アルゴリズム 40−11
（1994．7．22）

巡回カメラマン問題とその自動光学的検査への応用

岩野和生* Prabhakar Raghavan** 玉木久夫**

* 日本 IBM 東京基礎研究所
** 米国 IBM ワトソン研究所

巡回カメラマン問題とは、平面上に互いに素な有限個の領域が与えられた時、それらの全ての領域の単位正方形による集合被覆を考え、それらの中心点をめぐる巡回路の経路長を最小に問題である。つまり、集合被覆と巡回セールスマン問題を組み合わせた問題になっている。この問題は、プリント基盤の自動光学的検査から派生したものである。

当論文では、この問題に対する定数比率を持つ近似アルゴリズムを提唱する。また、各領域が異なる長方形の場合や、被覆に用いる正方形のサイズが可変だが、各領域毎に被覆に用いる正方形の最大サイズが決められている変種も考察する。最後の問題は、プリント基盤の自動光学的検査におけるズーム機能に相当する。

The Traveling Cameraman Problem, with Applications to Automatic Optical Inspection

Kazuo Iwano* Prabhakar Raghavan** Hisao Tamaki**

* Tokyo Research Laboratory, IBM Japan, Yamato, Kanagawa, Japan
** T.J. Watson Research Laboratory, IBM U.S., Yorktown, U.S.A.

We are given a finite set of disjoint regions in the plane. We wish to cover all the regions by unit squares, and compute a path that visits the centers of all the unit squares in the cover. Our objective is to minimize the length of this path. The problem arises in the automatic optical inspection of printed circuit boards and other assemblies.

We show that a natural heuristic yields a path of length at most a constant times optimal, whenever the problem of covering the regions by the minimum number of unit squares can be approximated to within a constant. We consider a generalization in which the regions may be covered by squares of different sizes, and for each input region we are given an upper bound on the size of the permissible square. This corresponds to a board inspection problem in which the camera may "zoom", and some parts are to be inspected with greater resolution than others. We show that a simple extension of our heuristic is provably good in this case as well.

1 Introduction

In the automatic optical inspection of printed circuit boards [5, 6, 7], a camera is positioned over the board, and is free to move in a plane parallel to the board. As it moves, it can take photographs of the board at various positions, and compares each photograph to a corresponding "master photograph". By taking a number of such photographs and comparing them to the masters, the board can be inspected for defects, violations of design rules [11], and mounting and soldering condition of components on the board [5, 7]. An important consideration in this automated process is the time taken to perform the entire inspection sequence. This time is proportional to the total distance moved by the camera, as well as the number of photographs taken (since each photograph takes a positive amount of time). In this paper we consider this problem, together with a generalization in which some parts may have to be photographed with greater resolution than others (i.e., covered by smaller squares). Due to lack of the space, we omit almost of all proofs in this paper and desribe them in the full version of this paper [9].

1.1 The Model

The input to the *traveling cameraman problem* (TCP) is a set of n disjoint regions in the plane, each of which we call a *feature*. The output is a set of axis-parallel unit squares whose union contains every feature, and a path through the centers of these squares. Our objective will be to minimize the length of this path; in Section 4 we discuss the more general objective function that takes into account the number of squares used in the solution. We use the L_∞metric to measure the length of the path. This corresponds to the fact that the camera is usually moved by two orthogonal drives and thus the greater of the x– and y–distances determines the time required for the move. Although some of our analyses are specialized to L_∞distance measure, our results also hold (with different constants) for other distance metrics.

The *traveling cameraman problem with zoom* (TCPZ) also has n disjoint regions (the features) in the input; in addition, there is a positive real s_i associated with the ith region, for all i. The output is a set of squares, such that the ith region is covered by the union of squares of side at most s_i. In addition, we are to construct a path through these squares, and we again wish to minimize its length.

1.2 Our Results

A trivial modification of the *NP*-completeness of the L_∞ TSP [12] yields:

Theorem 1 *The TCP is NP-complete.*

In Section 2, we consider the following natural heuristic for the TCP: we first find a minimum cover of all the features by unit squares, and then find a traveling salesman path through the centers of the squares in the cover. Although both the covering problem and the traveling salesman problem are themselves hard to solve exactly, we rely on the fact that both have very good approximation algorithms (especially in the cases of practical interest, where the features are axis-parallel rectangles).

