
A Priori Optimization

Author(s): Dimitris J. Bertsimas, Patrick Jaillet, Amedeo R. Odoni

Source: *Operations Research*, Vol. 38, No. 6 (Nov. - Dec., 1990), pp. 1019-1033

Published by: **INFORMS**

Stable URL: <http://www.jstor.org/stable/170968>

Accessed: 06/06/2011 09:04

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=informs>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



INFORMS is collaborating with JSTOR to digitize, preserve and extend access to *Operations Research*.

A PRIORI OPTIMIZATION

DIMITRIS J. BERTSIMAS

Massachusetts Institute of Technology, Cambridge, Massachusetts

PATRICK JAILLET

Ecole Nationale des Ponts et Chaussées, Noisy Le Grand, France

AMEDEO R. ODONI

Massachusetts Institute of Technology, Cambridge, Massachusetts

(Received November 1988; revision received August 1989; accepted September 1989)

Consider a complete graph $G = (V, E)$ in which each node is present with probability p_i . We are interested in solving combinatorial optimization problems on subsets of nodes which are present with a certain probability. We introduce the idea of a priori optimization as a strategy competitive to the strategy of reoptimization, under which the combinatorial optimization problem is solved optimally for every instance. We consider four problems: the traveling salesman problem (TSP), the minimum spanning tree, vehicle routing, and traveling salesman facility location. We discuss the applicability of a priori optimization strategies in several areas and show that if the nodes are randomly distributed in the plane the a priori and reoptimization strategies are very close in terms of performance. We characterize the complexity of a priori optimization and address the question of approximating the optimal a priori solutions with polynomial time heuristics with provable worst-case guarantees. Finally, we use the TSP as an example to find practical solutions based on ideas of local optimality.

This paper is concerned with a specific family of combinatorial optimization problems whose common characteristic is the explicit inclusion of probabilistic elements in the problem definitions, as we will explain. For this reason, we shall refer to them as probabilistic combinatorial optimization problems (PCOPs).

There are several motivations for investigating the effect of including probabilistic elements in combinatorial optimization problems. Among them, two are of particular importance. The first is the desire to define and analyze models which are more appropriate for those real-world problems in which randomness is not only present, but a major concern. There is a plethora of important and interesting applications of PCOPs, 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 an 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 or absence of subsets of the graph's set of nodes.

We next discuss the central theme of this paper, namely the idea of *a priori optimization*. In many applications, one finds that, after solving a given instance of a combinatorial optimization problem, it becomes necessary to repeatedly solve many other instances of the same problem. These other instances are usually variations of the instance solved originally. Yet, they may be sufficiently different from that original instance to necessitate a reconsideration of the entire problem on the part of the analyst.

The most obvious approach to deal with such cases is to attempt to optimally solve (or near-optimally with a good heuristic) every potential instance of the original problem. Throughout the paper, we call this approach the *reoptimization strategy* and denote it with the Greek letter Σ . This approach, however, suffers from several disadvantages. For example, if the combinatorial optimization problem considered is

Subject classifications: Networks/graphs: stochastic heuristics, traveling salesman. Probability: stochastic model applications. Transportation: vehicle routing.

NP-hard, one might have to solve exponentially many instances of a hard problem. Moreover, in many applications it is necessary to find a solution to each new instance quickly, but one might not have the required computing or other resources for doing so.

We propose to investigate a different strategy. Rather than reoptimizing every potential instance, we wish to find an a priori solution to the original problem and then update in a simple way this a priori solution to answer each particular instance or variation. Clearly, the natural questions to ask are: What is the measure of effectiveness of such an a priori solution? And, how does one update the a priori solution for each problem instance?

The above discussion is general, in the sense that it applies to any combinatorial optimization problem. In order to address these questions concretely, we restrict our attention to a class of network problems. Consider then a complete graph $G = (V, E)$ on n nodes on which an optimization problem is defined, for example, the traveling salesman problem. If every possible subset of the node set V may or may not be present on any given instance of the optimization problem, for example, on any given day, the traveling salesman may have to visit only a subset S of the nodes in V , then there are 2^n possible instances of the problem—all the possible subsets of V . Suppose that instance S has a probability $p(S)$ of occurring. Given a method U for updating an a priori solution f to the *full-scale* optimization problem on the original graph G , U will then produce for problem instance S , a feasible solution $t_f(S)$ with value (*cost*) $L_f(S)$. (In the case of the TSP, $t_f(S)$ would be a tour through the subset S of nodes and $L_f(S)$ the length of that tour.) Then, given that we have already selected the updating method U , the natural choice for the a priori solution f is to select f to minimize the expected cost

$$E[L_f] = \sum_{S \subseteq V} p(S)L_f(S) \quad (1)$$

with the summation over all subsets of V . In other words, we would like to minimize the *weighted average* over all problem instances of the values $L_f(S)$ obtained by applying the updating method U to the a priori solution f .

This choice of a measure of effectiveness for the a priori solution f that we seek, namely the expected cost (1), gives a reasonable answer to our first question. But what properties should the updating method U have? The most desirable property of U would be for $L_f(S)$ to be *close* to the value of the optimal solution $L_{OPT}(S)$ for every instance S . A less restrictive and more global property is to require $E[L_f]$ to be close to

the expected cost $E[\Sigma]$, over all problem instances of the reoptimization strategy

$$E[\Sigma] = \sum_{S \subseteq V} p(S)L_{OPT}(S). \quad (2)$$

In addition, U must be able to update efficiently the solution from one problem instance to the next.

In the following definitions of the updating methods U , the choices of U may initially seem arbitrary. But these choices will turn out to be natural ones. First, for every choice of U that we are proposing, the updating of the solution to a particular instance S can be done very easily. Moreover, these updating methods are well suited for applications. And finally, we prove in Section 2 that our a priori optimization strategies coupled with our particular choices of U are asymptotically very close (we conjecture equivalent) in terms of performance to the reoptimization strategies under reasonable probabilistic assumptions.

After this general discussion of the rationale behind the definitions, which follow, we informally describe the problems that we consider.

The Probabilistic Traveling Salesman Problem

The probabilistic traveling salesman problem (PTSP) is probably the most fundamental stochastic routing problem that can be defined. It is essentially a traveling salesman problem (TSP), in which the number of points to be visited in each problem is a random variable.

Consider a problem of routing through a set of n known points. On any given instance of the problem only a subset S that consists of $|S| = k$ out of n points ($0 \leq k \leq n$) must be visited. Suppose that the probability that instance S occurs is $p(S)$. As mentioned before, ideally we might like to reoptimize the tour for every instance, but, in many cases, we may not have the resources to do so, or even if we had them, reoptimization might turn out to be too time consuming. 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 (see Figure 1 for an illustration). 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 updating method U for the PTSP is therefore to visit the points on every problem instance in the same order as in the a priori tour, i.e., we simply skip those points that are not present in that problem instance.

The expectation is computed over all possible instances of the problem, i.e., over all subsets of the




Figure 1. The PTSP methodology.

vertex set $V = \{1, 2, \dots, n\}$. That is, given an a priori tour τ , if the problem instance $S (\subseteq V)$ will occur with probability $p(S)$ and will require covering a total distance $L_r(S)$ to visit the subset S of customers, that problem instance will receive a weight of $p(S)L_r(S)$ in the computation of the expected length. If we denote the length of the tour τ by L_r (a random variable), then our problem is to find an a priori tour τ_p through all n potential customers, which minimizes the quantity

$$E[L_r] = \sum_{S \subseteq V} p(S)L_r(S) \tag{3}$$

with the summation over all subsets of V .

