アルゴリズム 53−11
（1996. 9. 13）

平面点集合の k巡回路被覆問題：kが定数の場合の多項式時間近似スキーム

浅野哲夫　大阪電気通信大学
加藤直樹　神戸商科大学
玉木久夫　徳山豪　ＩＢＭ東京基礎研究所

あらまし：平面上に点集合 Sが与えられた時、固定された原点と Sの高々k点を通る巡回路を Sの k巡回路と呼ぶ。我々の目的は、Sのすべての点を k巡回路の集合で被覆しその総長を最小化することである。kが定数のとき、この問題に対して Haimovich と Rinnooy Kan による多項式時間近似スキームが知られている。本稿では別の考え方に基づく多項式時間近似スキームを与える。この結果は d が定数である限り d 次元空間における問題に拡張できる。対照的に、点集合が一般的な計量空間に置かれた場合にはこの問題は ＡＰＸ完全であることを示す。

キーワード：車両ルーティング、集合被覆、近似アルゴリズム

Covering points in the plane by k-tours: a polynomial time approximation scheme for fixed k

Tetsuo Asano
Dept. of Engr. Informatics
Osaka Electro-Communication University

Naoki Katoh
Dept. of Management Science
Kobe University of Commerce

Hisao Tamaki　Takeshi Tokuyama
IBM Tokyo Research Laboratory

Abstaract: Let S be a set of points in the plane. A *k-tour* through S is a tour in the plane that starts and ends at the fixed origin and visits at most k points of S. Our goal is to cover all the points of S by k-tours while minimizing the total length of the tours. When k is fixed, a polynomial time approximation scheme (PTAS) due to Haimovich and Rinnooy Kan is known. We give an alternative PTAS for fixed k. This result extends to the d-dimensional space, as long as d is a constant. In contrast, we show that the non-geometric version of this problem, in which the points are placed in a general metric space, is APX-complete for any fixed k.

Key words: vehicle routing, set cover, approximation

1 Introduction

Suppose a positive integer k, a set of n nodes V, a special node s called the *source* not in V, and the distance $w(u, v)$ between each pair in $V \cup \{s\}$ are given. We wish to cover all the nodes of V by tours, each of which contains s and at most k nodes of V, so that the total length of the tours is minimized. We call such tours *k-tours* and this problem the *k-tour cover problem.* Clearly, the n-tour cover problem is nothing but the celebrated traveling salesman problem (TSP).

A natural interpretation of the k-tour cover problem is to view it as a special case of the vehicle routing problem (VRP). VRP has been extensively studied over these thirty years in the field of operations research. (See the following references[Lun, Gol84, GA88, Fis]. In particular, [Lun] collects more than 300 papers related to VRP.) According to the constraints and objectives that arise in real applications, there is a wide variety of problem formulations. Also, many heuristic and exact algorithms have been proposed.

The simplest form of VRP is formulated as follows. A set of customers is given, each with a demand of goods to be delivered from a central depot. The goods are to be delivered by vehicles with an identical capacity. Each vehicle starts from the depot, delivers goods to a subset of the customers, and returns to the depot. The objective is to minimize the delivering cost incurred by vehicles (usually cost is measured by the traveling distance) so as to meet the demands of all customers. If the demands of all the customers are equal, say the unit quantity, and the capacity of each vehicle is k units, then this version of VRP is equivalent to our k-tour cover problem. Although this condition of unit demands may appear an oversimplification, we note that the general demand case can be reduced to the unit demand case if each demand is an integer and is allowed to be broken up into integral pieces to be satisfied by separate tours: we represent a customer with demand i by i customers with a unit demand each. The plausibility of such a model depends on particular applications.

The 2-tour cover problem can be solved easily: since the edge from s to each $v \in V$ must be included in any 2-tour cover of V, the problem reduces to finding a minimum weight matching among the nodes of V. When $k \geq 3$, the k-tour cover problem is NP-hard. Thus, our interest is directed to the approximability of the problem.

