A PRIORI OPTIMIZATION

Dimitris J. Bertsimas,
Patrick Jaillet
and
Amedeo R. Odoni

Sloan W.P. No. 3059-89-MS August 1989

A priori optimization

*Dimitris J. Bertsimas**

Sloan School of Management, Rm E53-359
Massachusetts Institute of Technology
Cambridge, MA 02139, USA

Patrick Jaillet

Centre de Mathématiques Appliquées
Ecole Nationale des Ponts et Chaussées
93167 NOISY LE GRAND CEDEX, France

*Amedeo R. Odoni**

Operations Research Center, Rm E40-165
Massachusetts Institute of Technology
Cambridge, MA 02139, USA

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

Abstract

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 present with a certain probability. We introduce the idea of *a priori* optimization as a strategy competitive to the strategy of re-optimization, under which the combinatorial optimization problem is solved optimally for every one of its instances. We consider four problems: the traveling salesman problem (TSP), the minimum spanning tree problem, the vehicle routing problem and the traveling salesman facility location problem. 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 re-optimization 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 of finding practical solutions based on ideas of local optimality.

AMS(MOS) Subject Classification. 05C45, 90C27, 60C05, 05-04

Key words. combinatorial optimization, a priori optimization, probabilistic analysis, heuristics.

Introduction

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 will be explained below. 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, as well. 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 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 solve repeatedly many other instances of the same problem. These other instances are usually just variations of the instance solved originally. Yet, they may be sufficiently different from that original instance to necessitate every time a re-consideration of the entire problem on the part of the analyst.

The most obvious approach in dealing with such cases is to attempt to solve optimally (or near-optimally with a good heuristic) every potential instance of the original problem. Throughout the paper, we call this approach the "re-optimization 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 $NP - hard$, 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 here a different strategy. Rather than re-optimizing 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/variation. Clearly, the natural questions to ask are: What is the measure of "effectiveness" of such an *a priori* solution? Once such a measure has been defined, how does one find the best *a priori* solution? And, how does one update the *a priori* solution for each particular 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 instance S has 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 so as to minimize the expected cost

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

with the summation being 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 re-optimization strategy:

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

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 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 re-optimization strategies under reasonable probabilistic assumptions.

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

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 instance 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 consisting 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 above, ideally we might like to re-optimize the tour for every instance, but in many cases we may not have the

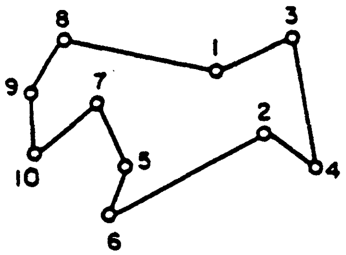

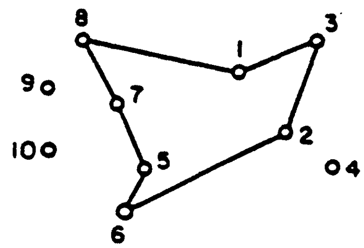

Figure 1: The PTSP methodology

resources to do so or, even if we had them, re-optimization 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 which 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 vertex set $V = \{1, 2, \ldots n\}$. That is, given an *a priori* tour τ, if problem instance $S(\subseteq V)$ will occur with probability $p(S)$ and will require 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

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

with the summation being 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 still are 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 Figure 2, where node 2 is kept on the tree in order to preserve connectness. 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

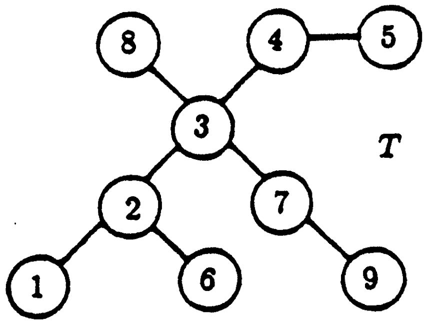

A priori tree

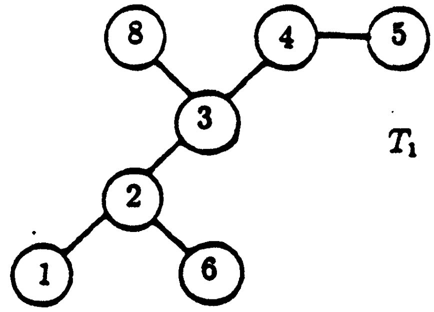

The resulting tree when the nodes 2,7,9 need not be visited.

