The Embroidery Problem

Esther M. Arkin* George W. Hart† Joondong Kim* Irina Kostitsyna†
Joseph S. B. Mitchell* Girishkumar R. Sabhnani† Steven S. Skiena†

Abstract

We consider the problem of embroidering a design pattern, given by a graph G, using a single minimum length thread. We give an exact polynomial-time algorithm for the case that G is connected. If G has multiple connected components, then we show that the problem is *NP-hard* and give a polynomial-time 2-*approximation* algorithm. We also present results for special cases of the problem with various objective functions.

1 Introduction

An embroidery is a decorative design sewn onto a fabric using one or more threads. The artist guides the thread with a needle as it alternates between the top and the bottom of the fabric. The exposed thread on the top of the fabric is the desired design; the thread on the bottom of the design is needed only to interconnect the needle holes as the design is sewn. We study the single-thread embroidery problem in which the goal is to minimize the total length of thread.

Figure 1: Embroidery of a girl with basket.

Model. We require that the complete embroidery must be done with a single continuous piece of thread and that the thread must form a cycle, returning to the starting point (where a knot will be tied). The *embroidery problem* is graph traversal optimization problem, as we now formally state.

Problem Statement. Given a *graph* $G(V, E)$, with *vertices* V and *edges* E embedded in the Euclidean plane, find a minimum-length closed tour T with alternating edge types (*front* and *back*), such that front edges exactly cover E (without repetition) and back edges form an arbitrary subset of the edges of the complete graph on V, with possible repetitions. We assume that V is a finite set of n points in the plane and that E is a set of m straight line segments joining pairs of points in V. The *length* of an edge is its Euclidean length; the total length of a tour or a set, X, of edges is denoted $|X|$.

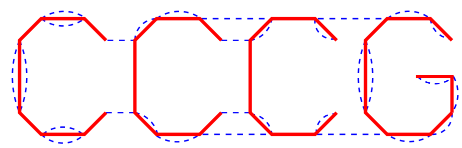

Figure 2: An embroidery graph, with red (solid) edges representing the front edges, E, of the embroidery design and blue (dashed) edges representing the back edges.

We refer to a tour T satisfying the above constraints as an *embroidery tour* for G. The front edges of T are denoted F, the back edges are denoted B. A single continuous piece of thread following T gives exactly the desired embroidery design $E = F$ (without repeating any edge) on the front of the cloth, and the back edges B of T represent "wasted" thread length. Since the edges F exactly cover E, the length of any feasible embroidery tour is simply $|T| = |E| + |B|$, so, for given E, exactly minimizing $|T|$ is equivalent to minimizing $|B|$. However, in terms of approximation ratio, the problem, OPT_T, of minimizing $|T|$ is different from the problem, OPT_B, of minimizing $|B|$.

We also consider the *Steiner* version of the embroidery problem in which we allow the set V to be augmented by a set of Steiner points that lie along edges E of the design; i.e., in the Steiner embroidery problem the set F of front edges must form an exact cover of the edges E, but each edge $e \in E$ may be (exactly) covered by a set of segments in F, with endpoints that may lie interior to e.

Related Work. The rural postman is most closely related to our problem: Given an undirected graph $G = (V, E)$ with edge weights, and a subset $E' \subseteq E$, find a closed walk of minimum weight traversing all edges of E' at least once. The stacker-crane problem is also similar, but the required edges to be traversed are directed. The main distinction between the embroidery

*Department of Applied Mathematics and Statistics, Stony Brook University, {estie,jdkim,jsbm}@ams.sunysb.edu
†Department of Computer Science, Stony Brook University, {ikost,george,gk,skiena}@cs.sunysb.edu

Graph G	OPT_T	OPT_B
Connected	poly-time Section 2.1	poly-time Section 2.1
Arbitrary	*NP-hard, 2-apx* Section 2.2	*NP-hard, 3-apx* Section 2.2
Indep Segments	*NP-hard, 1.5-apx,* Section 2.3	*NP-hard, 2-apx* Section 2.3

Table 1: Summary of results: No Steiner points allowed.

Graph G	OPT_T	OPT_B
Connected	poly-time Section 3.1	poly-time Section 3.1
Arbitrary	*NP-hard, 2-apx, PTAS* Section 3.2	*NP-hard, 3-apx* Section 3.2

Table 2: Summary of results with Steiner points.

problem and these related problems is that in the embroidery problem the tour is not allowed to traverse two of the required edges in a row; it must alternate between the front (specified) and back edges. The rural postman has a (Christofides-like) 3/2-approximation [3] and the stacker-crane has a 9/5-approximation [4]. Biedl [2] studies the special case of the embroidery problem in which only "cross-stitches" are used.