In recent years, a considerable amount of work has been done in classifying combinatorial optimization problems with respect to their approximability. (For literature, see a compendium by Crescenzi and Kann[CK].) For a combinatorial optimization problem, an algorithm is said to be an *α-approximation algorithm,* where $\alpha \geq 1$, if the value of the solution produced by the algorithm is always within a factor of α from the optimal solution. A problem *belongs to the class APX,* if it has a polynomial time c-approximation algorithm for some constant c. A problem has a *polynomial time approximation scheme* (PTAS), or *belongs to the class PTAS,* if it has a polynomial time $(1+\epsilon)$-approximation algorithm for every constant $\epsilon > 0$. A problem P is *APX-hard* if every problem Q in APX is PTAS-reducible to P, i.e., there is a reduction from Q to P that turns any PTAS of P into a PTAS of Q. It is *APX-complete* if it is in APX and is APX-hard. The recent developments in the theory of approximability [PY91, ALMSS92, KMSV94, CT94] have established that any APX-hard problem does not admit a PTAS unless P = NP.

An interesting issue in the study of approximability is the role of geometry. For example, the set cover problem, which is believed to be not in APX[LY94, Fei96], has some geometric versions that admit PTAS[HM85], some geometric versions in APX ([Mat93, GL95]), and yet some other geometric versions that are not known to be in APX. Another prominent example highlighted recently is TSP. The metric TSP, i.e., TSP with a distance function satisfying the triangle inequality, has been known to be APX-hard[PY93]. On the other hand, it had been open for a long time whether TSP with Euclidean distance admits a PTAS, until very recently Arora [Aro96] has answered the question in the affirmative (for the 2-dimensional case), showing that geometry indeed helps in approximating TSP.

The goal of this manuscript is to point out a similar phenomenon on the k-tour cover problem, when k is fixed. Haimovich and Rinnooy Kan [HR85] constructed a PTAS for the geometric version while we show that the general metric version is APX-complete. We also give a PTAS based on an approach different from that of the above authors.

It is interesting to note that the approximability question on the k-tour problem is resolved in a similar fashion in both of the opposite extremes: when k is a constant and when $k = n$ (Arora's TSP result). In fact, Arora's technique can be extended to give a PTAS for the k-tour cover problem for $k = \Omega(n)$. The question is open for a general k in between.

2 A PTAS for the Euclidean k-tour cover problem

In this section, we consider the k-tour cover problem where the weight function is the Euclidean distance in the plane. Thus, an instance of the problem is given by a triple (V, s, ψ), where s is the source, V is the set of nodes to be covered, and $\psi : V \cup \{s\} \rightarrow \mathbf{R}^2$ is an embedding of the nodes into the plane. We assume that the source is always mapped to the *origin* of the plane. Distinguishing between a node v and its location $\psi(v)$ in the plane will turn out useful, since our technique forces us to deal heavily with multiple nodes placed on a single point.

The rest of this section is devoted to an alternative proof of the following theorem due to Haimovich and Rinnooy Kan.

Theorem 2.1 *When k is fixed, the Euclidean k-tour cover problem admits a PTAS.*

The constants in the following subsections will in general depend on k as well as on ϵ, the bound on the relative error of the approximate solution with respect to the optimal.

2.1 Overview of the algorithm

Suppose an Euclidean instance (V, s, ψ) of the k-tour problem is given. We normalize the distance so that the smallest enclosing circle of the points of $\psi(V)$ centered at the origin has radius 1. Our goal is to find a k-tour cover of V which approximates the optimal solution within a relative error of ϵ.

Our approach is to first perturb the given embedding ψ into a highly degenerate embedding ψ' so that the number of distinct points in the image $\psi'(V)$ is $O(\log n)$. We keep the perturbation small so that the weight of any k-tour cover after the perturbation closely approximates its weight before the perturbation. Then we exactly solve the k-tour cover problem under this perturbed embedding, using a dynamic programming approach. The running time of this computation is exponential in $m = |\psi'(V)| = O(\log n)$, which is polynomial in n. A key in establishing this time bound is Lemma 2.1 which essentially states that, in any optimal k-tour cover of V under embedding ψ', at most $O(m)$ k-tours are non-trivial, i.e., visit more than one point of $\psi'(V)$.

The following subsections provide the details of the algorithm.

2.2 Perturbing the embedding

We first note that, in constructing the k-tour cover, we may ignore the nodes in V that are mapped into the vicinity of the origin or, more specifically, within distance $\epsilon/2n$ of the origin. This is because such nodes can be covered by 1-tours of weight at most ϵ/n each and thus at most ϵ in total. Since the length of the optimal k-tour cover of V is at least 2, separately dealing with these nodes cannot incur a relative error that is greater than $\epsilon/2$. In the following, we assume that ψ maps each node in V to a point at a distance of at least $\epsilon/2n$ from the origin and try to achieve an approximate solution with a relative error of at most $\epsilon/2$.