Figure 2: The PMST methodology

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 requiring 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 similarly to *a* with the sole difference that customers with no demand on a particular instance of

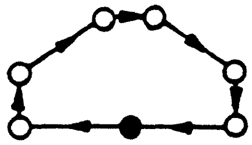

(i) An a priori route through 6 customers (each with a demand of zero or one unit) by a vehicle of capacity 2

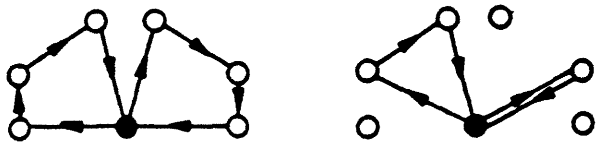

Class A

Class B

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

Figure 3: The PVRP methodology

the vehicle tour are simply skipped. An example of the PVRP under both methods can be seen in Figure 3.

<u>The Traveling Salesman Facility Location Problem</u>

We are given a set of n nodes (customer locations) on a network. Each day a subset S of customers make 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

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

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 more important theorem proofs are only sketchily outlined, with appropriate references given for interested readers. In Section 1 we review briefly the related research and we also 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 very close to the re-optimization 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 has undoubtedly been one of the fastest growing and most exciting areas in mathematical programming. Needless to say, the related scientific literature has been expanding at a very 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 of which 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 which 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 by Steele (1987).

Probabilistic combinatorial optimization problems.

In contrast to their deterministic counterparts, the professional literature on PCOPs to date is very 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 1970's [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 et al. (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 are considering in this paper.

The traveling salesman facility location problem has been considered by Eilon et al. (1971) and Burness and White (1976), where heuristic approaches are proposed. Recently in a series of papers, Berman and Simchi-Levi (1986, 1988a, 1988b) 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 (1989a) 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 as well as its 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 [see for example Colbourn (1987)] the nodes are usually assumed to be always reliable and the type of questions addressed are about the existence of paths among pairs of nodes. In the class of PCOPs the type of questions we are addressing as well as the motivation for their definition are different.

As noted earlier, PCOPs could prove highly 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 non-routing context, PTSP models can also 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 in analyzing such situations.

PVRPs are of course "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 also very common contexts 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 being 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 from 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 which 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 $i = 1, 2, \ldots, n$ 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 re-optimization and *a priori* optimization

In this section we characterize the asymptotic behavior of the re-optimization and the *a priori* strategies for the four problems we have 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 $X^{(n)} = (X_1, \ldots, X_n)$ be n points uniformly and independently distributed in the unit square. Let L^n_{TSP}, L^n_{MST}, $L^n_{STEINER}$, L^n_{VRP} 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 $E[\Sigma^n_{TSP}]$, $E[\Sigma^n_{MST}]$, $E[\Sigma^n_{STEINER}]$, $E[\Sigma^n_{VRP}]$, $E[\Sigma^n_{TSFLP}(i)]$ be the expectation of the TSP, MST, STEINER tree, VRP and TSFLP solutions obtained under the re-optimization 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 $E[L^n_{PTSP}], E[L^n_{PMST}], E[L^n_{PVRPa}], E[L^n_{PVRPb}], E[L^n_{PTSFLP}(i)]$ be the expectation of the a priori strategies, i.e. the expected length of the optimal a priori solution to the PTSP, PMST, PVRP under updating methods a and b, and PTSFLP defined on $X^{(n)}$.

