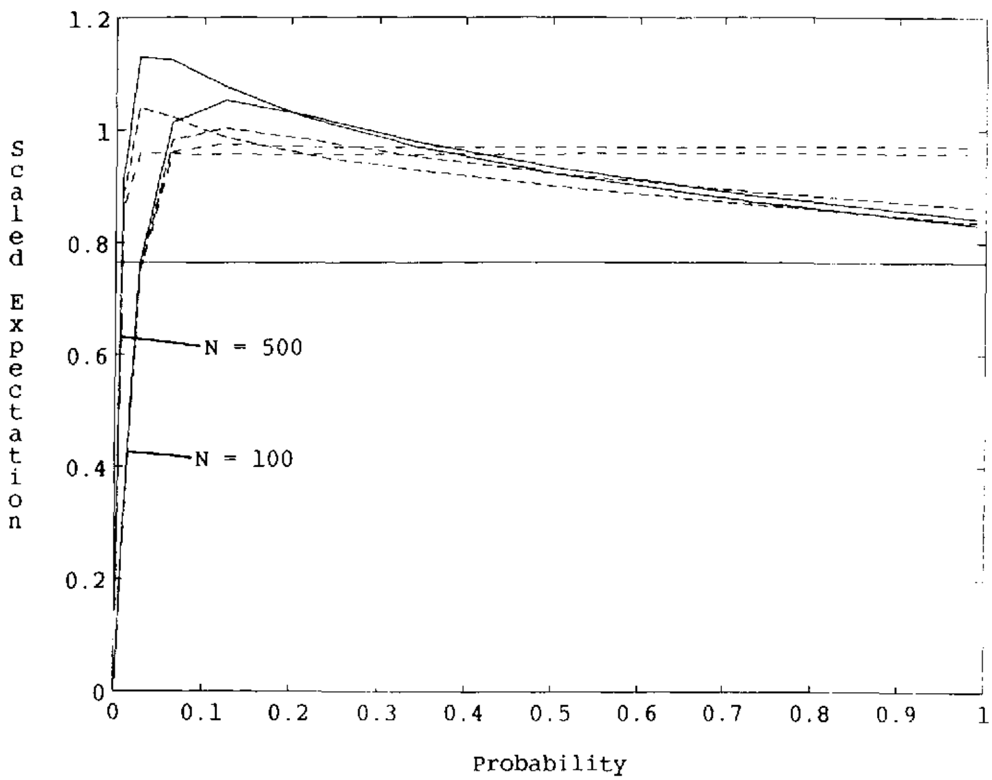

Figure 10. Performance of good TSP tours on the PTSP on problems with 100 and 500 nodes. The solid lines show 2-opt tours using random starting positions, the dot–dash lines show space-filling curve tours, and the dashed lines show 2-opt tours using space-filling curve tours as starting points

8. Concluding remarks

The idea of inserting probabilistic elements in combinatorial optimization problems and of having a *a priori* solution to the original instance of a combinatorial problem and also a real-time strategy to update the solution to the particular instance can be applied to many combinatorial optimization problems. There are several motivations for these considerations, among which two are of particular importance. The first is the desire to formulate and analyze models which are more appropriate for real-world problems, in which randomness is not only present but is a major concern. There is a plethora of important and interesting applications of probabilistic combinatorial optimization problems, especially in the context of strategic planning for collection and distribution services, communication and transportation systems, job scheduling, organizational structures, etc. For such applications, the probabilistic nature of the models makes them particularly attractive as mathematical abstractions of real-world systems.

The second motivation is interest in investigating the robustness (with respect to optimality) of optimal solutions to deterministic problems, when the instances for which these problems have been solved are modified. In our case, we confine the investigation to problems on graphs and the perturbation of a problem's instance is simulated by the presence of subsets of the given set of nodes. These ideas are further exploited in [5], where probabilistic variations of the vehicle routing problem, the minimum spanning tree problem and facility location problems are investigated.

Acknowledgments

The first author wishes to thank his Thesis advisor Professor Amedeo Odoni for several useful comments and interesting discussions and Professor Patrick Jaillet for his constructive remarks. The research of the first author was partially supported from the National Science Foundation under grant ECS-8717970. The second author was supported by US Air Force grant AFOSR-87-0102.

