

Transformation of the Generalized Traveling-Salesman Problem into the Standard Traveling-Salesman Problem

YAO-NAN LIEN

AT&T Bell Laboratories, Naperville, Illinois 60566

EVA MA

*IBM Research Division, T. J. Watson Research Center, Yorktown Heights,
New York 10598*

and

BENJAMIN W.-S. WAH

*Center for Reliable and High Performance Computing, Coordinated Science
Laboratory, University of Illinois, Urbana, Illinois 61801*

ABSTRACT

In the generalized traveling-salesman problem (GTSP), we are given a set of cities that are grouped into possibly intersecting clusters. The objective is to find a closed path of minimum cost that visits at least one city in each cluster. Given an instance G of the GTSP, we first transform G into another instance G' of the GTSP in which all the clusters are nonintersecting, and then transform G' into an instance G'' of the standard traveling-salesman problem (TSP). We show that any feasible solution of the TSP instance G'' can be transformed into a feasible solution of the GTSP instance G of no greater cost, and that any optimal solution of the TSP instance G'' can be transformed into an optimal solution of the GTSP instance G .

1. INTRODUCTION

In the *standard traveling-salesman problem* (TSP), we are given a set of cities and the distances (costs) between them; the objective is to find a tour (which is a closed path visiting each city exactly once) of minimum cost [1]. In the *generalized traveling-salesman problem* (GTSP), the cities are grouped

This research was partially supported by National Science Foundation Grants MIP 85-19649 and MIP 88-10584, and National Science Foundation PYI Award DCR84-51408.

into possibly intersecting clusters, and the objective is to find a g-tour (which is a closed path visiting at least one city in each cluster) of minimum cost [2].

For many real-world problems that are inherently hierarchical, the GTSP offers a more accurate model than the TSP. As an example, a traveling salesman may want to visit all his dealers in the country. To reduce the traveling cost, for each state, he would meet all the local dealers in only one out of several possible cities in the state. His objective is then to choose a set of cities, with one city for each state, and a route through them so that all the dealers can be visited in the minimum cost. As another example, a Post Office administrator may want to choose for each residential area a location out of several possible ones for a central mailbox in the area, as well as a route through these mailboxes so that the mail can be collected in the minimum cost. As a final example, in the layout of a ring network, the designer may have to select for each region a site out of several possible ones as the concentrator, as well as a loop through these concentrators so that all the regions can be connected at the minimum cost. These three problems can all be formulated as the GTSP.

The GTSP, of which the TSP is a special case, is obviously NP-hard. Polynomial-time heuristics are thus needed for solving this problem, especially when the problem size is large. In this paper, we study those cases of the GTSP for which the distance measures satisfy the triangular inequality. We develop a set of techniques to transform the GTSP into the TSP. With this transformation, polynomial-time heuristics for the TSP can be applied directly to solve the GTSP.

An obvious way to transform the GTSP to the TSP is by decomposing a GTSP instance into multiple smaller size TSP instances, each of which is defined by a distinct set of cities, with one city from each cluster. The different TSP instances correspond to the different possible choices of the cities from the clusters. A tour of any of these TSP instances is obviously a g-tour of the original GTSP instance. Further, since we assume that in the GTSP instance the distance measure satisfies the triangular inequality, there always exists an optimal g-tour that visits only one node in each cluster. Such an optimal g-tour can be obtained by finding a tour of minimum cost among all the optimal tours of the corresponding TSP instances.

On the other hand, in the transformations that we develop in this paper, a GTSP instance is transformed into exactly one TSP instance of larger size. Any tour of this TSP instance can be transformed into a g-tour of the GTSP instance of no greater cost, and any optimal tour of this TSP instance can be transformed into an optimal g-tour of the original GTSP instance. Since this transformation method only requires solving a single

TSP instance of larger size, rather than solving many TSP instances of smaller size, this method is more efficient if polynomial-time heuristics are used to solve the TSP, but less efficient if exponential-time algorithms are used instead. Another important difference between the two transformation methods is that in the resulting TSP instance, the triangular inequality on distance measure is preserved under the first method, but not under the second.