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

Theorem 1

With probability 1 there are constants [Steele (1981)] β_{TSP}, β_{MST}, $\beta_{STEINER}$, such that

$$\lim_{n\to\infty} \frac{L^n_{TSP}}{\sqrt{n}} = \beta_{TSP}, \quad \lim_{n\to\infty} \frac{L^n_{MST}}{\sqrt{n}} = \beta_{MST}, \quad \lim_{n\to\infty} \frac{L^n_{STEINER}}{\sqrt{n}} = \beta_{STEINER}.$$

For the VRP [Haimovitch and Rinnooy Kan (1985)]

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

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

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

We now characterize the expectation of the re-optimization 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. We remark that in the following theorem the expectation is taken over all the possible 2^n instances of the problem and the probability 1 statement refers to the random locations of the points.

Theorem 2 (Bertsimas (1988), Jaillet(1985))

With probability 1

$$\lim_{n\to\infty} \frac{E[\Sigma^n_{TSP}]}{\sqrt{n}} = \beta_{TSP}\sqrt{p}, \quad \lim_{n\to\infty} \frac{E[\Sigma^n_{MST}]}{\sqrt{n}} = \beta_{MST}\sqrt{p},$$

$$\lim_{n\to\infty} \frac{E[\Sigma^n_{STEINER}]}{\sqrt{n}} = \beta_{STEINER}\sqrt{p},$$

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

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

$$\lim_{n\to\infty} \frac{E[\Sigma^n_{TSFLP}(i)]}{\sqrt{n}} = \beta_{TSP}\sqrt{p}, \quad \forall i,$$

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 $|S| \in [\lfloor np(1-\epsilon)\rfloor, \lceil np(1+\epsilon)\rceil]$. 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 re-optimization strategy in appendix A. □

Intuitively theorem 2 means that solutions under the re-optimization strategy behave asymptotically similarly to those of the corresponding combinatorial optimization problems but on np rather than n points. The asymptotic behavior of the VRP re-optimization strategy suggests that the strategy behaves like the TSP re-optimization strategy if the capacity Q is large, a property which 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 re-optimization strategy we observe that asymptotically the location component of the problem is unimportant, since the same asymptotic behavior is observed irrespectively of the location decision i.

We next characterize asymptotically the *a priori* optimization strategies.

<u>Theorem 3 (Bertsimas (1988), Jaillet(1985))</u>

With probability 1

$$\lim_{n\to\infty} \frac{E[L^n_{PTSP}]}{\sqrt{n}} = \beta_{TSP}(p), \quad \lim_{n\to\infty} \frac{E[L^n_{PMST}]}{\sqrt{n}} = \beta_{MST}(p),$$

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

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

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

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

$$\lim_{n\to\infty} \frac{E[L^n_{PTSFLP}(i)]}{\sqrt{n}} = \beta_{TSP}(p), \forall i,$$

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

$$\max\left(\frac{2p}{Q}\sum_{i=1}^{n} d(0,i), L_{TSP}\right) \le E[L_{PVRPa}] \le L_{TSP}(1-\frac{2}{n}-\frac{p}{Q})+2(2+\frac{np}{Q})\frac{\sum_{i=1}^{n} d(0,i)}{n},$$

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 re-optimization 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 very easily updated, while re-optimization strategies require the computation of an optimal solution for every problem instance. Yet, *a priori* strategies behave asymptotically equally well on average with re-optimization strategies on Euclidean problems. In addition, we conjecture that *a priori* and re-optimization strategies have exactly the same asymptotic performance almost surely, i.e. $\beta_{TSP}(p) = \beta_{TSP}\sqrt{p}$ and $\beta_{MST}(p) = \beta_{MST}\sqrt{p}$.

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 re-optimization 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 $(p_1, \ldots, p_n)$ of the probabilities of 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, 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, since, as noted earlier, these three problems are generalizations of well known $NP-complete$ problems. Membership in NP is seen to hold, since given a solution f we can compute $E[L_f]$ in $O(n^2)$. For example for the PTSP if the tour is $\tau = (1, 2, \ldots, n, 1)$ then by looking at the probability of every link being present we can derive (Jaillet (1988)) the following expression:

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

The case of the PMST is more difficult, because the PMST problem is not a generalization of a $NP-complete$ 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 (1989b)):

$$E[L_T] = \sum_{e \in T} c(e) \{1 - \prod_{i \in K_e} (1 - p_i)\} \{1 - \prod_{i \in V - K_e} (1 - p_i)\},$$

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 compute efficiently 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 applied to the PTSP? In order to find useful bounds for the routing problems (PTSP, PVRP) we assume below that the triangle inequality holds. We can then prove the following:

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

$$\frac{L_H}{L_D} \leq c, \quad \text{then} \quad \frac{E[L_H]}{E[L_p]} \leq \frac{c}{p}.$$

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 $E[L_p] \geq pL_p$. Combining these inequalities the result follows. For details see Bertsimas (1988). □

Theorem 5 suggests that if the coverage probability is large then constant guarantee heuristics for the deterministic problem still behave well for the corresponding probabilistic problem. But if $p \to 0$ the bound is not informative and indeed one can find examples with $p \to 0, np \to \infty$ for which $\frac{E[L_D]}{E[L_p]} \to \infty$, 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 $\frac{E[L_D]}{E[L_p]}$ tends to infinity, we can prove the following:

Theorem 6 (Bertsimas (1988))

For the PTSP with triangle inequality

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

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 (1982) 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 the spacefilling curve heuristic produces a tour $\tau_{SF}(S)$ if we run the heuristic on the instance S. Consider now the tour τ_{SF} produced by the heuristic on the original instance 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

$$\frac{E[L_{\tau_{SF}}]}{E[L_{\tau_p}]} \leq \frac{E[L_{\tau_{SF}}]}{E[\Sigma_{TSP}]} = O(\log n). \tag{6}$$

$$\frac{E[L_{\tau_{SF}}]}{E[L_{PVRPb}]} \leq \frac{E[L_{\tau_{SF}}]}{E[\Sigma_{VRP}]} = Q + O(\log n).$$

Sketch of the Proof:

In Platzman and Bartholdi (1983) 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 $\tau_{SF}(S)$ is the tour produced by the PTSP strategy at instance S then

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

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 re-optimization strategies from a worst-case perspective. For the Euclidean PTSP, since $E[L_{\tau_p}] \leq E[L_{\tau_{SF}}]$,

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

Platzman and Bartholdi (1983) conjecture that the spacefilling curve heuristic is a constant-guarantee heuristic. Unfortunately Bertsimas and Grigni (1989) refuted 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 of the locations of the customers use the spacefilling curve heuristic to find the *a priori* tour τ_{SF}.

2. Compute $h(i, \tau_{SF})$ with a vector of probabilities $(p_1, \ldots, p_{i-1}, 1, p_{i+1}, \ldots, p_n)$, for every node i.

3. Select the point i_{SF} that minimizes $h(i, \tau_{SF})$. 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 (1989a))

If $p_i = \Omega(1/\log n)$ for all i, then

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

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 $h(i, \tau_{SF})$ 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 $h(i, \tau_{SF})$ and $h(i+1, \tau_{SF})$ is due to the corresponding probability vectors, which differ solely in the ith and $(i+1)th$ position, we can calculate $h(i+1, \tau_{SF})$ in $O(n)$ given $h(i, \tau_{SF})$, 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 $j = 1, \ldots, n$. Thus, we can compute $h(1, \tau_{SF})$ in $O(n^2)$ and then compute $h(i+1, \tau_{SF})$ from $h(i, \tau_{SF})$ 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 5/2 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 characterized sharply its asymptotic behavior, so that for random problems we know that the expected length of the optimal solution would be close to $\beta_{TSP}\sqrt{np}$. This can be used as a "benchmark" to compare the performance of various heuristics.

In our numerical experiments we have obtained near-optimal solutions to Euclidean PTSPs by means of two different types of heuristics. The first of them 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, $E[L_{\tau_{SF}}]$, 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 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.

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 equation (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 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 we have used 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 converged 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 become excessive. At that point we were 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 in order to minimize the effects of statistical fluctuations.