Let $\delta = \epsilon/4k$ and $L = 1 + \log_{1-\delta}(\epsilon/2n)$. Draw L circles $C_1, C_2, \ldots, C_L$ centered at the origin, with geometrically reducing radii 1, $(1-\delta), \ldots, (1-\delta)^{L-1} = \epsilon/2n$. Also draw $M = 2\pi/\delta$ rays from the origin so that the angle between two consecutive rays is δ. Let Γ be the set of the "grid points", i.e., the intersections of the above circles and rays. Note that $m = |\Gamma| = LM = O(\log n)$. Define ψ' by setting $\psi'(v)$ to be the point of Γ closest to $\psi(v)$, for each $v \in V$.

The embedding ψ' approximates ψ well in the following sense. For each $v \in V$, let r_v denote the distance of $\psi(v)$ from the origin and d_v the distance between $\psi(v)$ and $\psi'(v)$. Then, from the construction of the grid we have $d_v \leq \delta r_v$ for each $v \in V$. Recall here that our assumptions put $\psi(v)$ for every $v \in V$ in the annulus bounded by C_1 and C_L.

Let Z be an arbitrary k-tour cover of V. Then, the difference of the weights of Z under the embedding ψ and under the perturbed embedding ψ' is at most $2\sum_{v\in V} d_v \leq 2\delta\sum_{v\in V} r_v$. Since the weight of Z under the embedding ψ is at least $\frac{2}{k}\sum_{v\in V} r_v$, this difference amounts to a relative error of at most $k\delta = \epsilon/4$. Let Z_0 be a k-tour cover of V that is optimal under the perturbed embedding ψ'. Then, the weight of Z_0 is within relative error of $\epsilon/4$ compared to the optimal solution with respect to the original embedding ψ. When we measure the weight of Z_0 with respect to ψ, there is an additional relative error of at most $\epsilon/4$. Therefore, Z_0 is an approximate solution to the original problem with a relative error of at most $\epsilon/2$.

2.3 A characterization lemma

The following lemma, which characterizes the structure of the optimal k-tour cover of V with respect to the perturbed embedding ψ', forms a basis of our dynamic programming computation. Call an edge of a k-tour *lateral* if neither of its endnodes is the source.

Lemma 2.1 *Any optimal k-tour cover of V with respect to the embedding ψ' contains at most $O(m)$ lateral edges with positive length, where $m = |\Gamma|$.*

Intuitively, what this lemma says is plausible: in our situation where each grid point contains a large number of nodes (on average), most k-tours in an optimal k-tour cover visit only one grid point touching k nodes in one shot. The proof, however, is not trivial, especially when it is done for a general constant k. We prove this lemma for $k = 3$ below and omitt the general proof. Although the general case is considerably more involved, most of the proof ideas can be found in this special case.

Proof: (For $k = 3$). In this proof, we assume that each 3-tour in the optimal 3-tour cover of V under consideration visits exactly 3 nodes of V. If this is not the case, i.e., some of the 3-tours visit less than 3 nodes, then we may add appropriate copies of some nodes of V so that our assumption is satisfied; the result for the modified node set readily translates into the result for the original node set V.

Let Z be an arbitrary optimal 3-tour cover of V with respect to the embedding ψ'. For each 3-tour s, v_1, v_2, v_3, s of Z, we call v_1 and v_3 its *knees* and v_2 its *mid-node*. We also say that the gridpoint $\psi'(v_2)$ *owns* these knees v_1 and v_3. A knee v owned by p is *external* if $\psi'(v) \neq p$. Since each lateral edge with positive length (when $k = 3$) is between a mid-node and an external knee, it suffices to show that the number of external knees owned by a single gridpoint is bounded by a constant. Although we can in fact show the upper bound of 6 by a method specialized to $k = 3$, we prove the result for a larger constant using an approach that extends to the case where k is a general constant.