2. PRELIMINARIES

Unless stated otherwise, variables denote positive integers, and graphs $G=(V, E)$ are complete and directed, with $|V| \geq 3$ and with a distance measure d from the set E to the set of reals that satisfies the triangular inequality (that is, for any three distinct nodes $u, v, w \in V, d(u, v) + d(v, w) \geq d(u, w)$). Given any two paths $P_1 = v_1 \rightarrow v_2 \rightarrow \dots \rightarrow v_i$ and $P_2 = u_1 \rightarrow u_2 \rightarrow \dots \rightarrow u_j$ in G , we use $P_1 \rightarrow P_2$ to denote the path $v_1 \rightarrow v_2 \rightarrow \dots \rightarrow v_i \rightarrow u_1 \rightarrow u_2 \rightarrow \dots \rightarrow u_j$. For all positive integers k , we use $[k]$ to denote the set $\{0, 1, \dots, k\}$, and $[k]^+$ to denote the set $\{1, 2, \dots, k\}$.

Let $G=(V, E)$ be a graph for which $V = \{v_0, v_1, \dots, v_n\}$. The nodes in V are grouped into $m+1$ possibly intersecting, nonempty subsets, C_0, C_1, \dots, C_m , such that $C_0 = \{v_0\}$; for all $i \in [m]^+, C_0 \cap C_i = \emptyset$; and $\bigcup_{i=0}^m C_i = V$. We call these subsets *clusters*. For all $i \in [m]$, we call a cluster C_i an *intersecting cluster* if there exists some $j \in [m]$ such that $j \neq i$, and $C_i \cap C_j \neq \emptyset$; and a *nonintersecting cluster* otherwise. By definition, C_0 is a nonintersecting cluster. A path $v_{i_0} \rightarrow v_{i_1} \rightarrow \dots \rightarrow v_{i_l}$ in G is called a *tour* if (i) $l = n + 1$; (ii) $i_0 = i_l = 0$; and (iii) for all $j, k \in [l]^+, j \neq i_k$. The path is called a *g-tour* if (i) $m + 1 \leq l \leq n + 1$; (ii) $i_0 = i_l = 0$; (iii) for all $j, k \in [l]^+, j \neq i_k$; and (iv) for all $k \in [m]^+, C_k \cap \{v_{i_1}, v_{i_2}, \dots, v_{i_{l-1}}\} \neq \emptyset$. The cost of a tour or a g-tour is given by $\sum_{k=0}^{l-1} d(v_{i_k}, v_{i_{k+1}})$. By definition, a tour is a special case of a g-tour.

For the case in which all clusters in G are nonintersecting, we call an edge (v_i, v_j) an *intercluster edge* if v_i and v_j belong to two distinct clusters and an *intracluster edge* otherwise; we call a g-tour *nonredundant* if the tour visits each cluster exactly once. These three terms are defined only on those graphs in which all the clusters are nonintersecting.

Given a graph G with a distance measure d , the TSP is the problem of finding a tour of minimum cost among all possible tours in G . Given that the nodes in G are also grouped into clusters, the GTSP is the problem of finding a g-tour of minimum cost among all possible g-tours in G . For the case in which each cluster in G consists of only one distinct node, the GTSP degenerates into the TSP.

3. THE TRANSFORMATIONS

Given an instance G of the GTSP with possibly intersecting clusters, we transform it into an instance G'' of the TSP through two transformation steps. In the first step, G is transformed into an instance G' of the GTSP *without* intersecting clusters such that every g-tour in G' can be transformed into a g-tour in G of no greater cost, and every optimal g-tour in G' can be transformed into an optimal g-tour in G . We call this transformation the *I-N transformation* (intersecting cluster to nonintersecting cluster). In the second step, we transform G' into an instance G'' of the TSP such that there is a one-to-one correspondence between the tours in G'' and the nonredundant g-tours in G' , and such that the tour and the nonredundant g-tour in any corresponding pair have the same cost. Since, as shown below, there always exists a nonredundant g-tour in G' that is optimal, an optimal g-tour in G' can be obtained from an optimal tour in G'' . We call this transformation the *G-S transformation* (generalized TSP to standard TSP). Under these two transformations, any tour in G'' can be transformed into a g-tour in G of no greater cost, and any optimal tour in G'' can be transformed into an optimal g-tour in G . If all the clusters in G are nonintersecting, then the I-N transformation can be omitted, and the G-S transformation can be applied directly with $G' = G$.

As shown in the next two subsections, the graph G' is also a complete graph, with a distance measure satisfying the triangular inequality; neither of these properties, however, holds for the graph G'' .

3.1. THE I-N TRANSFORMATION

Let $G = (V, E)$ be a graph with clusters C_0, C_1, \dots, C_m , some of which are intersecting. The nodes in V are denoted by $0, 1, \dots, n$. For all $x \in [n]$, let $S_x = \{i \mid i \in [m], \text{ and } x \in C_i\}$; S_x is the set of indices of all the clusters containing the node x . We have $C_0 = \{0\}$; C_0 is a nonintersecting cluster; and $S_0 = \{0\}$. For all $x \in [n]$, let $q_x = |S_x|$.