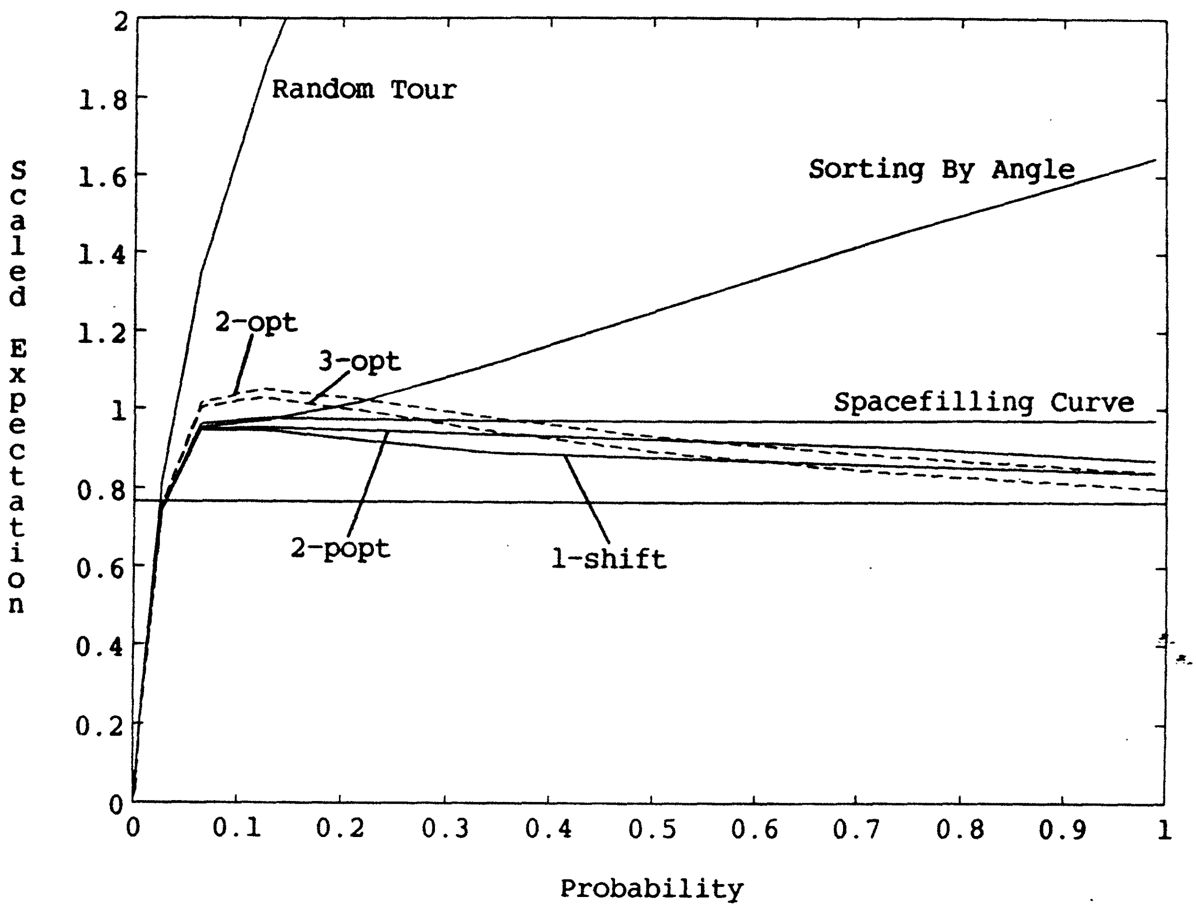

Figure 4: A summary of results for several PTSP heuristics on 100-node problems scaled by $\sqrt{np}$. Solutions obtained via the 2-opt and 3-opt TSP algorithms (dashed lines) are shown for comparison. The horizontal line shows the value of $\beta_{TSP} \approx .765$. The heuristics are 1) random tour, 2) angular sorting, 3) spacefilling curve, 4) 2-p-opt and 5) 1-shift.

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 re-optimization, 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, etc.

It was shown that for all problems defined here a priori and re-optimization 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 queueing phenomena arise which are interesting in themselves. A step in this direction is taken in Bertsimas and van Ryzin (1989), 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 very interesting aspects introduced by it and finally the excellent potential for deriving new results and solution procedures and for applying them to many important contexts.