Let $p \in \Gamma$. We say that two knees owned by p are *paired by Z* if they belong to the same 3-tour in Z. It is important to observe that the particular pairing of the knees owned by p is not essential. For example, suppose that p owns four knees v_1, v_2, u_1, and u_2 and that Z pairs v_1 with v_2 and u_1 with u_2. In other words, Z has 3-tours s, v_1, v_0, v_2, s and s, u_1, u_0, u_2, s, where v_0 and u_0 are nodes placed on p by ψ'. A variant of Z that pairs v_1 with u_1 and v_2 with u_2, i.e., that has 3-tours, say, s, v_1, v_0, u_1, s and s, v_2, u_0, u_2, s instead of the above two, is clearly *equivalent* to Z in the sense that it covers the same set of nodes and has the same weight. In general we may create a variant of Z in a similar fashion that pairs the knees owned by p in an arbitrarily prescribed way.

Now, suppose p owns more than 18 external knees. We will derive a contradiction by locally modifying Z to produce a 3-tour cover of V with weight strictly smaller than that of Z. By a simple averaging argument, there must be a cone of angle smaller than $\pi/3$ with apex at p, that contains three external knees v_0, v_1, and v_2 owned by p. Denote $\psi'(v_i)$ by q_i, $0 \leq i \leq 2$, and assume that q_0 is the closest to p among the three points. We may assume, by the above observation on pairing, that v_1 and v_2 are paired by Z. That is, there is a 3-tour $T : s, v_1, u, v_2, s$ of Z where $\psi'(u) = p$. Let $U : s, v, u_0, v_0, s$ be the 3-tour of Z that contains knee u_0. We modify Z by trading u and v_0 between T and U. More precisely, let T' be the 3-tour s, v_1, v_0, v_2, s and let U' be the 3-tour s, v, u_0, u, s. See Figure 1. Then, let Z' be obtained from Z by replacing T and U by T' and U'. Clearly Z' covers the same set of nodes as Z.

It remains to compare their weights. Clearly, the weight of U' is no greater than that of U. The comparison of the weights of T and T' amounts to comparing $\overline{pq_1} + \overline{pq_2}$ and $\overline{q_0q_1} + \overline{q_0q_2}$, because the edges from the source are common to both T and T'. To compare $\overline{pq_1}$ with $\overline{q_0q_1}$, let r be the point on

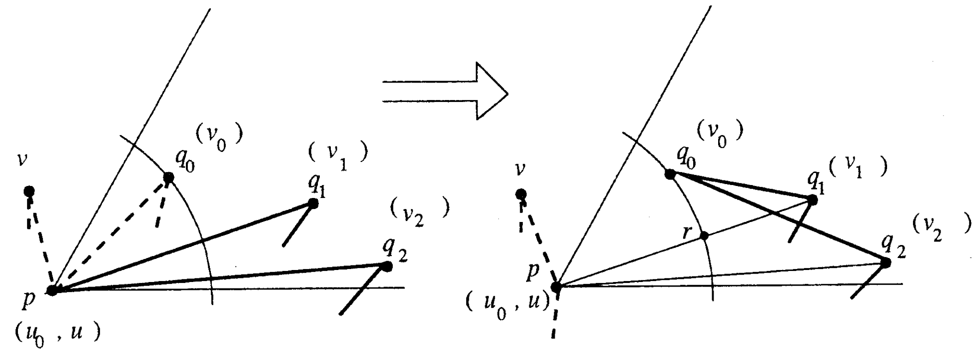

Figure 1: Reducing the weight of the 3-tour cover

the segment pq_1 such that $\overline{pq_0} = \overline{pr}$. Because q_0 and q_1 are in a cone of angle smaller than $\pi/3$, we have $\overline{q_0r} < \overline{pr}$ and hence $\overline{q_0q_1} < \overline{pq_1}$. Similarly, we have $\overline{q_0q_2} < \overline{pq_2}$. Therefore, the weight of T' is strictly smaller than that of T, a contradiction. □

Remark 1 In the above proof for the case $k = 3$, we did not need the fact that the distance between two points of Γ cannot be arbitrarily small. For the general case, where k is an arbitrary constant, we do use this fact.

Remark 2 We suspect that a stronger version of Lemma 2.1 holds: for any optimal k-tour cover with respect to embedding ψ', each grid point (individually rather than merely on average) is incident to $O(1)$ lateral edges with positive length. While such a stronger version would simplify the dynamic programming computation described in the next subsection, the proof would be probably more involved, if possible at all.