We transform G into a graph $G' = (V', E')$ with $m + 1$ nonintersecting clusters C'_0, C'_1, \dots, C'_m . For each node x in V , we create q_x nodes in V' . These nodes are called the *replicas* of x , and are denoted by $u_{x,i}$, for all $i \in S_x$, with $u_{x,i}$ in C'_i . The edges in E' are defined as follows.

(i) For all $x \in [n]$, and for all distinct $i, j \in S_x$, we create an edge $(u_{x,i}, u_{x,j})$ of zero cost.

(ii) For all distinct $x, y \in [n]$, corresponding to each edge (x, y) in E , we create for all $i \in S_x$ and all $j \in S_y$, the edges $(u_{x,i}, u_{y,j})$ in E' of the same cost.

By the definition of E' , the graph G' is a complete graph. For all $x \in [n]$, let h_x be a path of length $q_x - 1$ that connects all replicas of the node x in some arbitrary, fixed order, that is, $h_x = u_{x,i_0} \rightarrow u_{x,i_1} \rightarrow \dots \rightarrow u_{x,i_{q_x-1}}$, where for all distinct $k, l \in [q_x - 1]$, $i_k \neq i_l$ and $i_k, i_l \in S_x$. We call h_x the *internal subtour* of the replicas of x . An internal subtour has zero cost. For the case that $q_x = 1$, the internal path consists of only one node, and has length 0.

An example of applying the I-N transformation on a graph G to obtain the graph G' is given in Figure 1. Note that both G and G' are complete graphs, but to simplify the figure, most of their edges are omitted. Figure 1 also shows a g-tour in G and its corresponding g-tour in G' (the correspondence between the g-tours in G and the g-tours in G' is established in Theorem 6).

Lemma 1 shows that the distance measure on E' also satisfies the triangular inequality.

LEMMA 1. *The distance measure d (from the set E' to the set of reals) satisfies the triangular inequality.*

Proof. Let a, b , and c be three arbitrary, distinct nodes in V' . We want to show that $d(a, b) + d(b, c) \geq d(a, c)$. There are three cases.

Case 1. The three nodes are the replicas of the same node.

In this case, the edges among a, b , and c all have zero cost. Thus, $d(a, b) + d(b, c) = d(a, c)$.

Case 2. Two of the nodes are replicas of the same node, while the third node is the replica of a different node.

Without loss of generality, we assume that a and b are replicas of the same node, say x , and c is the replica of a different node, say y , such that $x, y \in [n]$ and $x \neq y$. We have $d(a, b) = 0$, $d(a, c) = d(b, c) = d(x, y)$. Therefore, $d(a, b) + d(b, c) = d(a, c)$.

Case 3. a, b , and c are replicas of three distinct nodes.

Let a, b , and c be the replica of x, y , and z , respectively, where $x, y, z \in [n]$, and $x \neq y \neq z$. Since the graph G satisfies the triangular inequality, we have

$$d(a, b) + d(b, c) = d(x, y) + d(y, z) \geq d(x, z) = d(a, c).$$

Since the triangular inequality is satisfied in all three cases, the lemma follows. ■

Lemma 2 shows that if there is a g-tour in G' that visits two replicas of the same node in a nonconsecutive order, then we can always construct




Fig. 1. Example of I-N transformation.

another g-tour in G' of no greater cost, in which the two replica are visited consecutively in some arbitrary order, and except for these two replicas, the two g-tours visit the same sequence of nodes.

LEMMA 2. Assume that $T' = u_{0,0} \rightarrow u_{x_1, i_1} \rightarrow u_{x_2, i_2} \rightarrow \dots \rightarrow u_{x_{k-1}, i_{k-1}} \rightarrow u_{x_k, i_k} \rightarrow u_{x_{k+1}, i_{k+1}} \rightarrow \dots \rightarrow u_{x_{l-1}, i_{l-1}} \rightarrow u_{x_l, i_l} \rightarrow u_{x_{l+1}, i_{l+1}} \rightarrow \dots \rightarrow u_{x_r, i_r} \rightarrow u_{0,0}$, where $m \leq r \leq n$, is a g-tour in G' such that $k, l \in [r]^+$, $l > k + 1$, and $x_k = x_l$.