Appendix A: Proof of theorem 2 for the PTSP

Let W be the number of nodes present and

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

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

$$E[\Sigma^n_{TSP}(X^{(n)})] = \sum_{k=0}^{n}\binom{n}{k}p^k(1-p)^{n-k}h_k = \sum_{k=0}^{n} Pr\{W=k\}h_k.$$

Fix $\epsilon > 0$. Then

$$E[\Sigma^n_{TSP}(X^{(n)})] = \sum_{k=0}^{\lfloor np(1-\epsilon)\rfloor - 1} Pr\{W=k\}h_k + \sum_{k=\lceil np(1+\epsilon)\rceil+1}^{n} Pr\{W=k\}h_k +$$

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

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

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

From the Chernoff bound we have

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

The contribution of the first two terms is then

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

For $\lfloor np(1-\epsilon)\rfloor \le k \le \lceil np(1+\epsilon)\rceil$ we apply theorem 1 and obtain that with probability 1

$$\forall\epsilon>0, \ni k_\epsilon : \forall S, \quad \text{with} \quad |S| = k \ge k_\epsilon \quad -\epsilon \le \frac{L_{TSP}(X^{(n)};S)}{\sqrt{k}} - \beta_{TSP} \le \epsilon \Rightarrow$$

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

In addition,

$$\sum_{k=\lfloor np(1-\epsilon)\rfloor}^{\lceil np(1+\epsilon)\rceil} Pr\{W = k\} = Pr\{|W - np| \leq np\epsilon\} > 1 - 2\delta^n.$$

Therefore,

$$(\beta_{TSP} - \epsilon)(1 - 2\delta^n) < \sum_{k=\lfloor np(1-\epsilon)\rfloor}^{\lceil np(1+\epsilon)\rceil} Pr\{W = k\}\frac{h_k}{\sqrt{k}} < (\beta_{TSP} + \epsilon)1,$$

from which

$$(\beta_{TSP}-\epsilon)(1-2\delta^n)\sqrt{p(1-\epsilon)} < \sum_{k=\lfloor np(1-\epsilon)\rfloor}^{\lceil np(1+\epsilon)\rceil} Pr\{W = k\}h_k/\sqrt{n} < (\beta_{TSP}+\epsilon)\sqrt{p(1+\epsilon)}.$$

Combining the above bounds, we find that almost surely $\forall\epsilon > 0$, $\forall n \geq \frac{k_\epsilon}{p(1-\epsilon)}$

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

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 $E[L^n_{PTSP}] = E[L^n_{\tau_p}]$. We will first prove that with probability 1 $\lim_{n\to\infty} E[L^n_{\tau_p}(X^{(n)})]/\sqrt{n}$ exists. In order to do this we check whether the functional

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

is a subadditive monotone Euclidean functional [Steele (1981)].

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

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

   and it is linear, i.e.

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

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, i.e.

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

4. $f(X^{(n)})$ is subadditive, i.e. if $Q_i, i = 1, \ldots m^2$ 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 in proving this property. Consider the following algorithm:

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

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).

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

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

$$E[L_\tau] = \sum_{i=1}^{m^2} f_1(X^{(n)} \cap rQ_i) + L_{\tau^*},$$

where $f_1(X^{(n)} \cap rQ_i)$ 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 probability p of being present. If we turn a "black" node into a "white" node (a node which has 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 expected length not smaller than $E[L_{\tau_p}]$, since by definition τ_p is the optimal PTSP. Then

$$E[L_{\tau_p}] \leq \sum_{i=1}^{m^2} f_1(X^{(n)} \cap rQ_i) + L_{\tau^*}.$$

It is well known [Larson and Odoni (1981)] 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 $f_1(X^{(n)} \cap rQ_i)$ with $f(X^{(n)} \cap rQ_i)$ or equivalently $E[L_{\tau_i}]$ with $E[L_{\tau_i}|$a node is black]. Without loss of generality assume that the optimal PTSP through the points $X^{(n)} \cap rQ_i$ is $(x_1, x_2, \ldots, x_{k_i}, x_1)$ 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 $f_1(X^{(n)} \cap rQ_i) \leq f(X^{(n)} \cap rQ_i) + 2(1-p)\max_{1<j\leq k_i} |x_1 - x_j|$. Since $\max_{1<j\leq k_i} |x_1 - x_j| \leq \sqrt{2}r/m$ we finally get