The Probabilistic Minimum Spanning Tree Problem

The probabilistic minimum spanning tree (PMST) problem is a natural extension of the classical minimum spanning tree problem. Given a set of n nodes on a network, a subset S of the n nodes is present on any particular instance of the problem with probability $p(S)$. We wish to find a priori a spanning tree through the n nodes which is used as follows: On any given instance of the problem, the a priori tree is retraced deleting only the nodes that are not present, provided the deletion of those nodes does not disconnect the tree. In this way, there would be nodes which will not be present but are still included in the tree. Thus, the updating method U is to include all nodes in the instance S and also those nodes in $V - S$ which are necessary to prevent the resulting tree from becoming disconnected. An example of the PMST can be found in Figure 2. Note that the problem has some Steinerish properties. This can be illustrated in the figure, where node 2 is kept on the tree to preserve connectedness. The problem of finding an a priori spanning tree of minimum expected length over all possible problem instances is the PMST problem.

The Probabilistic Vehicle Routing Problem

Consider a standard VRP but with demands which are probabilistic in nature rather than deterministic. The problem is then to determine a fixed set of routes of minimal expected total length, which corresponds to the expected total length of the fixed set of routes plus the expected value of extra travel distance that might be required. The extra distance will be due to the possibility that demand on one or more routes may occasionally exceed the capacity of a vehicle and force it to go back to the depot before continuing on its route.

The following two solution-updating methods can be defined. Under method a the vehicle visits all the points in the same fixed order as under the a priori tour, but serves only customers that require service during that particular problem instance. The total expected distance traveled corresponds to the fixed length of the a priori tour plus the expected value of the additional distance that must be covered whenever the demand on the route exceeds vehicle capacity. Method b is defined similar to a with the sole difference that customers with no demand on a particular instance of the vehicle tour are simply skipped. An example of the PVRP under both methods can be seen in Figure 3.

The Traveling Salesman Facility Location Problem

We are given a set of n nodes (customer locations) on a network. Each day a subset S of customers makes a request for service with probability $p(S)$. By a specific time of each day, a service unit receives the list of calls for that day and starts a traveling salesman tour using the underlying network that visits all the customer locations in the list. The objective is to find an optimal location i for the service unit, so that the expected distance traveled

$$E[\Sigma_{\text{TSFLP}}(i)] = \sum_{S \subseteq V} p(S)L_{\text{OPT}}(S \cup i)$$






Figure 2. The PMST methodology.



(i) An a priori route through 6 customers (each with a demand of zero or one unit) by a vehicle of capacity 2



(ii) The two strategies when only the second, third and fifth customers have a non-zero demand.

Figure 3. The PVRP methodology.

is minimized. This problem is called the traveling salesman facility location problem (TSFLP).

The difficulty of having to compute the optimal tour for every instance can be overcome by using an a priori tour τ_p and then follow the PTSP approach described before, i.e., skip customer locations with no demand. The problem is then to find a node i_p and an a priori tour τ_p to minimize the expected distance traveled using the PTSP approach, i.e., to minimize

$$h(i, \tau) \triangleq \sum_{S \subseteq V} p(S) L_\tau(S \cup i). \tag{4}$$

The problem of finding simultaneously an optimal location i_p and an optimum a priori tour τ_p is called the probabilistic traveling salesman facility location problem (PTSFLP).

Throughout the paper the emphasis is on concepts and results rather than detailed derivations. To keep the length of the presentation within reasonable limits, all but the most important theorem proofs are sketchily outlined, with appropriate references given for interested readers. In Section 1, we review briefly the related research and outline potential areas of application for the idea of a priori optimization. In Section 2, we prove that the a priori strategies we are proposing are asymptotically close to the reoptimization strategies for all the problems we have defined. This gives an indication of the importance of the a priori optimization idea. In Section 3, we address the complexity of finding the best a priori solutions for all PCOPs we have defined. In Section 4, we examine the question of finding good approximations from a

theoretical point of view, and in Section 5 we use the PTSP as an example to illustrate how to find good practical approximations. The final section contains some concluding remarks.

1. LITERATURE REVIEW AND APPLICATIONS

During the last decade, combinatorial optimization undoubtedly has been one of the fastest growing and most exciting areas in mathematical programming. The related scientific literature has been expanding at a rapid pace. Examples of particular relevance to this paper are the three excellent review volumes: on the traveling salesman problem (Lawler et al. 1985), on routing and scheduling (Bodin et al. 1983), and on vehicle routing (Golden and Assad 1988); each offers several hundreds of references.

Research at the interface between probability theory and combinatorial optimization spans a period of over 30 years and in recent years has been at the center of much activity. The dominant trends of this interplay that are relevant to this paper can be summarized as follows.

Probabilistic analysis of combinatorial optimization problems in the Euclidean plane. Research in this area was initiated by the pioneering paper of Beardwood, Halton and Hammersley (1959). After a period of more than 15 years and motivated by the significant advances in theoretical computer science, Karp (1977) used their main result to propose a partitioning heuristic, which constitutes an ϵ -approximation algorithm for the TSP in the Euclidean plane.

In the last decade, the asymptotic properties of many combinatorial optimization problems in the Euclidean plane have been investigated. The most general analysis in this direction is due to Steele (1981), who developed the theory of Subadditive Euclidean functionals to obtain very sharp limit theorems for a broad class of combinatorial optimization problems.

Probabilistic analysis on problems with random lengths. In the last decade there have been numerous papers dealing with the behavior of combinatorial optimization problems when the costs involved are taken from a probability distribution. Interest in this area intensified after the pioneering paper of Karp (1979) on the TSP and the attempts to explain probabilistically the success of the simplex method for linear programming. Of particular relevance to this paper are the papers on the minimum spanning tree problem by Frieze (1985) and Steele (1987).

Probabilistic combinatorial optimization problems. In contrast to their deterministic counterparts, the professional literature on PCOPs is sparse. Jaillet (1985, 1988) introduced the PTSP, examined some of its combinatorial properties and proved asymptotic theorems in the plane. A summary of these results as well as a discussion on the applications of the PTSP and the PVRP are contained in Jaillet and Odoni (1988). Bertsimas (1988) introduced the framework of a priori optimization and studied the problems considered in this paper.

Except for an isolated result in the 1970s (Tillman 1969), VRPs with stochastic elements in their definitions have received attention only recently. Stewart and Golden (1983), Dror and Trudeau (1986), Laporte and Louveau (1987) and Laporte, Louveau and Mercure (1987) use techniques from stochastic programming to solve optimally small problems and find bounds for them. The definitions of these problems are different from the ones we consider in this paper.

The traveling salesman facility location problem has been considered by Burness and White (1976), where heuristic approaches are proposed. In a series of papers, Berman and Simchi-Levi (1986, 1988a, b) and Simchi-Levi and Berman (1988) solved the problem on a tree network and proposed a heuristic of relative worst error $\frac{1}{2}$ for the general network case as well as for the Euclidean and the rectilinear metric. Bertsimas (1989) improved on their results by proving that the relative worst error is $\frac{1}{2}(1 - p)$, where p is the coverage probability.

To our knowledge, the PMST problem has never been examined before in the literature despite its intrinsic interest and applicability.

A final remark has to do with the relationship between network reliability theory and the class of PCOPs we are considering. In network reliability theory the nodes are usually assumed to be reliable and the types of questions addressed are about the existence of paths among pairs of nodes. In the class of PCOPs, the types of questions that we address as well as the motivation for their definition are different.