We first show that our heuristic is not necessarily optimal (even if we could solve the covering and TSP problems exactly); in fact, the heuristic may produce a solution whose cost is three times that of the optimal solution, even if all the features are points in the plane. We then show that it always produces a solution that is at most 6 times optimal plus a small constant (Theorem 7), provided that the associated covering and TSP problems can be solved exactly. The crux of our proof is to establish Lemma 5, a geometric lemma with a simple statement whose proof turns out to be surprisingly difficult; this proof is given in the full paper. By adapting an algorithm due to Hochbaum and Maass [8], we show that the covering problem can be approximated to within a factor of $1+\epsilon$ provided all the features are axis-parallel rectilinear regions. When combined with the fact that the TSP in the plane can be approximated to within a constant factor [3], we have a constant-factor approximation for the TCP for axis-parallel rectilinear features.

In Section 3 we turn to the TCPZ. Let S and s be the largest and the smallest values among the s_i in the input. We show that a simple modification of our heuristic always produces a solution within $O(\log(S/s))$ of the optimal. Extensions are discussed in Section 4.

1.3 Related Previous Work

Arkin and Hassin [1] study the *geometric covering salesman problem:* a variant of the TSP in which each city is a region in the plane, and the tour is required to touch at least one point of every region. They give approximation algorithms for several cases of this problem, most notably the case when every region is a line segment parallel to one of k fixed directions. The geometric covering salesman problem and the TCP are generalizations of a common problem in two different directions. If we assume that each city in the geometric covering problem is an axis-parallel unit square then this problem is equivalent to the special case of the TCP in which every feature is a point. We note, however, that the cases arising in practice require the study of more general features. Arkin, Fekete, Mitchell and Piatko [2] study related covering tour problems; their focus is on the following covering problem. Given a region in the plane, they wish to compute a tour such that every point in the region is within unit distance of the tour. The TCP may be viewed as a generalization of their *lawnmower problem*, one in which there are several lawns. Current and Schilling [4] consider a graph-theoretic version of this problem: given a directed graph with lengths on the edges, they wish to compute a short path such that every vertex of the graph is within some given distance S from the path. Like us, they solve their problem by first solving a suitable covering problem, and then computing a short tour through the vertices in the cover.

Although these previous works touch upon our problem indirectly, none of them directly deals with our problem. To our knowledge, there has been no previous algorithmic work related to the TCPZ.

2 The Basic Algorithm and Analysis

We assume a fixed x, y coordinate system according to which the L_∞distance is measured. For a point p on the plane and for a positive real x, let $square_x(p)$ denote the axis-parallel square region of side length x centered at p. For brevity, we write $square(p)$ for $square_1(p)$. Let F be the set of given features and let $n = |F|$. Each feature in F is a simply-connected region on the plane and we assume that these regions are mutually disjoint. We denote by $\hat{F}$ the region that is the union of all features in F. A *covering algorithm* takes F as input and outputs a set of points P such that $\hat{F} \subseteq \bigcup_{p \in P} square(p)$. We call such a set P a *square covering* of F. A square covering P of F is *minimal* if $square(p)$ for each $p \in P$ is non-empty, i.e. it intersects with some feature in F. We assume, quite reasonably, that the covering algorithms we consider output a minimal covering. We denote by $SC(F)$ the cardinality of the optimal (smallest) square covering of F.

By a *path* we mean a curve on the plane. Let *dist*(p, q) denote the L_∞distance between two points p and q, i.e. the greater of $|x_p - x_q|$ and $|y_p - y_q|$ where (x_p, y_p) and (x_q, y_q) are the coordinates of p and q respectively. The length of a path C, denoted by *length*(C), is defined as the integral of the local distance over the entire path. A more formal definition is given in the full paper.