Then $T^* = u_{0,0} \rightarrow u_{x_1,i_1} \rightarrow u_{x_2,i_2} \rightarrow \dots \rightarrow u_{x_{k-1},i_{k-1}} \rightarrow a \rightarrow b \rightarrow u_{x_{k+1},i_{k+1}} \rightarrow \dots \rightarrow u_{x_{l-1},i_{l-1}} \rightarrow u_{x_{l+1},i_{l+1}} \rightarrow \dots \rightarrow u_{x_r,i_r} \rightarrow u_{0,0}$, where $\{a, b\} = \{u_{x_k,i_k}, u_{x_l,i_l}\}$, is also a g-tour in G' of no greater cost.

Proof. We first assume that $a = u_{x_k,i_k}$ and $b = u_{x_l,i_l}$. Since G' is a complete graph, T^* is a closed path in G' ; since T^* visits the same set of nodes as T' , T^* is a g-tour in G' . Further, $x_k = x_l$ implies that u_{x_k,i_k} and u_{x_l,i_l} are the replicas of the same node in G . Thus, we have $d(u_{x_k,i_k}, u_{x_l,i_l}) = 0$ and $d(u_{x_l,i_l}, u_{x_{k+1},i_{k+1}}) = d(u_{x_k,i_k}, u_{x_{k+1},i_{k+1}})$. Further, by Lemma 1, we have

$$d(u_{x_{l-1},i_{l-1}}, u_{x_{l+1},i_{l+1}}) \leq d(u_{x_{l-1},i_{l-1}}, u_{x_l,i_l}) + d(u_{x_l,i_l}, u_{x_{l+1},i_{l+1}}).$$

Therefore, T^* has a cost no greater than the cost of T' .

The proof for the case with $a = u_{x_l,i_l}$ and $b = u_{x_k,i_k}$ is similar, and thus omitted. ■

Lemma 3 shows that given any two replicas of the same node, if there exists a g-tour in G' that visits only one of the two replicas, then we can always construct another g-tour in G' of the same cost, in which both replicas are visited constructively in some arbitrary order, and except for the additional replica, the two g-tours visit the same sequence of nodes.

LEMMA 3. Let $T' = u_{0,0} \rightarrow u_{x_1,i_1} \rightarrow u_{x_2,i_2} \rightarrow \dots \rightarrow u_{x_{k-1},i_{k-1}} \rightarrow u_{x_k,i_k} \rightarrow u_{x_{k+1},i_{k+1}} \rightarrow \dots \rightarrow u_{x_r,i_r} \rightarrow u_{0,0}$ be a g-tour in G' , where $m \leq r \leq n$, and $k \in [r]^+$. Assume that there exists some $j \in S_{x_k}$ such that $u_{x_k,j}$ is not in T' . Then $T^* = u_{0,0} \rightarrow u_{x_1,i_1} \rightarrow u_{x_2,i_2} \rightarrow \dots \rightarrow u_{x_{k-1},i_{k-1}} \rightarrow a \rightarrow b \rightarrow u_{x_{k+1},i_{k+1}} \rightarrow \dots \rightarrow u_{x_r,i_r} \rightarrow u_{0,0}$, where $\{a, b\} = \{u_{x_k,i_k}, u_{x_k,j}\}$, is also a g-tour in G' of the same cost.

Proof. We first assume that $a = u_{x_k,i_k}$ and $b = u_{x_k,j}$. Since G' is a complete graph, and since all the nodes visited by T' are visited by T^* , T^* is also a g-tour in G' . Since $d(u_{x_k,i_k}, u_{x_k,j}) = 0$ and $d(u_{x_k,i_k}, u_{x_{k+1},i_{k+1}}) = d(u_{x_k,j}, u_{x_{k+1},i_{k+1}})$, T^* has the same cost as T' 's.

The proof for the case with $a = u_{x_k,j}$ and $b = u_{x_k,i_k}$ is similar, and thus omitted. ■

Lemma 4 shows that given an arbitrary g-tour T' in G' , we can always transform it into another g-tour T^* in G' such that (i) T^* consists only of a sequence of internal subtours, (ii) there is an internal subtour h_x in T^* connecting all the replicas of x if and only if T' contains at least one of the replicas of x , and (iii) the internal subtours appear in T^* in the same order as the corresponding replicas first appear in T' .