As noted earlier, PCOPs could prove useful in many application contexts in which the explicit consideration of randomness is essential. For instance, the PTSP arises in practice whenever a company, on any given day, is faced with the problem of collections (deliveries) from (to) a random subset of some known global set of customers in an area and does not wish to, or, simply, cannot redesign the tours from scratch every day. Examples in this category include a "hot meals"

delivery system described by Bartholdi et al. (1983), routing of forklifts in a cargo terminal or in a warehouse and, interestingly, the daily delivery of mail to homes and businesses by postal carriers everywhere. In fact, it was this last application that led to the initial formulation of the PTSP by the third author. Jaillet and Odoni (1988) describe in considerable detail an application in a strategic planning context in which a package distribution company has decided to begin service in a particular area. After carrying out a market survey and identifying a set of potential major customers who, during any single time period, have a significant probability of requiring a visit, the company wishes to estimate the resources necessary to serve these customers. The PTSP then provides a model for computing approximately the expected amount of travel that will be required per time period and, by implication, the number of vehicles, drivers, etc.

In a nonrouting context, PTSP models also can be of interest in many situations in which an ordering of entities of any type has to be found and that sequence has to be preserved even when some of the entities may be absent. One such example can be given from the area of job-shop scheduling: Consider the problem of loading n jobs on a machine at which a changeover cost is incurred whenever a new job is loaded. With any given ordering of the n jobs on the machine, we can then associate a total changeover cost. Any given ordering of the n jobs may also impose specific long-term requirements on the job shop, such as a set of tasks to be performed before and after the processing of the jobs on the machine. These requirements may often be difficult to modify on a daily basis so that, if on a given day some jobs need not be processed, the relative ordering previously specified for the remaining jobs is nonetheless left unmodified. The PTSP is again relevant to analyze such situations.

PVRPs are *constrained* cases of PTSPs and thus arise in the same collection and distribution contexts as PTSPs, whenever the vehicle capacity Q becomes a practically significant issue. The capacity Q may be expressed in terms of a maximum allowable vehicle load, maximum number of stops, maximum distance per tour or some other physical or statutory limitation. For instance, in the case of the delivery of cash by a bank to a set of automatic teller machines spatially distributed throughout a city, Q might be the upper bound on the amount of money that a vehicle might carry for safety reasons. The uncertainty in this problem is due to the fact that each machine may or may not require a visit during any given time period,

depending on the amount of money it dispenses. Similar applications of the PVRP can be found in most problems that combine inventory and routing considerations.

Probabilistic traveling salesman location problems arise similarly in the complex but very common context in which facility location, routing and, possibly, inventory-related decisions must be made simultaneously. Note the difference between these problems and the classical *median* (or *minisum*) and *center* (or *minimax*) problems in facility location theory. In the case of (P)TSFLPs, once a facility is located, demands are visited through tours; therefore, the facility location problem must be central relative to the *ensemble* of the demand points, as ordered by the, yet unknown, tour through all of them. By contrast, in the classical problems the facility (or facilities) must be located by considering distances to *individual* demand points, thus making the problem more tractable.

Examples of applications of the PMST are less obvious, but important nonetheless. The problem arises in many cases where a set of points must be connected through an underlying tree structure, with only portions of that structure activated with each problem instance. For example, in a communications context the active demand points would be centers that seek to communicate with each other on each problem instance and the activated portion of the underlying communications network would be the minimum tree necessary to establish communications between every possible pair of *active* demand points. Similar examples can be drawn from transportation and circuit design.

A more unusual application of the PMST problem is in the area of organizational structures. For instance, a rather intriguing paradigm might be the following: Suppose the n points that we wish to interconnect represent our agents or spies in a foreign country. They will undertake in the future a series of missions, each mission involving a different subset of agents. A mission, in our context, is an instance of the problem. We are looking for an a priori organizational structure in which, for obvious reasons, each agent will know only the people immediately above or below him/her in the structure; this implies a spanning tree-like structure. The probability p_i associated with point i is the a priori probability that agent i will have to participate in any random mission undertaken by the network. For any given mission, only that part of the organization that is necessary to interconnect all the agents participating in that particular mission is activated. The distance between points i and j is interpreted as

the cost or risk of exposure incurred when agents i and j must communicate or work with each other. Given p_i for $i = 1, 2, \dots, n$ and the distance matrix for all possible pairs (i, j) , the PMST gives the organizational structure which, in the expected value sense, minimizes the risk of exposure of the network on a random mission.

2. ASYMPTOTIC COMPARISON OF REOPTIMIZATION AND A PRIORI OPTIMIZATION

In this section, we characterize the asymptotic behavior of the reoptimization and the a priori strategies for the four problems we defined in the Introduction, if the locations of the points are uniformly and independently distributed in the Euclidean plane. This comparison is important in order to assess the promise and potential usefulness of the a priori strategies.

Let $X^{(n)} = (X_1, \dots, X_n)$ be n points uniformly and independently distributed in the unit square. Let L_{TSP}^n , L_{MST}^n , L_{STEINER}^n and L_{VRP}^n be the length of the TSP, MST, STEINER tree and the VRP (the depot being the point $(0, 0)$ and the vehicle capacity being Q) defined on $X^{(n)}$, respectively.

Let $E[\Sigma_{\text{TSP}}^n]$, $E[\Sigma_{\text{MST}}^n]$, $E[\Sigma_{\text{STEINER}}^n]$, $E[\Sigma_{\text{VRP}}^n]$, $E[\Sigma_{\text{TSFLP}}^n(i)]$ be the expectation of the TSP, MST, STEINER tree, VRP and TSFLP solutions obtained under the reoptimization strategies defined on $X^{(n)}$. Note that for the case of the TSFLP the expectation depends on the node i selected as the server's location node.

Let $E[L_{\text{PTSP}}^n]$, $E[L_{\text{PMST}}^n]$, $E[L_{\text{PVRP}a}^n]$, $E[L_{\text{PVRP}b}^n]$, $E[L_{\text{PTSFLP}}^n(i)]$ be the expectation of the a priori strategies, i.e., the expected length of the optimal a priori solution to the PTSP, PMST and PVRP under updating methods a and b , and PTSFLP defined on $X^{(n)}$.

It is well known that we can sharply characterize the solutions to the deterministic problems.

Theorem 1. *With probability 1 there are constants (Steele 1981) β_{TSP} , β_{MST} , β_{STEINER} , such that*

$$\lim_{n \rightarrow \infty} \frac{L_{\text{TSP}}^n}{\sqrt{n}} = \beta_{\text{TSP}}$$

$$\lim_{n \rightarrow \infty} \frac{L_{\text{MST}}^n}{\sqrt{n}} = \beta_{\text{MST}}$$

$$\lim_{n \rightarrow \infty} \frac{L_{\text{STEINER}}^n}{\sqrt{n}} = \beta_{\text{STEINER}}.$$

For the VRP (Haimovitch and Rinnooy Kan 1985)

$$\lim_{n \rightarrow \infty} \frac{Q_n E[L_{\text{VRP}}^n]}{n} = 2E[r] \quad \text{if } Q_n = o(\sqrt{n})$$

$$\lim_{n \rightarrow \infty} \frac{E[L_{\text{VRP}}^n]}{\sqrt{n}} = \beta_{\text{TSP}} \quad \text{if } Q_n = \Omega(\sqrt{n})$$

where $E(r)$ is the expected radial distance from the depot to a point in $X^{(n)}$.

We now characterize the expectation of the reoptimization strategy for each problem assuming that each of the n points is present with the same constant probability p , which is called the *coverage probability*. In the following theorem, the expectation is taken over all possible 2^n instances of the problem and the probability 1 statement refers to the random locations of the points.