A *TSP algorithm* inputs a set of points P and returns a path that visits all points of P, which we call a *TSP path* for P. We denote by $TSP(P)$ the length of the optimal (shortest) TSP path for P. A *TCP path* for a set of features F is a path such that the set of points on the path is a square covering of F. Given a set of features F, let $TCP(F)$ denote the length of the shortest TCP path for F.

Our algorithm for TCP depends on a covering algorithm and a TSP algorithm as subroutines.

Algorithm : Apply a covering algorithm to F and let the result be P. Then apply a TSP algorithm and obtain a TSP path for P. Output this path together with P.

It is clear that the result of this algorithm is a TCP path for F, provided that the covering algorithm and the TSP algorithm we use are correct. We want to show that if the covering and TSP algorithms give good approximations to the optimal solutions of their respective subproblems then our algorithm gives a good approximation to the optimal solution of the entire problem. We first give a lower bound on such a performance guarantee of our algorithm and then give an upper bound.

2.1 A lower bound on the approximation ratio

In this subsection, we show that there is an instance of TCP such that our algorithm gives a path whose length is three times that of the optimal solution, even if the covering and the TSP subroutines give exact optimal solutions.

Theorem 2 *For every integer $L \geq 2$, there is a set of features F such that*

1. $TCP(F) = L$,
2. *there is a unique optimal square covering P of F, and*
3. $TSP(P) = 3L$.

Even if we restrict the features to be points, it is easy to see that the above lower bound still holds to arbitrary precision:

Corollary 3 *For every positive integer L and positive ϵ, there is a set of points F with $n = |F| = O(L)$ such that*

1. $TCP(F) = L$ *and*
2. *for any optimal square covering P of F, $TSP(P) \geq 3L - \epsilon$.*

2.2 An upper bound

Now we turn to an upper bound on the approximation ratio of our algorithm. We begin by bounding the size of the optimal covering in terms of the length of the optimal TCP path.

Lemma 4 *For any set F of features, we have*

$$SC(F) \leq 3\lceil TCP(F)\rceil + 1.$$

Note that this lemma is tight in view of the lower bound example in the proof of Theorem 2.

Once we have Lemma 4, it is not difficult to obtain *some* constant upper bound on the approximation ratio of our algorithm. The following geometric lemma arises in our effort to get a smaller constant. Given a path C and positive d, let $locus_d(C)$ denote the region defined by $locus_d(C) = \{p \mid \exists q \in C : dist(p,q) \leq d\}$. Note that $locus_d(C)$ may contain *holes*, regions enclosed by but not belonging to $locus_d(C)$.

Lemma 5 *Let C be a path. Then there exists a closed path C' with length at most $2length(C) + 8d$ such that C' contains all the boundaries of $locus_d(C)$, including those around holes.*

The proof of this lemma is given in the full paper. The proof would be considerably easier if we could assume there were no holes. With holes, the proof would still be easy if we only had to bound the total length of the boundaries. The complexity of the proof arises from the necessity to thread all these boundaries with a path of the allowed length.

Lemma 6 *Let F be an arbitrary set of features and let P be an arbitrary minimal square covering of F. Then, $TSP(P) \leq 3TCP(F) + |P| + 4.5$.*

Theorem 7 *Suppose our covering subroutine gives a solution whose cardinality is at most c_1 times that of the optimal and our TSP subroutine gives a path with length c_2 times that of the optimal. Then, our algorithm gives a TCP path with length less than $c_2(3c_1+3)L + c_2(4c_1+4.5)$, where L is the length of an optimal TCP path.*

The following corollary follows from the $(1+\epsilon)$–approximate covering algorithm given in the next subsection, and the 3/2–approximate TSP algorithm of Christofides [3].

Corollary 8 *Suppose either each feature in F is a point or each feature in F is an axis-parallel rectilinear polygon. Then, for each fixed positive ϵ, there is a TCP algorithm that produces a path of length at most $(9+\epsilon)$ times the optimal and runs in time polynomial in $n = |F|$.*

2.3 Approximate covering of axis-parallel rectilinear polygons

The goal of this section is to extend the polynomial time approximate covering algorithm due to Hochbaum and Maass [8] to the case where the features are rectilinear polygons (rather than points).