$$f_1(X^{(n)} \cap rQ_i) \leq f(X^{(n)} \cap rQ_i) + 2(1-p)\sqrt{2}r/m.$$

Therefore, we can conclude that

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

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 $\beta_{TSP}(p)$ such that with probability 1

$$\lim_{n\to\infty} E[L^n_{\tau_p}(X^{(n)})]/\sqrt{n} = \beta_{TSP}(p).$$

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

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

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

References

[1] Bartholdi J.J. and Platzman L.K., (1982), "An $O(N \log N)$ Planar Traveling Salesman Heuristic Based on Spacefilling Curves", Oper. Res. Let., 1, 121-125.

[2] Bartholdi J.J., Platzman L.K., Collins R. Lee and Warden W.H., (1983), "A Minimal Technology Routing System for Meals on Wheels", Interfaces, 13, 1-8.

[3] Beardwood J., Halton J. and Hammersley J., (1959), "The Shortest Path Through Many Points", Proc. Camb. Phil. Soc., 55, 299-327.

[4] Berman O. and Simchi-Levi D., (1986), "Minisum Location of a Traveling Salesman", Networks, 16, 329-354.

[5] Berman O. and Simchi-Levi D., (1988a), "Finding the Optimal *a priori* Tour and Location of a Traveling Salesman with Non-Homogeneous Customers", Transp. Science, 22, 148-154.

[6] Berman O. and Simchi-Levi D., (1988b), "A Heuristic Algorithm for the Traveling Salesman Location Problem on Networks", Oper. Res., 36, 478-484.

[7] Bertsimas D., (1988), "Probabilistic Combinatorial Optimization Problems", Ph.D thesis, Operations Research Center, MIT, Technical report 194.

[8] Bertsimas D., (1989a), "On probabilistic traveling salesman facility location problems", to appear in Transportation Science.

[9] Bertsimas D., (1989b), "The probabilistic minimum spanning tree problem", to appear in Networks.

[10] Bertsimas D. and Grigni M. (1989), "Worst case examples for the spacefilling curve heuristic for the Euclidean traveling salesman problem", to appear in Operations Research Letters.

[11] Bertsimas D. and van Ryzin G. (1989), "The Dynamic Traveling Repairman Problem", Sloan School of Management, MIT, working paper 3036-89, sumbitted for publication to Operations Research.

[12] Bodin L., Golden B., Assad A. and Ball M., (1983), "Routing and Scheduling of Vehicles and Crews, the State of the Art", Comp. Oper. Res., 10, 69-211.

[13] Burness R. and White J., (1976), "The Traveling Salesman Location Problem", Transp. Science, 10, 348-360.

[14] Dror M. and Trudeau P., (1986), "Stochastic Vehicle Routing with Modified Savings Algorithm", Eur. Jour. Oper. Res., 23, 228-235.

[15] Frieze A.M., (1985),"On the Value of a Minimum Spanning Tree Problem", Discr. Appl. Math., 10, 47-56.

[16] Garey M. and Johnson D., (1979), Computers and Intractability: A Guide to the Theory of NP-Completeness, Freeman, San Francisco.

[17] Golden G. and Assad A., (1988), Vehicle Routing; Methods and Studies, North Holland, Amsterdam.

[18] Haimovitch M. and A. Rinnooy Kan, (1985), "Bounds and Heuristics for Capacitated Routing Problems", Math. Oper. Res., 10, 527-542.

[19] Jaillet P., (1985), "Probabilistic Traveling Salesman Problems", (Ph.D. Thesis) Technical Report No. 185, Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA.