Theorem 2. (Jaillet 1985, Bertsimas 1988). *With probability 1*

$$\lim_{n \rightarrow \infty} \frac{E[\Sigma_{\text{TSP}}^n]}{\sqrt{n}} = \beta_{\text{TSP}} \sqrt{p}$$

$$\lim_{n \rightarrow \infty} \frac{E[\Sigma_{\text{MST}}^n]}{\sqrt{n}} = \beta_{\text{MST}} \sqrt{p}$$

$$\lim_{n \rightarrow \infty} \frac{E[\Sigma_{\text{STEINER}}^n]}{\sqrt{n}} = \beta_{\text{STEINER}} \sqrt{p}$$

$$\lim_{n \rightarrow \infty} \frac{Q_n E[\Sigma_{\text{VRP}}^n]}{n} = 2E[r]p \quad \text{if } Q_n = o(\sqrt{n})$$

$$\lim_{n \rightarrow \infty} \frac{E[\Sigma_{\text{VRP}}^n]}{\sqrt{n}} = \beta_{\text{TSP}} \sqrt{p} \quad \text{if } Q_n = \Omega(\sqrt{n})$$

$$\lim_{n \rightarrow \infty} \frac{E[\Sigma_{\text{TSFLP}}^n(i)]}{\sqrt{n}} = \beta_{\text{TSP}} \sqrt{p}, \quad \text{for any } i$$

where $E(r)$ is the expected radial distance.

Proof. The main idea in the proof is that the principal contribution to $E[\Sigma^n]$ comes from the sets S with $|S| \in [lnp(1 + \epsilon) \downarrow, \lceil np(1 + \epsilon) \rceil]$. The reason is that the number of points present is given by a binomial distribution with parameters n , p , and hence, the probability mass function is concentrated within ϵ of np . In this range of $|S|$ we can apply Theorem 1 to obtain Theorem 2. We illustrate the idea with respect to the TSP reoptimization strategy in Appendix A.

Intuitively Theory 2 means that solutions under the reoptimization strategy behave asymptotically simi-

larly to those of the corresponding combinatorial optimization problems, but on np rather than n points. The asymptotic behavior of the VRP reoptimization strategy suggests that the strategy behaves like the TSP reoptimization strategy if the capacity Q is large, a property that is quite intuitive. If the capacity Q is small, the vehicle has to make many trips back to the depot, so that the radial collection term $(2E[r]np/Q_n)$ rather than the routing component dominates. For the TSFLP reoptimization strategy we observe that asymptotically the location component of the problem is unimportant, because the same asymptotic behavior is observed irrespective of the location decision i .

We next characterize asymptotically the a priori optimization strategies.

Theorem 3. (Jaillet 1985, Bertsimas 1988). *With probability 1*

$$\lim_{n \rightarrow \infty} \frac{E[L_{\text{PTSP}}^n]}{\sqrt{n}} = \beta_{\text{TSP}}(p)$$

$$\lim_{n \rightarrow \infty} \frac{E[L_{\text{PMST}}^n]}{\sqrt{n}} = \beta_{\text{MST}}(p)$$

$$\lim_{n \rightarrow \infty} \frac{Q_n E[L_{\text{PVRPa}}^n]}{n} = 2E[r]p \quad \text{if } Q_n = o(\sqrt{n})$$

$$\lim_{n \rightarrow \infty} \frac{E[L_{\text{PVRPa}}^n]}{\sqrt{n}} = \beta_{\text{TSP}} \quad \text{if } Q_n = \Omega(\sqrt{n})$$

$$\lim_{n \rightarrow \infty} \frac{Q_n E[L_{\text{PVRPb}}^n]}{n} = 2E[r]p \quad \text{if } Q_n = o(\sqrt{n})$$

$$\lim_{n \rightarrow \infty} \frac{E[L_{\text{PVRPb}}^n]}{\sqrt{n}} = \beta_{\text{TSP}}(p) \quad \text{if } Q_n = \Omega(\sqrt{n})$$

$$\lim_{n \rightarrow \infty} \frac{E[L_{\text{PTSFLP}}^n(i)]}{\sqrt{n}} = \beta_{\text{TSP}}(p) \quad \text{for any } i$$

where $E(r)$ is the expected radial distance.

Sketch of the Proof. We first prove that the PTSP and the PMST belong to the class of subadditive Euclidean functionals whose asymptotic behavior has been characterized by Steele (1981). Their value is almost surely asymptotic to $c\sqrt{n}$, where c depends on the functional.

For the PVRP and the PTSFLP we find tight upper and lower bounds from which we can characterize the asymptotic behavior. For the PVRP under method a ,

for example, we prove that

$$\begin{aligned} & \max\left(\frac{2p}{Q} \sum_{i=1}^n d(0, i), L_{\text{TSP}}\right) \\ & \leq E[L_{\text{PVRP}_Q}] \\ & \leq L_{\text{TSP}}\left(1 - \frac{2}{n} - \frac{p}{Q}\right) + 2\left(2 + \frac{np}{Q}\right) \frac{\sum_{i=1}^n d(0, i)}{n} \end{aligned}$$

where $d(0, i)$ denotes the distance between the depot and node i .

In order to illustrate the techniques we are using, we present in detail the argument for the PTSP in Appendix B.

Comparing Theorems 2 and 3 we can observe that the a priori and reoptimization strategies have very close asymptotic performance almost surely. The result may be considered surprising in view of the fact that a priori strategies require the computation of only one solution and are easily updated, while reoptimization strategies require the computation of an optimal solution for every problem instance. Yet, a priori strategies behave asymptotically equally well on average with reoptimization strategies on Euclidean problems. In addition, we conjecture that a priori and reoptimization strategies have exactly the same asymptotic performance almost surely, i.e., $\beta_{\text{PTSP}}(p) = \beta_{\text{TSP}}\sqrt{p}$ and $\beta_{\text{PMST}}(p) = \beta_{\text{MST}}\sqrt{p}$.

3. THE COMPLEXITY OF A PRIORI OPTIMIZATION

In the previous section, we showed that in terms of performance a priori strategies are attractive compared with reoptimization strategies. In this section, we address the question of how difficult it is to find the optimal a priori solutions from a computational complexity perspective.

We first introduce the decision version of a PCOP. Given a complete graph $G = (V, E)$, $|V| = n$, a cost $d: E \rightarrow R$, a vector (p_1, \dots, p_n) of the probabilities of the presence of the vertices and a bound B , does there exist a structure f (a tour, a tree, a route, a tour and a vertex for the PTSP, PMST, PVRP and PTSFLP, respectively) such that

$$E[L_f] \leq B?$$

We can then characterize the complexity of a priori strategies as follows.

Theorem 4. (Bertsimas 1988). *The decision version of all four PCOPs is NP-complete.*

Sketch of the Proof. For the cases of the PTSP, PVRP and PTSFLP we only need to show membership in NP because, as noted earlier, these three problems are generalizations of well known NP-complete problems. Membership in NP is seen to hold, because given a solution f we can compute $E[L_f]$ in $O(n^2)$. For example, for the PTSP if the tour is $\tau = (1, 2, \dots, n, 1)$ then by looking at the probability of every link present we can derive (Jaillet 1988) the expression

$$\begin{aligned} E[L_\tau] &= \sum_{i=1}^n \sum_{j=i+1}^n d(i, j) p_i p_j \prod_{k=i+1}^{j-1} (1 - p_k) \\ &+ \sum_{j=1}^n \sum_{i=1}^{j-1} d(j, i) p_i p_j \\ &\cdot \prod_{k=j+1}^n (1 - p_k) \prod_{k=1}^{i-1} (1 - p_k). \end{aligned} \tag{5}$$

The case of the PMST is more difficult because the PMST problem is not a generalization of an NP-complete problem, since the MST is solved by a greedy algorithm in $O(n^2)$. Membership in NP holds because of the following closed form expression for the expected length $E[L_T]$ of a given a priori tree T (Bertsimas 1990)

$$\begin{aligned} E[L_T] &= \sum_{e \in T} c(e) \left\{ 1 - \prod_{i \in K_e} (1 - p_i) \right\} \\ &\cdot \left\{ 1 - \prod_{i \in V - K_e} (1 - p_i) \right\} \end{aligned}$$

where $K_e, V - K_e$ are the subsets of nodes contained in the two subtrees obtained from T by removing the edge e from T .