Summary of Results. Table 1 summarizes our results on the OPT_T and OPT_B problems for different types of input embroidery graphs G: (i) connected, (ii) arbitrary, with possibly many connected components, and (iii) an independent set of edges – no two edges of E share an endpoint (however, the line segments that embed E may cross arbitrarily). Table 2 lists our results for the Steiner embroidery problem.

2 Embroidery Without Steiner Points

An embroidery tour T alternates between front edges and back edges. Hence $\forall v \in V$ the number of back edges incident to v must be exactly equal to the number of front edges incident to v. See Theorem 1.

Theorem 1 *T is an embroidery tour for $G(V,E)$ if and only if $G(V,T)$ is connected and $\forall v \in V$: $d_F(v) = d_B(v)$, where $d_F(v)$ is the degree of vertex v in $G(V,F)$.*

Proof. *If:* Since T is an embroidery tour (using single continuous thread) $G(V,T)$ must be connected. If there exists a vertex v such that $d_B(v) < d_F(v)$, by the pigeon-hole-principle on entry and exit type of edges on v, T must have two consecutive front edges $e_i, e_j \in F$ sharing v, hence contradicting that T is embroiderable. A similar contradiction holds if $d_B(v) > d_F(v)$.

Only If: Since $G(V,T)$ is connected and $d_T(v)$ is even there exists an Euler tour in $G(V,T)$. We show how to construct an Euler tour that alternates edges from F and B. First, we show that $G(V,T)$ must contain an edge-alternating circuit. Start an edge-alternating walk $W = \{a, \ldots, v, x \ldots, y, v\}$ from an arbitrary vertex, until a vertex v repeats. This defines an edge-alternating circuit, unless edges (v,x) and (y,v) belong to the same side. But if so, W can be continued, as there remains at least one unused alternate side edge incident on v. Thus, we can decompose $G(V,T)$ into a set of edge-disjoint, alternating circuits. Any two such circuits (say c_1 and c_2) incident on a common vertex v' can be merged to form a larger alternating circuit, since both c_1 and c_2 contain front and back edges at v'. Repeated merging operations reduce the set of alternating circuits to a single alternating tour. □

2.1 One Connected Component

If G is connected, then the embroidery problem can be solved as follows: Find a minimum-length set of back edges B such that the degree requirement $\forall v \in V$: $d_B(v) = d_E(v)$ is satisfied. By Theorem 1 $E \cup B$ is an embroidery tour for G. Since B is minimum-length, the resulting tour is optimal. Thus, the selection of an optimal B is exactly the minimum-weight b-*matching* problem on V, with vertex weights (degrees) $b(v) = d_E(v)$, $\forall v \in V$, which is solvable in polynomial time [1].

2.2 Multiple Connected Components

Consider now an arbitrary design G, with possibly many connected components. As before, we can compute a minimum-weight b-matching, with vertices weighted by the degrees, $d_E(v)$; however, an optimal b-matching does not result in a set B of back edges that yields a complete solution, since the graph $G(V, E \cup B)$ may be disconnected.

In fact, we show that it is NP-hard to solve OPT_T or OPT_B exactly, using a simple reduction from Euclidean TSP ([6]):

Theorem 2 *The embroidery problem (either OPT_T or OPT_B) is NP-hard for arbitrary graphs G, with many connected components.*

Approximating OPT_T. We turn now to approximating OPT_T. We define a new graph $G'(V', E')$, where V' is the set of connected components in $G(V,E)$, and E' is the set of edges in the complete graph on V'. For each edge $e(i,j) \in E'$, the weight of the edge $w(i,j) = \min_{u \in V_i, w \in V_j} dist(u,w)$. Let MST be a minimum spanning tree of G'.

Now initialize B to contain a copy of front edges E and two copies of each MST edge. Note that each MST

edge is a minimum-weight edge connecting the appropriate vertices in the two different components. Let $T_{apx} = E \cup B$. Hence, $\forall v \in V, d_{T_{apx}}(v) \geq 2 \cdot d_E(v)$, implying

$$|T_{apx}| = 2 \cdot |E| \; + \; 2 \cdot |MST|.$$

Theorem 3 *T_{apx} can be converted to an embroidery tour for $G(V, E)$ with length $\leq 2 \cdot OPT_T$.*

Proof. By the definition of T_{apx}, $G(V, T_{apx})$ is connected (since it uses the MST edges to connect between different connected components) and $d_{T_{apx}}(v)$ is even. Also since $\forall v \in V, d_{T_{apx}}(v) \geq 2 \cdot d_E(v)$, we can find an Euler tour in $G(V, T_{apx})$, such that there are no consecutive front edges. We also try to "avoid" consecutive back edges, by choosing to leave a node on a front edge, if it was entered on a back edge, if possible. Note that T_{apx} may have consecutive back edges $e_i(v_i, v), e_j(v, v_j) \in B$ at a vertex v, in which case we can shortcut using $e'(v_i, v_j)$ and update $T_{apx} := T_{apx} \cup \{e'\} \setminus \{e_i, e_j\}$ without increasing $|T_{apx}|$ (by triangle inequality). Since all front edges touching v had already been used, we know that this shortcut does not disconnect v from the tour. Thus we can convert T_{apx} to a tour containing alternate front and back edges to make it an embroidery tour without increasing its cost.