Theorem 9 *[8] Let F be a set of points on the plane, with $|F| = n$. For any fixed positive ϵ, there exists an algorithm that runs in time polynomial in n and gives a square covering P of F such that $|P| \leq (1+\epsilon)SC(F)$.*

Their algorithm is based on the following elegant lemma. Let F be the given set of points. Fix a positive integer l and, for each $1 \leq k \leq l$ and each integer i, let $F_{k,i}$ denote the subset of F consisting of those points between two lines $x = il + k$ and $x = (i+1)l + k$. Let $SC_k(F)$ denote $\sum_i SC(F_{k,i})$. Note that $SC_k(F)$ corresponds to the square covering obtained by taking the optimal covering of $F_{k,i}$ for each i independently.

Lemma 10 *(Shifting lemma, [8]) For some k, $1 \leq k \leq l$, $SC_k(F) \leq (1 + 1/l)SC(F)$.*

Applying this lemma twice, horizontally and then vertically, Theorem 9 reduces to the following lemma; we choose l so that $(1 + 1/l)^2 \leq 1 + \epsilon$.

Lemma 11 *[8] Let l be a fixed positive integer and let F be an arbitrary set of n points in an axis-parallel square of side length l. Then an exact optimal square covering of F can be obtained in time polynomial in n.*

The shifting lemma does not depend on the number of points to be covered, or even the finiteness of the number of points. Therefore, we can use this lemma when F is a set of general features instead of points, simply by replacing F by $\hat{F}$.

Thus, we only have to provide an extension to Lemma 11, namely:

Lemma 12 *Let l be a fixed positive integer and let F be an arbitrary set of axis-parallel rectilinear polygonal regions contained in an axis-parallel square of side length l. Let m be the total number of line segments forming the boundaries of the features in F. Then an exact optimal square covering of F can be obtained in time polynomial in m.*

Now we have the desired extension of Theorem 9.

Theorem 13 *Let F be a set of features, where each feature is an axis-parallel rectilinear region, such that the number of line segments defining the features of F is m and such that all the features of F are contained in a region with area M. For any fixed positive ϵ, there exists an algorithm that runs in time polynomial in $m + M$ and gives a square covering P of F such that $|P| \leq (1+\epsilon)SC(F)$.*

3 The Zoom Case

As before, let F be the set of given features and let $n = |F|$. In addition, for the TCPZ, we are given a positive real s_i associated with the ith feature $f_i \in F$. We refer to s_i as the *frame size* of f_i. We are required to cover f_i by a square of side at most s_i.

Algorithm : 1. First, we round each s_i down to the nearest power of 2: thus s_i is rounded down to $2^{\lfloor \log_2 s_i \rfloor}$. Let $r_1 \leq \cdots \leq r_k$ be the distinct values (powers of 2) that result.

2. For $i = 1, \ldots, k$, we use a covering algorithm to cover all features f_i whose frame size was rounded to r_i, using squares of side r_i. (while covering with squares of side r_i, we omit from consideration those features covered by squares of size r_j for $j < i$),
3. We then use a TSP algorithm to connect the centers of the squares that result.

Informally, our algorithm first covers the features that are to be covered by the smallest squares, and then proceeds to larger and larger squares. Some features may be covered by a square smaller than necessary, and this is acceptable. Moreover, in the application of automatic optical inspection, it is preferable to use a smaller square. As before, we apply a TSP routine to join the centers of the resultant squares. We bound the quality of this approximation. Let $R = r_k/r_1$. Let $TCPZ(F)$ denote the length of the optimal TCPZ path through the features F.

Theorem 14 *There is a positive constant c such that the length of the TCPZ path produced by the above algorithm is at most* $(c \log R)(TCPZ(F) + 1)$.

For $i = 1, \ldots, k$, let F_i be the set of features covered in the ith iteration of step 2 of the algorithm. Let P_i be the set of centers of squares of side r_i in a minimum cover of F_i. By essentially repeating the argument used to prove Theorem 7, we have:

Lemma 15 *For* $1 \le i \le k$, $TSP(P_i) \le 6TCPZ(F) + 8.5$.