[20] Jaillet P., (1988), "A priori Solution of a Traveling Salesman Problem in which a Random Subset of the Customers are Visited", Oper. Res., 36 (6).

[21] Jaillet P. and Odoni A., (1988), The Probabilistic Vehicle Routing Problem, in Vehicle Routing; Methods and Studies, edited by B.L. Golden and A.A. Assad, North Holland, Amsterdam.

[22] Karp R., (1977), "Probabilistic Analysis of Partitioning Algorithms for the Traveling Salesman in the Plane", Math. Oper. Res., 2, 209-224.

[23] Karp R., (1979), "A Patching Algorithm for the Non-symmetric Traveling Salesman Problem", SIAM Jour. Comp., 8, 561-573.

[24] Laporte G. and Louveau F., (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, Montreal, Canada.

[25] Laporte G., Louveau F. and Mercure H., (1987), "Models and Exact Solutions for a Class of Stochastic Location-Routing Problems", G-87-14, Ecole des Hautes Etudes Commerciale, University of Montreal, Montreal, Canada.

[26] Larson R. and Odoni A., (1981), Urban Operations Research, Prentice-Hall, Englewood Cliffs, NJ.

[27] Lawler E.L., Lenstra J.K., A. Rinnooy Kan and Shmoys D.B. (eds.), (1985), The Traveling Salesman Problem: A Guided Tour of Combinatorial Optimization, Wiley, Chichester.

[28] Lin S., (1965),"Computer Solutions of the Traveling Salesman Problem", Bell System Tech. J., 2245-2269.

[29] Platzman L.K. and Bartholdi J.J., (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, GA.

[30] Simchi-Levi D. and Berman O., (1988), "Heuristics and Bounds for the Traveling Salesman Location Problem on the Plane", Oper. Res. Let., 5, 243-248.

[31] Steele J.M., (1981), "Subadditive Euclidean Functionals and Nonlinear Growth in Geometric Probability, Ann. Prob., 9, 365-376.

[32] Steele J.M., (1987), "On Frieze's $\zeta(3)$ Limit for Lengths of Minimal Spanning Trees", Disc. Appl. Math., 18, 99-103.

[33] Stewart W. and Golden B., (1983), "Stochastic Vehicle Routing: A Comprehensive Approach", Eur. Jour. Oper. Res., 14, 371-385.

[34] Tillman F., (1969), "The Multiple Terminal Delivery Problem with Probabilistic Demands", Transp. Science, 3, 192-204.

M.I.T. Sloan School of Management
Working Paper Series

Papers by Dimitris Bertsimas
Assistant Professor of Management Science

Paper #	Date	Title/Author(s)
3099	12/89	"Transient and Busy Period Analysis of the GI/G/1 Queue: Part II, Solution as a Hilbert Problem," Bertsimas, D, Keilson, J., Nakazato, D., and Zhang, H.
3098	12/89	"Transient and Busy Period Analysis of the GI/G/1 Queue: Part I, The Method of Stages," Bertsimas, D., and Nakazato, D.
3059	8/89	"A Priori Optimization," Bertsimas, D., and Odoni, A.
3058	8/89	"An Asymptotic Determination of the Minimum Spanning Tree and Minimum Matching Consonants In Geometrical Probability," Bertsimas, D., and Van Ryzin, G.
3036	6/89	"The Dynamic Traveling Repairman Problem," Bertsimas, D., and Ryzin, G.
2105		"On the Spacefilling Curve Heuristic for the Euclidean Traveling Salesman Problem," Bertsimas, D.,and Grigni, M.
2104		"Probabilistic Analysis of the 1-Tree Relaxation for the Euclidean Traveling Salesman Problem," Goemans, M., and Bertsimas, D.
2068	9/88	"Two Taveling Salesman Facility Location Problems," Bertsimas, D.
2067	9/88	"The Probabilistic Vehicle Routing Problem," Bertsimas, D.
2057	8/88	"The Probabilistic Minimum Spanning Tree Problem Part I: Complexity and Combinatorial Properties," Bertsimas, D.

Please Use Form On Reverse When Ordering Papers