2.4 Dynamic programming

Our dynamic programming computation is rather standard. However, an entirely naive approach may result in a $m^{O(m)}$ running time rather than the desired $2^{O(m)}$. We sketch below how this can be avoided.

Let Q be an arbitrary subset of Γ, the set of grid points. Given a k-tour cover Z of V, define the Q-*part* of Z to be the set of subpaths of the k-tours of Z induced by $Q \cup \{s\}$. The interface between the Q-part of Z and the remaining part may be specified by the following parameters.

1. $f_{p,i}$, for each $p \in Q$ and integer i, $1 \leq i \leq k-1$: $f_{p,i}$ is the number of k-tours of Z whose ith visited node is mapped to p by ψ' and whose $(i+1)$st visited node is mapped by ψ' to a point in $\Gamma \setminus Q$.

2. $g_{p,i}$, for each $p \in Q$ and integer i, $2 \leq i \leq k$: $g_{p,i}$ is the number of k-tours of Z whose ith visited node is mapped to p by ψ' and whose $(i-1)$st visited node is mapped by ψ' to a point in $\Gamma \setminus Q$.

We call the tuple of the above parameter values the Q-*interface* of Z. Note that if Z and Z' are two k-tour covers of V that have identical Q-interfaces then their Q-parts are interchangeable: Z with its Q-part replaced by the Q-part of Z' is a k-tour cover of V. Given a Q-interface I, call a set of paths a *partial solution conforming to* I if it is a Q-part of some k-tour cover of V. A dynamic programming approach naturally suggests itself: starting from a small subset Q to larger ones, compute the minimum weight partial solution conforming to each of all possible Q-interfaces.

The number of parameters to be specified in a Q-interface is at most $2km = O(m)$. Moreover, due to the characterization lemma in the previous subsection, the sum of all the parameter values is at most $O(m)$, if the Q-interface is that of an optimal k-tour cover. Therefore, for each Q, the number of distinct Q-interfaces that must be considered in the dynamic programming computation is $2^{O(m)}$: the number of ways to distribute $O(m)$ indistinguishable balls to $O(m)$ bins.

Our dynamic programming computation proceeds as follows. List the grid points of Γ in an arbitrary fixed order as $p_1, \ldots, p_m$. Let $Q_i = \{p_1, \ldots, p_i\}$ be the set of first i points of Γ.

1. For each Q_1-interface I that must be considered, compute the unique partial solution conforming to I (if any). Set $i = 1$.

2. Repeat the following until $i = m$. For each Q_{i+1}-interface I that must be considered, compute the partial solution conforming to I (if any) that has the minimum weight. This is done by going through all the Q_i-interfaces I' with associated optimal partial solution for the interface and computing the best way to extend the partial solution for I' to a partial solution for I. Increment i by 1.

Once we have completed the computation, the optimal k-tour cover can be obtained as the partial solution conforming to the Γ-interface that have all parameter values set to 0.

We claim that the running time of the above computation is $2^{O(m)}$. The running time of the first step is clearly bounded by $2^{O(m)}$ since the number of Q_1-interfaces to be considered is at most $2^{O(m)}$. In the induction step, since the number of Q_i-interfaces and the number of Q_{i+1}-interfaces to be considered are both $2^{O(m)}$, it suffices to show that, for each pair of a Q_i-interface I' and a Q_{i+1}-interface I, the best way to extend a partial solution for I' to one for I can be computed in $2^{O(m)}$ steps. But this can be done by checking all the possible ways of placing $O(m)$ edges between p_{i+1} and the points of Q_i (which are fewer than m): at most $2^{O(m)}$ possibilities.

3 General metric case: APX-completeness

In this section, we consider the k-tour cover problem in the general metric case. An instance of the k-tour cover problem is given by a triple (V, s, w), where s is the source, V is the set of nodes to be covered, and $w : V \cup \{s\} \times V \cup \{s\} \to \mathbf{R}$ is a symmetric weight function that satisfies the triangle inequality: $w(v, u) = w(u, v) \leq w(u, t) + w(t, v)$.

We first recall the following result due to Haimovich and Rinnooy Kan that the metric k-tour cover problem is approximable within a factor of $2.5 - 1.5/k$. This result, unlike others in this paper, works even when k depends on n. For a tour T, we denote by $w(T)$ the weight of the tour, i.e., the sum of the weights of the adjacent node pairs in the tour.