LEMMA 4. Let $T' = u_{0,0} \rightarrow u_{x_1,i_1} \rightarrow u_{x_2,i_2} \rightarrow \dots \rightarrow u_{x_r,i_r} \rightarrow u_{0,0}$ be a g-tour in G' , where $m \leq r \leq n$. Let s , for $1 \leq s \leq r$, be the number of distinct elements

visited replicas,

$u_{x_{k-1},i_{k-1}} \rightarrow u_{0,0}$,
 $x_k = x_l$.

in the sequence x_1, x_2, \dots, x_r , and let y_1, y_2, \dots, y_s be these s distinct elements in the same order as they first appear in the sequence x_1, x_2, \dots, x_r . Then $T^* = u_{0,0} \rightarrow h_{y_1} \rightarrow h_{y_2} \rightarrow \dots \rightarrow h_{y_s} \rightarrow u_{0,0}$ is a g-tour in G' of no greater cost.

Proof. By Lemma 2, we can construct from T' a g-tour \tilde{T} in G' of no greater cost such that (i) \tilde{T} visits the same set of nodes as T' , (ii) for all $i \in [s]^+$, those replicas of y_i that are in T' appear consecutively in some arbitrary order in \tilde{T} , and (iii) for all $i, j \in [s]^+$, if $i < j$, then the replicas of y_i precede the replicas of y_j in \tilde{T} .

By Lemma 3, we can next construct from \tilde{T} a g-tour \hat{T} in G' of no greater cost such that (i) for all $i \in [s]^+$, \hat{T} visits *all* the replicas of y_i , with these replicas appearing consecutively in some arbitrary order in \hat{T} , and with the union of the replicas of y_1, y_2, \dots, y_s being all the nodes visited by \hat{T} ; and (ii) for all $i, j \in [s]^+$, if $i < j$, then the replicas of y_i precede the replicas of y_j in \hat{T} .

Since for all $i \in [s]^+$, the order in which the replicas of y_i appear in \hat{T} is arbitrary, by making this order the same as the order in which these replicas appear in the internal subtour h_{y_i} , we have $\hat{T} = T^*$. Therefore, T^* is a g-tour in G' of a cost no greater than T' 's. ■

Lemma 5 establishes the correspondence between the g-tours in G and those g-tours in G' , each of which consists only of a sequence of internal subtours.

LEMMA 5. *There is a g-tour $u_{0,0} \rightarrow h_{y_1} \rightarrow h_{y_2} \rightarrow \dots \rightarrow h_{y_s} \rightarrow u_{0,0}$ in G' , where $1 \leq s \leq n$, if and only if there is a g-tour $u_{0,0} \rightarrow y_1 \rightarrow y_2 \rightarrow \dots \rightarrow y_s \rightarrow u_{0,0}$ in G of the same cost.*

Proof. The lemma follows from the following two properties: (i) any internal tour in G' has zero cost; and (ii) the cost of any edge (a, b) in G is the same as the cost of any edge in G' that connects a replica of a to a replica of b . ■

Theorem 6 states that any g-tour in G' can be transformed into a g-tour in G of no greater cost, and any optimal g-tour in G' can be transformed into an optimal g-tour in G .

THEOREM 6. *Let $T' = u_{0,0} \rightarrow u_{x_1, i_1} \rightarrow u_{x_2, i_2} \rightarrow \dots \rightarrow u_{x_r, i_r} \rightarrow u_{0,0}$ be a g-tour in G' , where $m \leq r \leq n$. Let s , where $1 \leq s \leq r$, be the number of distinct elements in the sequence x_1, x_2, \dots, x_r , and let y_1, y_2, \dots, y_s be these s distinct elements in the same order as they first appear in the sequence x_1, x_2, \dots, x_r . Then $T = u_{0,0} \rightarrow y_1 \rightarrow y_2 \rightarrow \dots \rightarrow y_s \rightarrow u_{0,0}$ is a g-tour in G of no greater cost. Further, if T' is an optimal g-tour in G' , then T is also an optimal g-tour in G .*

Proof. Let $T^* = u_{0,0} \rightarrow h_{y_1} \rightarrow h_{y_2} \rightarrow \dots \rightarrow h_{y_s} \rightarrow u_{0,0}$. By Lemma 4, T^* is also a g-tour in G' of a cost no greater than T' 's; by Lemma 5, $T = u_{0,0} \rightarrow y_1 \rightarrow y_2 \rightarrow \dots \rightarrow y_s \rightarrow u_{0,0}$ is a g-tour in G of the same cost as T^* 's.