Also, $OPT_T = |T_{opt}| \geq |E| + |MST|$, since T_{opt} must cover all the edges in E and must also span all the connected components and by definition of MST, it is the cheapest way to connect the disconnected components. Thus, $2 \cdot OPT_T \geq 2 \cdot (|E| + |MST|) \geq |T_{apx}|$. □

Approximating OPT_B. We start by finding a minimum-weight *b-matching* M_b of V with weight $b(v) = d_E(v), \forall v \in V$. Then for all connected components $V_1, V_2, \ldots, V_k$ in graph $G(V, E \cup M_b)$, we find the MST on graph $G'(V', E')$, as we did in Section 2.2. Now add a copy of each M_b edge and two copies of each MST edge to B. Let $T_{apx} = E \cup B$. Again, $\forall v \in V, d_{T_{apx}}(v) \geq 2 \cdot d_E(v)$. Thus,

$$|B| = |M_b| \; + \; 2 \cdot |MST|.$$

Theorem 4 *T_{apx} can be converted to an embroidery tour for $G(V, E)$ with back edges of length $\leq 3 \cdot OPT_B$.*

Proof. By similar arguments as in the proof of Theorem 3, we can convert T_{apx} to an embroidery tour with $|B| \leq |M_b| + 2 \cdot |MST|$. Now, $OPT_B = |B_{opt}| \geq |M_b|$, since in T_{opt}, B_{opt} is one b-matching satisfying $b(v) = d_E(v)$ and M_b is a minimum-weight b-matching. Also $OPT_B \geq |MST|$, since T_{opt} must span all of the connected components of $G(V, E)$. Note that MST here is a minimum spanning tree on the connected components of $G(V, E \cup M_b)$, which has smaller cost as compared to the minimum spanning tree on connected components of $G(V, E)$. Thus, $3 \cdot OPT_B \geq |M_b| + 2 \cdot |MST| \geq |B|$. □

2.3 Independent Segments

In the case that the edges E do not share endpoints (i.e., they form a set of possibly intersecting line segments), OPT_T can be approximated using the 3/2-approximation algorithm for the rural postman: If the approximating tour uses two consecutive back edges, then we simply shortcut, replacing the two edges with one shorter back edge. This results in a 3/2-approximation for OPT_T.

3 Embroidery with Steiner Points

It may be possible to use a shorter thread if we allow a front edge to be split into two or more subsegments by placing *Steiner* points judiciously along it. In fact, by placing a Steiner point arbitrarily close to an endpoint (vertex) of a front edge, we can make the length of back edges arbitrarily close to zero; see Figure 3. We say that such a Steiner point *doubles* the vertex where it is placed.

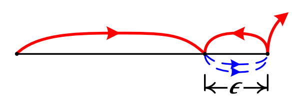

Figure 3: Placing a *Steiner* point near a vertex.

3.1 One Connected Component

Lemma 5 *There exists an optimal embroidery tour T_{opt} (allowing Steiner points) for a connected $G(V, E)$ that does not have Steiner points on edges other than those near endpoints that double vertices.*

Proof. If s is a *Steiner* point interior to an edge $e \in E$, with back edges $e(a, s), e(s, b) \in B$ incident to s, then we can simply replace these two edges with a single (back) edge $e(a, b)$ (and remove Steiner point s) without increasing the cost of tour $|T_{opt}|$. See Figure 4. □

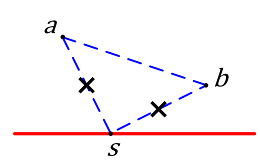

Figure 4: An optimal tour T exists without having a Steiner point s interior to an edge.

Lemma 6 *There exists an optimal tour T_{opt} (allowing Steiner points) for a connected $G(V, E)$ that does not have two back edges $e_i, e_j \in B$ incident to a common vertex $v \in V$.*

Proof. Similar to the proof of Lemma 5. □

Since an optimal embroidery tour T_{opt} is an *Euler* cycle (by definition of T), the sum of front and back edge degrees for each vertex is even. Thus, all odd-degree vertices $v \in V$ (if any present) have one back outgoing edge, and even-degree vertices do not have any outgoing back edges (using Lemma 5, 6). Therefore, an optimal solution is a union of front edges and back edges constituting minimum-weight perfect matching edges built on odd-degree vertices. The problem hence reduces to finding a minimum-weight perfect matching in a complete graph (of odd-degree vertices in this case), which can be solved in time $O(n^3)$.