We have proved the completeness of the PMST by a reduction from the problem EXACT COVER BY 3-SETS, which is NP-complete (see Garey and Johnson 1979). For details see Bertsimas (1988).

Thus, although we can efficiently compute the expected length of any given a priori solution to a PCOP, it is still NP-hard to find an optimal a priori solution.

4. THEORETICAL APPROXIMATIONS TO OPTIMAL A PRIORI SOLUTIONS

In the previous section, we found that it is still NP-hard to obtain optimal a priori solutions to the PCOPs. In this section, we address the question of approximating the optimal a priori solutions with polynomial time heuristics, whose worst case behavior we can characterize.

The first natural question to address is how heuristic approaches to the deterministic problem perform when applied to the corresponding probabilistic problem. For example, what is the performance of the well known Christofides heuristic for the TSP (see Larson and Odoni 1981) if it is applied to the PTSP? In order to find useful bounds for the routing problems (PTSP, PVRP) we assume that the triangle inequality holds. We can then prove the following theorem.

Theorem 5. (Bertsimas 1988). *Let L_D be the length of the optimal solution to the deterministic TSP, MST or VRP and let L_H be the length of a heuristic solution to the same problem. Let p be the coverage probability and $E[L_p]$ the expected length of the optimum a priori solution to the corresponding PCOP. If the heuristic has the property that*

$$\frac{L_H}{L_D} \leq c$$

then

$$E\left[\frac{L_H}{E[L_p]}\right] \leq \frac{c}{p}.$$

Sketch of the Proof. In all cases, we show that $E[L_f] \leq L_f$. Here we use the triangle inequality in the case of the two routing problems. Also, $E[L_p] \geq pL_D$. Combining these inequalities the result follows. For the details see Bertsimas (1988).

Theorem 5 suggests that if the coverage probability is large, the constant guarantee heuristics for the deterministic problem still behave well for the corresponding probabilistic problem. But if $p \rightarrow 0$ the bound is not informative and indeed one can find examples with $p \rightarrow 0$, $np \rightarrow \infty$ for which $(E[L_D]/E[L_p]) \rightarrow \infty$, that is, even if $c = 1$, the optimal deterministic solution is an arbitrarily bad approximation to the optimal a priori solution. As an indication of the rate at which the ratio $(E[L_D]/E[L_p])$ tends to infinity, we can prove the following theorem.

Theorem 6. (Bertsimas 1988). *For the PTSP with triangle inequality*

$$\frac{E[L_D]}{E[L_p]} = O(\sqrt{n}).$$

We next investigate the existence of constant guarantee heuristics for the routing problems we are considering. We restrict our attention to Euclidean problems and examine the spacefilling curve heuristic, first introduced by Kakutani (see *The Collected Work*

of S. Kakutani, Vol II, p. 444, 1966) and proposed by Platzman and Bartholdi (1983) for the Euclidean TSP. The spacefilling curve heuristic can be described 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 $f: R^2 \rightarrow R$ is called the Sierpinski curve (for details on the computation of $f(x, y)$, see Bartholdi and Platzman 1982).
2. Sort the numbers $f(x_i, y_i)$ and visit the corresponding initial points (x_i, y_i) in that order, producing a tour τ_{SF} .

The key property of the spacefilling curve heuristic that makes its analysis for the PTSP possible is the following: Consider an instance S of the problem. Suppose that the spacefilling curve heuristic produces a tour $\tau_{SF}(S)$ if we run the heuristic on the instance S . Consider the tour τ_{SF} produced by the heuristic on the original instance of the problem, i.e., when all points are present. 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 $\tau_{SF}(S)$, because sorting has the property of preserving the order in which the points in S will be visited by the spacefilling curve, which is exactly the property of the PTSP strategy as well. Based on this critical observation we can then analyze the spacefilling curve heuristic.

Theorem 7. (Bertsimas 1988). *For the Euclidean PTSP and PVRP under method b the spacefilling curve heuristic produces a tour τ_{SF} with the property*

$$\begin{aligned} \frac{E[L_{\tau_{SF}}]}{E[L_{\tau_p}]} &\leq \frac{E[L_{\tau_{SF}}]}{E[\Sigma_{TSP}]} = O(\log n). \\ \frac{E[L_{\tau_{SF}}]}{E[L_{PVRP_b}]} &\leq \frac{E[L_{\tau_{SF}}]}{E[\Sigma_{VRP}]} = Q + O(\log n). \end{aligned} \tag{6}$$

Sketch of the Proof. In Platzman and Bartholdi it is proven that the length of the spacefilling curve heuristic satisfies

$$\frac{L_{\tau_{SF}}}{L_{TSP}} = O(\log n).$$

Consider an instance S of the problem. If the spacefilling curve heuristic is applied to the instance S , it will similarly produce a tour $\tau_{SF}(S)$ with length

$$\frac{L_{\tau_{SF}}(S)}{L_{TSP}(S)} = O(\log |S|) = O(\log n).$$

But since $\tau_{\text{SF}}(S)$ is the tour produced by the PTSP strategy at instance S then

$$\begin{aligned} \frac{E[L_{\tau_{\text{SF}}}] }{E[\Sigma_{\text{TSP}}]} &= \frac{\sum_{S \subseteq V} p(S) L_{\tau_{\text{SF}}}(S)}{\sum_{S \subseteq V} p(S) L_{\text{TSP}}(S)} \\ &\leq \frac{\sum_{S \subseteq V} p(S) O(\log n) L_{\text{TSP}}(S)}{\sum_{S \subseteq V} p(S) L_{\text{TSP}}(S)} = O(\log n). \end{aligned}$$

Note that this 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 spacefilling curve heuristic ignores the probabilistic nature of the problem but surprisingly produces a tour which is globally, in every instance, close to the optimal. A similar argument holds for the PVRP under method b .

As a corollary to Theorem 7 we can compare the PTSP and the reoptimization strategies from a worst case perspective. For the Euclidean PTSP, since $E[L_{\tau_p}] \leq E[L_{\tau_{\text{SF}}}]$

$$\frac{E[L_{\tau_p}]}{E[\Sigma_{\text{TSP}}]} = O(\log n).$$

Platzman and Bartholdi conjecture that the spacefilling curve heuristic is a constant-guarantee heuristic. Unfortunately, Bertsimas and Grigni (1989) refute the conjecture by exhibiting an example in which the $O(\log n)$ bound is tight.

For the PTSFLP for which node i needs a visit with probability p_i we consider the following location heuristic.

Spacefilling Curve Location Heuristic

1. Given the coordinates the locations of the customers use the spacefilling curve heuristic to find the a priori tour τ_{SF} .
2. Compute $h(i, \tau_{\text{SF}})$ with a vector of probabilities $(p_1, \dots, p_{i-1}, 1, p_{i+1}, \dots, p_n)$ for every node i .
3. Select the point i_{SF} that minimizes $h(i, \tau_{\text{SF}})$. Location i_{SF} and the tour τ_{SF} are the proposed solutions to the PTSFLP.

Using similar techniques with Theorem 7 we can analyze the worst case error of the heuristic.

Theorem 8. (Bertsimas 1989). *If $p_i = \Omega(1/\log n)$ for all i , then*

$$\frac{h(i_{\text{SF}}, \tau_{\text{SF}})}{E[\Sigma_{\text{TSFLP}}(i^*)]} = O(\log n) \quad (7)$$

where i^* is the optimal location for the TSFLP.