Theorem 3.1 *A metric instance of the k-tour cover problem can be approximated within a factor of $2.5 - 1.5/k$ in polynomial time.*

Now we turn to the hardness of the metric k-tour cover problem.

Theorem 3.2 *Let $k \geq 3$ be fixed. Then, the k-tour cover problem is APX-complete.*

The proof is the reduction from the following problem. For given graphs G and H, an H-matching of G is a set of vertex-disjoint subgraphs of G each of which is isomorphic to H.

> **Maximum H-Matching-B:** Let H be a fixed graph and B some positive constant. The maximum H-matching-B problem is, given a graph G with maximum degree at most B, to find a H-matching of G with the largest possible cardinality.

The following theorem is due to Kann [Kan94]:

Theorem 3.3 *If H has a connected component containing at least three vertices then there is a constant B such that the maximum H-matching-B problem is APX-hard.*

Proof of Theorem 3.2.: Fix $k \geq 3$ and let $H = C_k$ be a chain consisting of k vertices. Let B be a constant in Theorem 3.3 so as to make the maximum H-matching-B problem APX-hard for this choice of H. We reduce this problem to the k-tour cover problem.

Let an instance $G = (V, E)$ of the maximum H-matching-B problem be given. From G, we construct a complete weighted graph $U(G)$ satisfying triangle inequality, such that an approximate solution to the k-tour cover problem on $U(G)$ gives an approximate solution to the H-matching problem on G. We assume that G is connected; otherwise, we work on connected components of G separately. An immediate consequence of this assumption and the bound B on the degree of G is that the cardinality of the optimal solution to the H-matching problem is $\Omega(n)$.

Let the vertices of V enumerated as $V = \{z_1, ..., z_n\}$. The node set $U(V)$ of $U(G)$ consists of the source s and $\{v_i, u_i : i = 1, 2, .., n\}$, where u_i is a multiple vertex with multiplicity $k-1$. Let r be a real number in the range $1 < r < 3/2$ whose value will be specified later. The edge weight function w is determined such that $w(v_i, v_j) = r$ if $(z_i, z_j) \in E$, and $w(s, u_i) = 1/4$ and $w(u_i, v_i) = 3/4$ for $i = 1, 2, .., n$. The weights of the other edges are the shortest path weight generated from the weights defined above. Indeed, $w(s, v_i) = 1$ for $i = 1, 2, .., n$, and $w(v_i, v_j) = 2$ if $(z_i, z_j) \notin E$.

We focus on the following three types of tours. A *straight tour* is a round trip tour from the source to some v_i that picks the $(k-1)$ copies of u_i on its way. A *mini-tour* is a round trip tour to u_i picking its $(k-1)$ copies. An *h-chain tour*, where $2 \leq i \leq k$, is based on a chain $z_{t_1}, .., z_{t_h}$ of length h in G and visits $v_{t_1}, \ldots, v_{t_h}$ in this order, without picking any u_i. Given any k-tour cover of $U(G)$, we can transform the cover into a "canonical form" that contains only tours of the above three types without increasing the weight: first remove redundant visits to each v_i by short-cutting and then add mini-tours to cover u_i's that have become unvisited due to the short-cutting. Thus, we may assume in the following that any k-tour cover is in the canonical form in this sense.

We compare an h-chain tour, together with the associated h mini-tours, to the h straight tours that cover the same set of vertices. The weight of the h-chain tour plus mini-tours is $2 + (h-1)r + h/2$ while the weight of the straight tours is $2h$. We choose r so that the former weight is smaller than the latter if and only if $h = k$. It can be easily verified that $r = (3k-7)/(2k-4)$ satisfies this requirement. Thus, for this choice of r, any optimal k-tour in the canonical form contains only straight tours and k-chain tours with accompanying mini-tours.

Suppose the cardinality of an optimal solution of the H-matching problem for G is l. Then, an optimal k-tour cover of $U(G)$ contains l k-chain tours. Suppose we have an approximate solution to this optimal k-tour cover with relative error of ϵ'. Since $l \geq \Omega(n)$ as noted before, this k-tour cover must contain $l' > (1-\epsilon)l$ k-chain tours if ϵ' is taken sufficiently small depending on ϵ. From these k-chain tours we can trivially obtain a H-matching of cardinality l'. Thus, our construction gives the required PTAS reduction. □

4 Open problems

As noted in the introduction, the approximability of the k-tour cover problem for a general k depending on n is an interesting open question. Probably it is not very difficult to show that the metric version is APX-complete. Constructing a PTAS for the Euclidean version seems to be a more challenging problem. It does not seem likely that the Arora's approach for TSP, that works for $k = \Omega(n)$, and our approach for a constant k can be combined to produce a PTAS for a general k.