3.2 Multiple Connected Components

Clearly, the same NP-hardness reduction for the non-Steiner version applies also if we allow Steiner points.

Approximating OPT_T. The idea is very similar to Section 2.2, except that the graph $G'(V', E')$ (defined over different components) has edge weights $w(i, j) = \min_{u \in G(V_i,E), w \in G(V_j,E)} dist(u, w), \forall e(i, j) \in E'$ (where u, w are edges). We refer to this minimum spanning tree on this new graph $G'(V', E')$ as MST_{St}. As before, we add a copy of front edges, F in this case (since each edge e from E that contains one or more Steiner points, gets split and is put as two ore more segments in F) and two copies of each MST_{St} edge to B. Note that $|F| = |E|$, as F exactly covers E. Let $T_{apx} = F \cup B$. Thus,

$$|T_{apx}| = 2 \cdot |E| \ + \ 2 \cdot |MST_{St}|$$

Theorem 7 *T_{apx} can be converted to an embroidery tour (allowing Steiner points) for $G(V, E)$ with length $\leq 2 \cdot OPT_T$.*

Proof. We note that every time we introduce a Steiner point, we create a new vertex v' with $d_F(v') = 2$. Since the introduction of a Steiner point is only because of some MST_{St} edge and because we double the MST_{St} edge, $d_{T_{apx}}(v') \geq 2 \cdot d_E(v')$ for all new Steiner points v'. Excluding other details (which are similar to those in Theorem 3), T_{apx} can be converted to an embroidery tour without increasing its cost.

Also, as before, $OPT_T = |T_{opt}| \geq |E| + |MST_{St}|$. Thus, $2 \cdot OPT_T \geq 2 \cdot (|E| + |MST_{St}|) \geq |T_{apx}|$. □

Approximating OPT_B. This idea is also very similar to Section 2.2, except that, instead of M_b, it uses a perfect matching M on odd-degree vertices in $G(V, E)$. It also uses MST_{St} defined in Section 3.2. We add a copy of each edge of M and two copies of each MST_{St} edge to B. Let $T_{apx} = F \cup B$, where F is the front edge cover of Steiner point split edges in E. Then,

$$|B| = |M| \ + \ 2 \cdot |MST_{St}|.$$

Theorem 8 *T_{apx} can be converted to an embroidery tour (allowing Steiner points) for $G(V, E)$ with back edges of length $\leq 3 \cdot OPT_B$.*

Proof. We apply shortenings to B as in Lemmas 5, 6. The result is a perfect matching of odd-degree vertices of $G(V, E)$. Thus, $OPT_B \geq |M|$, since the cost of any perfect matching is at least as much as the cost of the minimum weight perfect matching. Also $OPT_B \geq |MST_{St}|$, since T_{opt} must span all of the connected components of $G(V, E)$. Thus, $3 \cdot OPT_B \geq |M| + 2 \cdot |MST_{St}| \geq |B|$. □

3.3 A PTAS

By using the m-guillotine method for geometric network approximation [5], we obtain a PTAS for the problem:

Theorem 9 *The embroidery problem with Steiner points and an arbitrary input graph G has a PTAS for OPT_T.*

3.4 OPT_{St} vs OPT_{NSt}

We analyze how much one actually gains by allowing Steiner points to be inserted on front edges:

Theorem 10 $OPT_{St} \geq \frac{1}{2} OPT_{NSt}$.

Acknowledgments

We thank Bonnie Skiena for suggesting the problem. We thank the anonymous referees for their helpful suggestions. This work was partially supported by grants from the National Science Foundation (EIA-0325123, CCF-0431030, DBI-0444815, CCF-0528209, CCF-0729019), NASA Ames, and Metron Aviation.

References

[1] R. P. Anstee. A polynomial algorithm for b-matchings: an alternative approach. *Information Processing Letters*, 24(3):153–157, 1987.

[2] T. Biedl, J. D. Horton, and A. Lopez-Ortiz. Cross-stitching using little thread. In *17th Canadian Conference of Computational Geometry*, pages 196–199, 2005.

[3] H. A. Eiselt, M. Gendreau, and G. Laporte. Arc routing problems, part II: The rural postman problem. *Operations Research*, 43(3):399–414, 1995.

[4] G. N. Frederickson, M. S. Hecht, and C. E. Kim. Approximation algorithms for some routing problems. *SIAM Journal on Computing*, 7(2):178–193, 1978.

[5] J. S. B. Mitchell. Guillotine subdivisions approximate polygonal subdivisions: A simple polynomial-time approximation scheme for geometric TSP, k-MST, and related problems. *SIAM Journal on Computing*, 28:1298–1309, 1999.

[6] C. H. Papadimitriou. The Euclidean traveling salesman problem is NP-complete. *Theoretical Computer Science*, 4:237–244, 1977.