The final question concerns the heuristic's running time. Step 1 can be performed in $O(n \log n)$. A straightforward implementation of Step 2 can be performed in $O(n^3)$, since we can calculate $h(i, \tau_{\text{SF}})$ for each i in $O(n^2)$ from (5) (it is the expected length of the a priori tour τ_{SF}). By noticing that the only difference between calculating $h(i, \tau_{\text{SF}})$ and $h(i+1, \tau_{\text{SF}})$ is due to the corresponding probability vectors, which differ solely in the i th and $(i+1)$ th position, we can calculate $h(i+1, \tau_{\text{SF}})$ in $O(n)$ given $h(i, \tau_{\text{SF}})$, since only the contribution of $O(n)$ distances is different in the two expectations, namely the contributions of the edges $d(i, j)$, $d(i+1, j)$ for $j = 1, \dots, n$. Thus, we can compute $h(1, \tau_{\text{SF}})$ in $O(n^2)$, and then compute $h(i+1, \tau_{\text{SF}})$ from $h(i, \tau_{\text{SF}})$ in $O(n)$. The total computation takes $O(n^2)$ time. Step 3 clearly takes $O(n)$ time. As a result, the overall heuristic can be implemented in $O(n^2)$ time.

For the PVRP under updating method a , Bertsimas (1988) proposes an $O(n^3)$ heuristic which produces a route with expected length $\frac{1}{2}$ from the optimal solution.

5. PRACTICAL APPROXIMATIONS TO OPTIMAL A PRIORI SOLUTIONS

In this section, we briefly discuss some of our experience in trying to find useful heuristic solutions to PCOPs using the a priori optimization approach. We use the Euclidean PTSP as an example, since we have sharply characterized its asymptotic behavior, so we know that the expected length of the optimal solution would be close to $\beta_{\text{TSP}} \sqrt{np}$ for random problems. This can be used as a benchmark to compare the performance of various heuristics.

In our numerical experiments we obtained near-optimal solutions to Euclidean PTSPs by means of two different types of heuristics. The first is the spacefilling curve heuristic, while the second is based on seeking local optimality. Our implementation of the spacefilling curve heuristic 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, this is even faster than the computation of the expected length of that tour, $E[L_{\tau_{\text{SF}}}]$, which requires $O(n^2)$ time. Since the computed tour τ_{SF} is independent of the probabilities p_i , the spacefilling 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

$$E[L_\tau] \leq \min_{\tau' \in S(\tau)} E[L_{\tau'}]. \tag{8}$$

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 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.

Lin (1965) used an iterative improvement algorithm for the TSP based on what he called the λ -opt local neighborhood. For a given tour τ that consists of n links between nodes, the neighborhood $S_\lambda(\tau)$ consists of those tours that differ from τ by no more than λ links. For $\lambda = 2$, this is the set of tours that can be obtained by reversing a section of τ ; for $\lambda = 3$, it is the set of tours obtained 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.

Unlike the TSP case, the expected length $E[L_\tau]$ in the PTSP sense depends on all $(n^2 - n)/2$ independent elements of the distance matrix. We cannot, therefore, speak of some links leaving and others entering the tour; rather, it is only the weight given to each of the $d(i, j)$ by (5) 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 (referred to heretofore as "2- p -opt") 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.

Another neighborhood that we tried consists of 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 , it is a subset of S_3 , and yields much better results than S_2 in our experiments.

A summary of the behavior of each of the heuristics

is shown in Figure 4. The spacefilling curve solutions were used as starting positions for the 2- p -opt 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- p -opt 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- p -opt 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 spacefilling curve algorithm, followed by 3-opt if p is fairly large, and then finish by applying 1-shift. The threshold point below which 3-opt ceases to be helpful is uncertain and probably depends strongly on the specifics of the problem. For problems with more than a few hundred nodes both the running time and the memory required for the distance matrix and the auxiliary matrices begin to be excessive. At that point we were

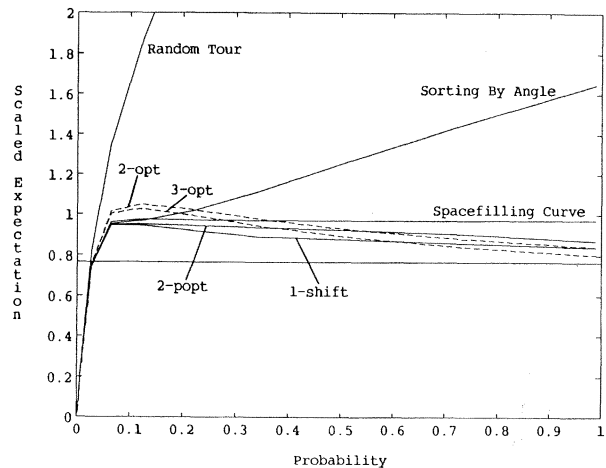


Figure 4. A summary of results for several PTSP heuristics on 100-node problems scaled by \sqrt{np} . Solutions obtained via the 2-opt and 3-opt TSP algorithms (dashed lines) are shown for comparison. The horizontal line shows the value of $\beta_{TSP} \cong 0.765$. The heuristics are: 1) random tour,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

forced to switch to heuristics like the spacefilling curve algorithm, which do not require $O(n^2)$ memory.

In the calculations for Figure 4 results from 10 separate 100-node problems were averaged to minimize the effects of statistical fluctuations. The locations of the nodes for each problem were chosen from a uniform distribution in the unit square, and the expected lengths $E[L_\tau]$ were scaled by \sqrt{np} . The asymptotic results of Section 2 would then lead us to expect that data from optimal tours would follow a horizontal line on the plot. Our heuristics confirm this behavior except when p is small. The reason is that if p is small, 100 points are not enough for the expected length of the optimal PTSP to reach its asymptotic value.

6. SOME CONCLUDING REMARKS

This paper has introduced the idea of a priori optimization, an approach which may be competitive, especially in many practical contexts, with the strategy of reoptimization, under which every possible instance of the problem is solved to optimality. A priori optimization strategies were applied to four problems: the TSP, the VRP, the MST and the TSFLP. In all cases, the a priori strategies have potential areas of application in such fields as communications, transportation, routing, VLSI design, scheduling, strategic and organizational planning.

It was shown that for all problems defined here a priori and reoptimization strategies have on average very close asymptotic behavior, a property that further underscores the importance of studying a priori strategies. We then characterized the complexity of the two types of strategies and proposed heuristics for the PCOPs.

Further generalizations of these ideas include stochastic demands which are not only binary (demand of one unit with a certain probability), but can be any random variable. This generalization is especially important in the case of the vehicle routing problem. Another important extension is the inclusion of a dynamic component in the problems, i.e., demands are generated over time according to a stochastic process. In this case queuing phenomena arise which are interesting in themselves. A step in this direction is taken in Bertsimas and van Ryzin (1990), in which the authors analyze a dynamic version of the traveling repairman problem.

The paper has attempted to indicate the wide range of questions that can be addressed with respect to the idea of a priori optimization, the novel and interesting aspects introduced by it, and finally, the excellent

potential for deriving new results and solution procedures and for applying them to important contexts.