References

[1] Bartholdi, J.J., and Platzman, L.K., "An $O(N \log N)$ planar traveling salesman heuristic based on spacefilling curves", *Operations Research Letters* 1 (1982) 121–125.
[2] Bartholdi, J.J., Platzman, L.K., Collins, R. Lee, and Warden, W.H., "A minimal technology routing system for meals on wheels" *Interfaces* 13 (1983) 1–8.
[3] Beardwood, J., Halton, J., and Hammersley, J., "The shortest path through many points", *Proceedings of the Cambridge Philosophical Society* 55 (1959) 299–327.
[4] Berman, O., and Simchi-Levi, D., "Finding the optimal a priori tour and location of a traveling salesman with non-homogeneous customers", *Transportation Science* 22 (1988) 148–154.
[5] Bertsimas, D., "Probabilistic combinatorial optimization problems", Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA, 1988.
[6] Bertsimas, D., "The probabilistic minimum spanning tree", *Networks* 20 (1990) 245–275.
[7] Bertsimas, D., and Grigni, M., "On the spacefilling curve heuristic for the Euclidean Traveling Salesman Problem", *Operations Research Letters* 8 (1989) 241–244.
[8] Christofides, N., "Worst-case analysis of a new heuristic for the Traveling Salesman Problem", Carnegie-Mellon University, Pittsburgh, 1976.
[9] Dror, M., Laporte, G., and Trudeau, P., "Vehicle routing with stochastic demands: Properties and solution framework", *Transportation Science* (1989) 166–176.
[10] Garey, M., and Johnson, D., *Computers and Intractability: A Guide to the Theory of NP-Completeness*, Freeman, San Francisco, CA, 1979.
[11] Gomory, R.E., "Solving Linear Programming problems in integers", *Proceedings of Symposia in Applied Mathematics* 10 (1963) 211–215.
[12] Grötchel, M., and Padberg, M.W., "Polyhedral theory", in: E.L. Lawler, J.K. Lenstra, A.H.G. Rinnooy Kan and D.B. Shmoys (eds.), *Traveling Salesman Problem*, Wiley, Chichester, 1985, 251–306.
[13] Held, M., and Karp, R., "A dynamic programming approach to sequencing problems", *SIAM Review* 10 (1962) 196–210.

[14] Held, M., and Karp, R., "The Traveling Salesman Problem and minimum spanning trees", *Operations Research* 18 (1970) 1138–1162.
[15] Held, M., and Karp, R., "The Traveling Salesman Problem and minimum spanning trees: Part II", *Mathematical Programming* 1 (1971) 6–25.
[16] Jaillet, P., "Probabilistic Traveling Salesman Problems", (Ph.D. Thesis) Technical Report No. 185, Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA, 1985.
[17] Jaillet, P., and Odoni, A., "The probabilistic vehicle routing problem", in: B.L. Golden and A.A. Assad (eds.), Vehicle Routing: Methods and Studies, North-Holland, Amsterdam, 1988, 293–318.
[18] Jezequel, A., "Probabilistic vehicle routing problems", S.M. Thesis, Department of Civil Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts, 1986.
[19] Karp, R., "Probabilistic analysis of partitioning algorithms for the traveling salesman in the plane", *Mathematics of Operations Research* 2 (1977) 209–224.
[20] Karp, R., and Steele, J.M., "Probabilistic analysis of heuristic", in: E.L. Lawler, J.K. Lenstra, A.H.G. Rinnooy Kan and D.B. Shmoys (eds.), *Traveling Salesman Problem*, Wiley, Chichester, 1985.
[21] Larson, R., and Odoni, A., *Urban Operations Research*, Prentice-Hall, Englewood Cliffs, NJ, 1981.
[22] Lawler, E.L., Lenstra, J.K., Rinnooy Kan, A., and Shmoys, D.B. (eds.), *The Traveling Salesman Problem: A Guided Tour of Combinatorial Optimization*, Wiley, Chichester, 1985.
[23] Lin, S., and Kernighan, B., "An effective heuristic algorithm for the Traveling Salesman Problem", *Operations Research* 21 (1973) 498–516.
[24] Lin, S., "Computer solution of the Traveling Salesman Problem", *Bell System Technical Journal* 44 (1965) 2245–2269.
[25] Odoni, A., private communication, 1988.
[26] Platzman, L.K., and Bartholdi, J.J., "Spacefilling curves and the planar Traveling Salesman Problem", Report No. PDRC 83-02, School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA, 1983.
[27] Raghavan, P., "Probabilistic construction for deterministic algorithms approximating packing integer programs", *Journal of Computer and System Sciences* (1988) 1–14.
[28] Rosenkrantz, D., Stearns, R., and Lewis, P., "An analysis of several heuristics for the Traveling Salesman Problem", *SIAM Journal on Computing* 6 (1977) 563–581.
[29] Steele, J.M., "Subadditive Euclidean functionals and nonlinear growth in geometric probability", *The Annals of Probability* 9 (1981) 365–376.