Acknowledgments

We thank Madhu Sudan for helpful discussions.

References

[ALMSS92] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof verification and the intractability of approximation problems. *Proc. 33rd Annual IEEE Symposium on the Foundation of Computer Science,* 14-23, 1992.

[Aro96] S. Arora. Polynomial Time Approximation Schemes for Euclidean TSP and other Geometric Problems manuscript, 1996.

[CK] P. Crescenzi and V. Kann. A compendium of NP optimization problems ftp-able manuscript, 4th version, 1995.

[CT94] P. Crescenzi and L. Trevisan. On approximation scheme preserving reducibility and its applications. *Proc. 14th Annual Conference on Foundations of Software Technology and Theoretical Computer Science,* LNCS 880, Springer-Verlag, 330-341, 1994.

[Fei96] U. Feige. A threshold of $\ln n$ for approximating set cover. *Proc. 28th Annual ACM Symposium on the Theory of Computing*, 314-318, 1996.

[Fis] M.L. Fischer. Vehicle Routing. in *Network Routing,* Handbooks in Operations Research and Management Science, **8**, Ball, M. O., T. L. Magnanti, C. L. Monma and G. L. Nemhauser (Eds.), Elsevier Science, Amsterdam, 1-33, 1995.

[GA88] B.L. Golden and A. A. Assad (Eds.). *Vehicle Routing: Methods and Studies,* Studies in Manag. Science and Systems 16, North-Holland Publ., Amsterdam, 1988.

[GL95] M. Grötschel and L. Lovász. Combinatorial optimization in *Handbook of Combinatorics,* R.L.Graham, Grötschel and L. Lovász (Eds.), North Holland, 1541-1598, 1995.

[Gol84] B.L. Golden. Introduction to and Recent Advances in Vehicle Routing Methods. in *Transportation Planning Models,* M. Florian (Ed), Elsevier Sci. Publ. B. V. (North Holland), 383-418, 1984.

[HR85] M. Haimovich and A.H.G. Rinnooy Kan. Bounds and Heuristics for capacitated routing problems *Mathematics of Operations Research,* **10**(4), 527-542, 1985.

[HM85] D.S. Hochbaum and W. Maas. Approximation scheme for covering and packing problems in image processing and VLSI. *JACM,* **32**(1), 130-136, 1985.

[Joh74] D.S. Johnson. Approximation algorithms for combinatorial problems. *J. Comput. System Sci.* **9**, 256-278, 1974.

[Kan94] V. Kann. Maximum bounded H-matching is MAX SNP-complete. *Information Processing Letters,* **49**, 309-318, 1994.

[Lov75] L. Lovász. On the Ratio of Optimal Integral and Fractional Covers. *Discrete Mathematics* **13**, 383-390, 1975.

[Lun] K. Lund VRP References - Part I/II. WWW home page, http://www.imm.dtu.dk/documents/users/kl/vrp_1.html

[LY94] C. Lund and M. Yannakakis On the hardness of approximating minimization problems. *JACM,* **41**(5), 960-981, 1994.

[Mat93] J. Matoušek. Epsilon-nets and computational geometry Chapter III in *New Trends in Discrete and Computational Geometry,* Algorithms and Combinatorics **10**, J. Pach (Ed.), Springer-Verlag, 1993.

[KMSV94] S. Khanna, R. Motwani, M. Sudan, U. Vazirani. On syntactic versus computational views of approximability. *Proc. 35th Annual IEEE Symposium on the Foundation of Computer Science,* 819-830, 1992.

[PY91] C. H. Papadimitriou and M. Yannakakis. Optimization, approximation, and complexity classes. *J. Comput. System. Sci.* **43**, 425-440, 1991.

[PY93] C. H. Papadimitriou and M. Yannakakis. The traveling salesman problem with distances one and two. *Math. Oper. Res.* **18**, 1-11, 1993.