APPENDIX A

Proof of Theorem 2 for the PTSP

Let W be the number of nodes present and

$$h_k \triangleq \sum_{S:|S|=k} L_{\text{TSP}}(X^{(n)}; S) / \binom{n}{k}$$

where $L_{\text{TSP}}(X^{(n)}; S)$ is the length of the TSP on the set S . Then

$$\begin{aligned} E[\sum_{\text{TSP}}^n(X^{(n)})] &= \sum_{k=0}^n \binom{n}{k} p^k (1-p)^{n-k} h_k \\ &= \sum_{k=0}^n \Pr\{W = k\} h_k. \end{aligned}$$

Fix $\epsilon > 0$. Then

$$\begin{aligned} E[\sum_{\text{TSP}}^n(X^{(n)})] &= \sum_{k=0}^{\lfloor np(1-\epsilon) \rfloor - 1} \Pr\{W = k\} h_k \\ &+ \sum_{k=\lceil np(1+\epsilon) \rceil + 1}^n \Pr\{W = k\} h_k \\ &+ \sum_{k=\lceil np(1-\epsilon) \rceil}^{\lfloor np(1+\epsilon) \rfloor} \Pr\{W = k\} h_k. \end{aligned}$$

Since $L_{\text{TSP}}(X^{(n)}; S) < c\sqrt{|S|}$ for some constant c , then $h_k \leq c\sqrt{n}$. As a result

$$\begin{aligned} &\sum_{k=0}^{\lfloor np(1-\epsilon) \rfloor - 1} \Pr\{W = k\} h_k + \sum_{k=\lceil np(1+\epsilon) \rceil + 1}^n \Pr\{W = k\} h_k \\ &\leq c\sqrt{n} \Pr\{|W - np| > np\epsilon\}. \end{aligned}$$

From the Chernoff bound we have

$$\begin{aligned} \Pr\{|W - np| > np\epsilon\} &< 2 \left[\frac{e^\epsilon}{(1+\epsilon)^{1+\epsilon}} \right]^{np} = 2\delta^n, \\ &0 < \delta < 1. \end{aligned}$$

The contribution of the first two terms is then

$$\begin{aligned} &\sum_{k=0}^{\lfloor np(1-\epsilon) \rfloor - 1} \Pr\{W = k\} h_k + \sum_{k=\lceil np(1+\epsilon) \rceil + 1}^n \Pr\{W = k\} h_k \\ &< 2c\sqrt{n}\delta^n, \quad \delta < 1. \end{aligned}$$

For $\lfloor np(1-\epsilon) \rfloor \leq k \leq \lceil np(1+\epsilon) \rceil$ we apply Theorem 1 and obtain that with probability 1 for any

$\epsilon > 0, \exists k_\epsilon$: for any S , with

$$|S| = k \geq k_\epsilon - \epsilon \leq \frac{L_{\text{TSP}}(X^{(n)}; S)}{\sqrt{k}} - \beta_{\text{TSP}}$$

$$\leq \epsilon \Rightarrow -\epsilon \leq \frac{h_k}{\sqrt{k}} - \beta_{\text{TSP}} \leq \epsilon.$$

In addition

$$\sum_{k=\lfloor np(1-\epsilon) \rfloor}^{\lceil np(1+\epsilon) \rceil} \Pr\{W = k\}$$

$$= \Pr\{|W - np| \leq n\epsilon\} > 1 - 2\delta^n.$$

Therefore

$$(\beta_{\text{TSP}} - \epsilon)(1 - 2\delta^n)$$

$$< \sum_{k=\lfloor np(1-\epsilon) \rfloor}^{\lceil np(1+\epsilon) \rceil} \Pr\{W = k\} \frac{h_k}{\sqrt{k}} < (\beta_{\text{TSP}} + \epsilon)1$$

from which

$$(\beta_{\text{TSP}} - \epsilon)(1 - 2\delta^n)\sqrt{p(1-\epsilon)}$$

$$< \sum_{k=\lfloor np(1-\epsilon) \rfloor}^{\lceil np(1+\epsilon) \rceil} \frac{\Pr\{W = k\}h_k}{\sqrt{n}} < (\beta_{\text{TSP}} + \epsilon)\sqrt{p(1+\epsilon)}.$$

Combining the above bounds, we find that almost surely for any $\epsilon > 0$, for any $n \geq (k_\epsilon/p(1-\epsilon))$

$$(\beta_{\text{TSP}} - \epsilon)(1 - 2\delta^n)\sqrt{p(1-\epsilon)}$$

$$< \frac{E[\sum_{\text{TSP}}^n(X^{(n)})]}{\sqrt{n}} < (\beta_{\text{TSP}} + \epsilon)\sqrt{p(1+\epsilon)} + 2c\delta^n.$$

Since ϵ can be arbitrarily small we let $\epsilon \rightarrow 0$, and thus we prove the theorem.

APPENDIX B

Proof of Theorem 3 for the PTSP

Let τ_p be the optimum PTSP tour. Clearly $E[L_{\text{PTSP}}^n] = E[L_{\tau_p}^n]$. We will first prove that with probability 1 $\lim_{n \rightarrow \infty} E[L_{\tau_p}^n(X^{(n)})]/\sqrt{n}$ exists. In order to do this we check whether the functional

$$f(X^{(n)}) \triangleq E[L_{\tau_p}^n(X^{(n)})]$$

is a subadditive monotone Euclidean functional (Steele 1981).

1. $f(X^{(n)})$ is Euclidean, because clearly it is invariant under translation, that is

$$f(X^{(n)} + \xi) = f(X^{(n)})$$

and it is linear

$$f(aX^{(n)}) = af(X^{(n)}).$$

2. $f(X^{(n)})$ is monotone, because clearly

$$f(\{\xi\} \cup X^{(n)}) \geq f(X^{(n)}).$$

3. Clearly $f(X^{(n)})$ has finite variance

$$\text{Var}[f(X^{(n)})] < \infty.$$

4. $f(X^{(n)})$ is subadditive, i.e., if $Q_i, i = 1, \dots, m^2$ is a partition of the unit square in m^2 subsquares then

$$f(X^{(n)} \cap [0, r]^2)$$

$$\leq \sum_{i=1}^{m^2} f(X^{(n)} \cap rQ_i) + crm.$$

It is not clear that the subadditivity property holds for the PTSP. We will next concentrate on proving this property. Consider the following algorithm.

Step 1. For every nonempty subsquare Q_i construct the optimal PTSP tour τ_i for the points $X^{(n)} \cap rQ_i$.

Step 2. Select arbitrarily a point from $X^{(n)} \cap rQ_i$ in each nonempty subsquare and call it a representative. Consider the representatives as points always present (*black points*).

Step 3. Construct a TSP tour τ^* among the representatives.

Step 4. The PTSP tours τ_i and the τ^* create a closed walk τ , which connects all the points $X^{(n)}$.

The expected length of the tour τ is

$$E[L_\tau] = \sum_{i=1}^{m^2} f_i(X^{(n)} \cap rQ_i) + L_{\tau^*}$$

where $f_i(X^{(n)} \cap rQ_i)$ is the expected length of the tour τ_i in which one point, the representative, is always present (it is a *black node*) and all the others have a probability p of being present. If we turn a *black node* into a *white node* (a node which has a probability p of being present), the expected length of the closed walk clearly decreases and so it does if we also transform the closed walk into a tour. The resulting tour has an expected length not smaller than $E[L_{\tau_p}]$, since by definition τ_p is the optimal PTSP. Then

$$E[L_{\tau_p}] \leq \sum_{i=1}^{m^2} f_i(X^{(n)} \cap rQ_i) + L_{\tau^*}.$$

It is well known (Larson and Odoni) that

$$L_{\tau^*} \leq b\sqrt{m^2 r^2} = brm$$

that is, the optimal TSP tour among l points in an area A is less than $b\sqrt{lA}$ for some constant b . In our case $l \leq m^2$ and $A = r^2$. The question now is to relate $f_i(X^{(n)} \cap rQ_i)$ with $f(X^{(n)} \cap rQ_i)$, or equivalently, $E[L_{\tau_i}]$ with $E[L_{\tau_i} | \text{a node is black}]$. Without loss of

generality, assume that the optimal PTSP through the points $X^{(n)} \cap rQ_i$ is $(x_1, x_2, \dots, x_{k_i}, x_1)$ where $k_i = |X^{(n)} \cap rQ_i|$. If we consider x_1 to be the black node, then it is easy to prove that $f_i(X^{(n)} \cap rQ_i) \leq f(X^{(n)} \cap rQ_i) + 2(1-p)\max_{1 < j \leq k_i} |x_1 - x_j|$. Since $\max_{1 < j \leq k_i} |x_1 - x_j| \leq \sqrt{2}r/m$ we finally get

$$f_i(X^{(n)} \cap rQ_i) \leq f(X^{(n)} \cap rQ_i) + 2(1-p)\sqrt{2}r/m.$$

Therefore, we can conclude that

$$E[L_{r,p}] = f(X^{(n)} \cap [0, r]^2) \leq \sum_{i=1}^{m^2} f(X^{(n)} \cap rQ_i) + (b + \sqrt{2}(1-p)/p)rm$$

which means that the PTSP is subadditive.