Proof of Theorem 14: If we had a TSP algorithm that would find the optimal TSP path, our algorithm would find a solution of length $TSP(\cup_{i=1}^{k} P_i)$. By the triangle inequality, $TSP(\cup_{i=1}^{k} P_i) \le \sum_{i=1}^{k} TSP(P_i)$. By Lemma 15, $\sum_{i=1}^{k} TSP(P_i) \le k(6TCPZ(F) + 8.5)$. Noting that $k \le \log R = \log r_k/r_1$ now yields the result. □

4 Discussion

We begin by considering the case where we are interested in minimizing not only the path length but also the number of photographs taken. Thus a TCP algorithm in this case output, in addition to the path, a set of points on the path that forms a square cover of the given features. Our objective function here is the sum of the path length and the size of the square cover. Let us call this latter component of the cost the *photo cost.* We note that if we were to use an optimal cover in our algorithm of Section 2, our photo cost is no more than that of the optimal solution to the TCP. If we used an approximate covering algorithm that delivered a cover of size at most c_1 times the optimal, it follows that our photo cost is no more than c_1 times that of the optimal. Therefore, Theorem 7 and all its corollaries apply for any non-negative photo cost.

In some inspection problems, it may be required that each feature is completely contained in some single photograph. Out results apply to this variant too; in fact, the extension of Hochbaum-Maass theorem (Theorem 9) is rather trivial in this case.

In practice, we suspect that the covering heuristic of Section 2.3 would be too computationally intensive, and a variant of the greedy heuristic [10] might be preferable. In this case, we would lose the constant bound on the approximation ratio but the performance may turn out to be fairly good in practice.

Acknowledgment

We thank Don Coppersmith for many useful discussions leading to the proof of Lemma 5.

参考文献

[1] E. Arkin and R. Hassin. Approximation algorithms for the geometric covering salesman problem. Technical Report 968, School of Operations Research and Industrial Engineering, Cornell University, 1991. To appear: *Discrete Applied Math.*

[2] E. Arkin, S. Fekete, J. Mitchell and C. Piatko. Optimal Covering Tour Problems. Manuscript 1993.

[3] N. Christofides. Worst-case Analysis of a New Heuristic for the Traveling Salesman Problem. Technical Report, GSIA, Carnegie-Mellon Univ., 1976.

[4] J. Current and D. Schilling. The Covering Salesman Problem. *Transportation Science* 23, 208–213, 1989.

[5] P. Forte et al. Automatic Inspection of Electronic Surface Mount Assemblies. *Proceedings of SPIE - International Society of Optical Engineering, Intelligent Robots and Computer Vision XII: Algorithms and Techniques*, vol. 2055, 50–56, 1993.

[6] S. Hata. Vision Systems for PCB Manufacturing in Japan. *IECON '90, 16th Annual Conference of IEEE Industrial Electronics Society*, vol. 1, 792–797, 1990.

[7] M.E. Hernandes, J.R. Villalobos, and W.C. Johnson. Sequential Computer Algorithms for Printed Circuit Board Inspection. *Proceedings of SPIE - International Society of Optical Engineering, Intelligent Robots and Computer Vision XII: Active Vision and 3D Methods*, vol. 2056, 438–449, 1993.

[8] D.S. Hochbaum and W. Maass. Approximation Schemes for Covering and Packing Problems in Image Processing and VLSI. *Journal of the ACM*, 32:130–136, January 1985.

[9] K. Iwano, P. Raghavan, and H. Tamaki. The Traveling Cameraman Problem with Applications to Optical Inspection. IBM Research Report RC 19474, March, 1994.

[10] L. Lovász. On the ratio of optimal integral and fractional covers. *Discrete Math.*, 13:383–390, 1975.

[11] S.H. Oguz and L. Onural. An Automated System for Design Rule Based Visual Inspection of Printed Circuit Boards. *Proceedings of the 1991 IEEE International Conference on Robotics and Applications, vol. 3*, 2696–2701, 1991.

[12] C.H. Papadimitriou. The Euclidean Traveling Salesman Problem is NP-Complete. *Theoretical Computer Science*, 4:237–244, 1977.