We next show that if T' is an optimal g-tour in G' , then T is an optimal g-tour in G . Assume for contradiction that T' is optimal, but T is not optimal. This then implies that there is some other g-tour $\tilde{T} = u_{0,0} \rightarrow z_1 \rightarrow z_2 \rightarrow \dots \rightarrow z_i \rightarrow u_{0,0}$ in G of a cost less than the cost of T . Thus, by Lemma 5, there is a g-tour $u_{0,0} \rightarrow h_{z_1} \rightarrow h_{z_2} \rightarrow \dots \rightarrow h_{z_i} \rightarrow u_{0,0}$ in G' of the same cost as \tilde{T} 's, which is less than the cost of T' . This contradicts the assumption that T' is optimal. Therefore, T is an optimal g-tour in G . ■

3.2. THE G-S TRANSFORMATION

We now present the transformation of the GTSP with nonintersecting clusters to the TSP. Let G' be a graph with clusters C'_0, C'_1, \dots, C'_m , all of which are nonintersecting. For all $i \in [m]$, let $p_i = |C'_i|$. In this subsection, we use $v_{i,0}, v_{i,1}, \dots, v_{i,p_i-1}$ to denote the nodes in C'_i . We transform G' into a graph $G'' = (V'', E'')$, also with $m+1$ nonintersecting clusters $C''_0, C''_1, \dots, C''_m$. The nodes in V'' are obtained as follows.

- (i) For the node $v_{0,0}$ in C'_0 , we create two nodes $a_{0,0}$ and $c_{0,0}$ in C''_0 .
- (ii) For all $i \in [m]^+$ and all $j \in [p_i - 1]$, corresponding to each node $v_{i,j}$ in C'_i , we create three nodes $a_{i,j}, b_{i,j}$, and $c_{i,j}$ in C''_i ; these nodes are referred to, respectively, as the a -node, b -node, and c -node in G'' corresponding to $v_{i,j}$ in G' .
- (iii) For all $i \in [m]^+$, corresponding to each cluster C'_i in V' , we create a node e_i in C''_i .

The cluster C''_0 consists of the nodes $a_{0,0}$ and $c_{0,0}$, and for all $i \in [m]^+$, the cluster C''_i consists of the node e_i and all the a -nodes, b -nodes, and c -nodes corresponding to the nodes in C'_i . All the clusters are nonintersecting.

We next construct the edges in E'' . Given any two nodes u, v in G'' , we say that u is connected to v if the edge (u, v) is in E'' .

- (i) In the cluster C''_0 , the node $c_{0,0}$ is connected to $a_{0,0}$ with an edge of zero cost.
- (ii) For all $i \in [m]^+$, in the cluster C''_i , all nodes, except e_i , are connected into a cycle $a_{i,0} \rightarrow b_{i,0} \rightarrow c_{i,0} \rightarrow a_{i,1} \rightarrow b_{i,1} \rightarrow c_{i,1} \rightarrow \dots \rightarrow a_{i,p_i-1} \rightarrow b_{i,p_i-1} \rightarrow c_{i,p_i-1} \rightarrow a_{i,0}$; for all $j \in [p_i - 1]$, $b_{i,j}$ is connected to $a_{i,j}$; all the c -nodes are connected to e_i ; and e_i is connected to all the b -nodes. All these edges have zero costs.
- (iii) For all distinct $i, k \in [m]$, all $j \in [p_i - 1]$, and all $l \in [p_k - 1]$, corresponding to every intercluster edge $(v_{i,j}, v_{k,l})$ in E' , there is an intercluster edge $(a_{i,j}, c_{k,l})$ in E'' of the same cost.

Note that as opposed to the graphs G and G' , G'' is not complete. As a result, the distance measure on E'' does not satisfy the triangular inequality.

An example of applying the G-S transformation on the graph G' to obtain the graph G'' is given in Figure 2. To simplify the figure, most of the edges in G' and G'' are omitted. The figure also shows a nonredundant g-tour in G' and its corresponding tour in G'' (the correspondence




Fig. 2. Example of G-S transformation.

plete.
angular

between the nonredundant g-tours in G' and the tours in G'' is established in Theorem 10).

G' to
most of
redundance

Since every intercluster edge in E'' is directed from an a -node to a c -node, every tour in G'' enters a cluster through one of its c -nodes and leaves the cluster through one of its a -nodes.

For all $i \in [m]^+$ and for all $j \in [p_i - 1]$, let $s_{i,j} = a_{i,j} \rightarrow b_{i,j} \rightarrow c_{i,j}$, and let

$$t_{i,j} = c_{i,j} \rightarrow s_{i,(j+1)} \rightarrow s_{i,(j+2)} \rightarrow \dots \rightarrow s_{i,(j+p_i-1)} \rightarrow e_i \rightarrow b_{i,j} \rightarrow a_{i,j},$$

where the addition is modulo p_i . The path $t_{i,j}$ visits all the nodes in C_i'' . We call $t_{i,j}$ the *complete subtour* of the cluster C_i'' at $c_{i,j}$. The cluster C_i'' has p_i distinct complete subtours. For the special case of cluster C_0'' , we define its complete subtour to be $c_{0,0} \rightarrow a_{0,0}$. For example, in the graph G'' shown in Figure 2, the path $c_{1,1} \rightarrow a_{1,0} \rightarrow b_{1,0} \rightarrow c_{1,0} \rightarrow e_1 \rightarrow b_{1,1} \rightarrow a_{1,1}$ is the complete subtour of cluster C_1'' at $c_{1,1}$.