Monotone subadditive Euclidean functionals are almost surely asymptotic to $\beta\sqrt{n}$. In our case, there exists a constant $\beta_{\text{TSP}}(p)$ such that with probability 1

$$\lim_{n \rightarrow \infty} E[L_{r,p}^n(X^{(n)})]/\sqrt{n} = \beta_{\text{TSP}}(p).$$

Furthermore, the following bounds on $\beta_{\text{TSP}}(p)$ can be established (see Jaillet 1985)

$$\beta_{\text{TSP}}\sqrt{p} \leq \beta_{\text{TSP}}(p) \leq \min[\beta_{\text{TSP}}, 0.92\sqrt{p}]$$

i.e., $\beta_{\text{TSP}}(p) = \Theta(\sqrt{p})$. For details about the other problems see Bertsimas (1988).

ACKNOWLEDGMENT

The research of the first and third authors is partially supported by the National Science Foundation under grant ECS-8717970.

REFERENCES

- BARTHOLDI, J. J., AND L. K. PLATZMAN. 1982. On $O(N \log N)$ Planar Traveling Salesman Heuristic Based on Spacefilling Curves. *Opns. Res. Lett.* **1**, 121–125.
- BARTHOLDI, J. J., L. K. PLATZMAN, R. LEE COLLINS AND W. H. WARDEN. 1983. A Minimal Technology Routing System for Meals on Wheels. *Interfaces* **13**, 1–8.
- BEARDWOOD, J., J. HALTON AND J. HAMMERSLEY. 1959. The Shortest Path Through Many Points. *Proc. Camb. Philos. Soc.* **55**, 299–327.
- BERMAN, O., AND D. SIMCHI-LEVI. 1986. Minisum Location of a Traveling Salesman. *Networks* **16**, 329–354.
- BERMAN, O., AND D. SIMCHI-LEVI. 1988a. Finding the

- Optimal *a priori* Tour and Location of a Traveling Salesman With Non-Homogeneous Customers. *Trans. Sci.* **22**, 148–154.
- BERMAN, O., AND D. SIMCHI-LEVI. 1988b. A Heuristic Algorithm for the Traveling Salesman Location Problem on Networks. *Opns. Res.* **36**, 478–484.
- BERTSIMAS, D. 1988. Probabilistic Combinatorial Optimization Problems. Ph.D. Thesis, Operations Research Center, MIT, Cambridge, Mass.
- BERTSIMAS, D. 1989. On Probabilistic Traveling Salesman Facility Location Problems. *Trans. Sci.* **3**, 184–191.
- BERTSIMAS, D. 1990. The Probabilistic Minimum Spanning Tree Problem. *Networks* **20**, 245–275.
- BERTSIMAS, D., AND M. GRIGNI. 1989. Worst Case Examples for the Spacefilling Curve Heuristic for the Euclidean Traveling Salesman Problem. *Opns. Res. Lett.* **8**, 241–244.
- BERTSIMAS, D., AND G. VAN RYZIN. 1990. A Stochastic and Dynamic Vehicle Routing Problem in the Euclidean Plane. *Opns. Res.* (to appear).
- BODIN, L., B. GOLDEN, A. ASSAD AND M. BALL. 1983. Routing and Scheduling of Vehicles and Crews, the State of the Art. *Comput. Opns. Res.* **10**, 69–211.
- BURNES, R., AND J. WHITE. 1976. The Traveling Salesman Location Problem. *Trans. Sci.* **10**, 348–360.
- DROR, M., AND P. TRUDEAU. 1986. Stochastic Vehicle Routing With Modified Savings Algorithm. *Eur. J. Opnl. Res.* **23**, 228–235.
- FRIEZE, A. M. 1985. On the Value of a Minimum Spanning Tree Problem. *Discr. Appl. Math.* **10**, 47–56.
- GAREY, M., AND D. JOHNSON. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman, San Francisco.
- GOLDEN, G., AND A. ASSAD. 1988. *Vehicle Routing; Methods and Studies*. North Holland, Amsterdam.
- HAIMOVITCH, M., AND A. RINNOOY KAN. 1985. Bounds and Heuristics for Capacitated Routing Problems. *Math. Opns. Res.* **10**, 527–542.
- JAILLET, P. 1985. Probabilistic Traveling Salesman Problems. Technical Report No. 185, Operations Research Center, MIT, Cambridge, Mass.
- JAILLET, P. 1988. A priori Solution of a Traveling Salesman Problem in Which a Random Subset of the Customers Are Visited. *Opns. Res.* **36**, 929–936.
- JAILLET, P., AND A. ODONI. 1988. The Probabilistic Vehicle Routing Problem. In *Vehicle Routing; Methods and Studies*, B.L. Golden and A. A. Assad (eds.). North Holland, Amsterdam.
- KARP, R. 1977. Probabilistic Analysis of Partitioning Algorithms for the Traveling Salesman in the Plane. *Math. Opns. Res.* **2**, 209–224.
- KARP, R. 1979. A Patching Algorithm for the Non-symmetric Traveling Salesman Problem. *SIAM J. Comput.* **8**, 561–573.
- LAPORTE, G., AND F. LOUVEAU. 1987. Formulations and Bounds for the Stochastic Capacitated Vehicle

- Routing Problem With Uncertain Supplies. G-87-23, Ecole des Hautes Etudes Commerciale, University of Montreal, Canada.
- LAPORTE, G., F. LOUVEAU AND H. MERCURE. 1987. Models and Exact Solutions for a Class of Stochastic Location-Routing Problems. G-87-14, Ecole des Hautes Etudes Commerciale, University of Montreal, Canada.
- LARSON, R., AND A. ODONI. 1981. *Urban Operations Research*. Prentice-Hall, Englewood Cliffs, N.J.
- LAWLER, E. L., J. K. LENSTRA, A. RINNOOY KAN AND D. B. SHMOYS (eds.). 1985. *The Traveling Salesman Problem: A Guided Tour of Combinatorial Optimization*. Wiley, Chichester.
- LIN, S. 1965. Computer Solutions for the Traveling Salesman Problem. *Bell Syst. Tech. J.* 2245-2269.
- PLATZMAN, L. K., AND J. J. BARTHOLDI. 1983. Spacefilling Curves and the Planar Traveling Salesman Problem. Report No. PDRC 83-02, School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta.
- SIMCHI-LEVI, D., AND O. BERMAN. 1988. Heuristics and Bounds for the Traveling Salesman Location Problem on the Plane. *Opns. Res. Lett.* 5, 243-248.
- STEELE, J. M. 1981. Subadditive Euclidean Functionals and Nonlinear Growth in Geometric Probability. *Ann. Prob.* 9, 365-376.
- STEELE, J. M. 1987. On Frieze's $\zeta(3)$ Limit for Lengths of Minimal Spanning Trees. *Disc. Appl. Math.* 18, 99-103.
- STEWART, W., AND B. GOLDEN. 1983. Stochastic Vehicle Routing: A Comprehensive Approach. *Eur. J. Opnl. Res.* 14, 371-385.
- TILLMAN, F. 1969. The Multiple Terminal Delivery Problem With Probabilistic Demands. *Trans. Sci.* 3, 192-204.