LEMMA 7. *A tour visits every cluster in G'' using one of its complete subtours.*

Proof. Let T'' be an arbitrary tour in G'' . For the clusters C_0'' , since it only has two nodes $a_{0,0}$ and $c_{0,0}$, T'' must visit C_0'' using the path $c_{0,0} \rightarrow a_{0,0}$, which is the unique complete subtour of C_0'' .

We next prove the lemma for all the other clusters. In the following, we use \rightarrow to denote an intracluster edge and \Rightarrow to denote an intercluster edge. Let i be an arbitrary index in $[m]^+$. Let $c_{i,j}$, where $j \in [p_i - 1]$, be the first node in C_i'' visited by T'' . We show that T'' visits all the nodes in C_i'' using the complete subtour at $c_{i,j}$ by establishing the following three claims.

Claim 1. The path $\Rightarrow c_{i,j} \rightarrow \dots \rightarrow e_i \rightarrow b_{i,j} \rightarrow a_{i,j} \Rightarrow \dots$ is in T'' .

Since $c_{i,j}$ is the first node in C_i'' visited by T'' , and since $b_{i,j}$ has only two outgoing edges, one to $a_{i,j}$ and the other to $c_{i,j}$, T'' leaves $b_{i,j}$ to $a_{i,j}$; otherwise, T'' would have visited $c_{i,j}$ at least twice. This, in turn, implies that T'' enters $b_{i,j}$ from e_i because $b_{i,j}$ only has two incoming edges, one from $a_{i,j}$, and the other from e_i . Since $a_{i,j}$ has only one intracluster outgoing edge, which is to $b_{i,j}$, T'' leaves $a_{i,j}$ using an intercluster edge. Finally, since all the edges connected to e_i are intracluster, and all the outgoing edges of $c_{i,j}$ are intracluster, the path $\Rightarrow c_{i,j} \rightarrow \dots \rightarrow e_i \rightarrow b_{i,j} \rightarrow a_{i,j} \Rightarrow \dots$ is in T'' .

Claim 2. T'' enters and exits C_i'' exactly once.

By Claim 1, if T'' enters C_i'' at a particular c -node, then it enters the corresponding b -node from the e -node of the cluster. Since the e -node can only be visited once, the claim follows.

Claim 3. T'' visits C_i'' using the complete subtour at $c_{i,j}$.

For the case that $p_i = 1$, Claim 3 follows trivially from Claim 1. We thus assume that $p_i > 1$. In the following analyses, all arithmetics are modulo p_i . By Claim 1, T'' enters C_i'' at $c_{i,j}$, and exits C_i'' at $a_{i,j}$. By Claim 2, T'' only enters and exits C_i'' once. Therefore, T'' enters all the other c -nodes of the cluster using intracluster edges, and also exits all the other a -nodes of the cluster using intracluster edges. Since an a -node has only one intracluster outgoing edge, which is directed to its corresponding b -node, and a c -node has only one intracluster incoming edge, which is directed from its corresponding b -node, for all $k \in [p_i - 1]$ such that $k \neq j$, the path $a_{i,k} \rightarrow b_{i,k} \rightarrow c_{i,k}$ is in T'' .

Further, for all $k \in [p_i - 1]$, $c_{i,k}$ has only two outgoing edges, both being intracluster, with one to e_i and the other to $a_{i,k+1}$. Since the path $\dots \Rightarrow c_{i,j} \rightarrow \dots \rightarrow e_i \rightarrow b_{i,j} \rightarrow a_{i,j} \Rightarrow \dots$ is in T'' , T'' cannot leave $c_{i,j-1}$ to $a_{i,j}$, but must leave $c_{i,j-1}$ to e_i . This, in turn, implies that for all $k \in [p_i - 1]$ such that $k \neq j - 1$, T'' leaves $c_{i,k}$ to $a_{i,k+1}$. Therefore, T'' visits C_i'' using the complete subtour $c_{i,j} \rightarrow s_{i,j+1} \rightarrow s_{i,j+2} \rightarrow \dots \rightarrow s_{i,j+p_i-1} \rightarrow e_i \rightarrow b_{i,j} \rightarrow a_{i,j}$. ■

The previous lemma implies that every tour in G'' consists only of a sequence of complete subtours. Lemma 8 establishes a one-to-one correspondence between the nonredundant g -tours in G' and the tours in G'' .

LEMMA 8. *There is a nonredundant g -tour $v_{0,0} \rightarrow v_{i_1,j_1} \rightarrow v_{i_2,j_2} \rightarrow \dots \rightarrow v_{i_m,j_m} \rightarrow v_{0,0}$ in G' if and only if there is a tour $a_{0,0} \rightarrow t_{i_1,j_1} \rightarrow t_{i_2,j_2} \rightarrow \dots \rightarrow t_{i_m,j_m} \rightarrow c_{0,0} \rightarrow a_{0,0}$ in G'' of the same cost, where for all $k \in [m]^+$, $i_k \in [m]^+$, and $j_k \in [p_{i_k} - 1]$.*

Proof. The lemma follows from the following two properties: (i) all the intracluster edges in G'' have zero cost; and (ii) the cost of an edge (x, y) in G' is the same as the cost of the corresponding intercluster edge in G'' that connects the a -node of x to the c -node of y . ■

Given a g -tour T' in G' and a node x in T' , we say that x is a *redundant node* in T' if there exists some other node y in T' such that x and y belong to the same cluster in G' .

LEMMA 9. *There exists at least one optimal, nonredundant g -tour in G' .*

Proof. Since G' is a complete graph, there always exists some optimal g -tour T' in G' . If T' is nonredundant, we are finished; otherwise, we want to show that we can always construct an optimal, nonredundant g -tour from T' . Assume that T' is redundant. Let C_i' , where $i \in [m]^+$, be a

cluster in G' that is visited by T' at least twice. We have $T' = v_{0,0} \rightarrow \dots \rightarrow v_{i,j} \rightarrow \dots \rightarrow u \rightarrow v_{i,k} \rightarrow w \rightarrow \dots \rightarrow v_{0,0}$, where $j, k \in [p_i - 1]$, and $u, w \in V'$ (u may equal to $v_{i,j}$ and w may equal to $v_{0,0}$). Let $T^* = v_{0,0} \rightarrow \dots \rightarrow v_{i,j} \rightarrow \dots \rightarrow u \rightarrow w \rightarrow \dots \rightarrow v_{0,0}$ such that, except for omitting the node $v_{i,k}$, T^* visits the same sequence of nodes as T' . T^* is also a g-tour in G' , with one less redundant node than T' . Since the distance measure on G' satisfies the triangular inequality (Lemma 1), T^* has a cost no greater than T' 's. Further, since T' is optimal, T' and T^* have the same cost, implying that T^* is also optimal. By applying the procedure described above repeatedly for $r - (m + 1)$ times, where r is the length of T' , we can construct a sequence of optimal g-tours, each of which visits one less redundant node than its predecessor in the sequence, with the last one in the sequence being nonredundant. ■

THEOREM 10. *There is a one-to-one correspondence between the nonredundant g-tours in G' and the tours in G'' , with the tour and the nonredundant g-tour in any corresponding pair having the same cost. Furthermore, an optimal g-tour in G' can be obtained from an optimal tour in G'' .*

Proof. By Lemma 7, every tour in G'' consists only of a sequence of complete subtours. Thus, by Lemma 8, we establish a one-to-one correspondence between the tours of G'' and the nonredundant tours of G' , with the tour and the nonredundant g-tour in any corresponding pair having the same cost. Further, by Lemma 9, there exists at least one nonredundant g-tour in G' that is optimal. Therefore, an optimal g-tour in G' can be obtained from an optimal tour in G'' . ■

4. SUMMARY

By the I-N transformation and the G-S transformation presented in this paper, given an instance G of the GTSP with intersecting clusters, we can transform G into an instance G'' of the TSP such that any tour in G'' can be transformed into a g-tour in G of no greater cost, and any optimal tour in G'' can be transformed into an optimal g-tour in G .

REFERENCES

1. G. B. Dantzig, D. R. Fulkerson, and S. Johnson, Solution of a large-scale traveling-salesman problem, *Oper. Res.* 2:393-410 (1954).
2. S. S. Srivastava, S. Kumar, R. C. Garg, and P. Sen, Generalized traveling-salesman problem through n sets of nodes, *J. Canadian Oper. Res. Soc.*: 97-101 (1970).

Received 5 January 1988; revised 10 June 1988