Information Bounds Are Weak in the Shortest Distance Problem

RONALD L. GRAHAM

Bell Laboratories, Murray Hill, New Jersey

AND

ANDREW C. YAO AND F. FRANCES YAO

Stanford University, Stanford, California

ABSTRACT. In the all-pair shortest distance problem, one computes the matrix $D = (d_{ij})$, where d_{ij} is the minimum weighted length of any path from vertex i to vertex j in a directed complete graph with a weight on each edge. In all the known algorithms, a shortest path p_{ij} achieving d_{ij} is also implicitly computed. In fact, $\log_3(f(n))$ is an information-theoretic lower bound, where $f(n)$ is the total number of distinct *patterns* (p_{ij}) for n-vertex graphs. As $f(n)$ potentially can be as large as 2^{n^3}, it would appear possible that a nontrivial lower bound can be derived this way in the decision tree model. The characterization and enumeration of realizable patterns is studied, and it is shown that $f(n) \leq C^{n^2}$. Thus no lower bound greater than Cn^2 can be derived from this approach. It is proved as a corollary that the Triangular polyhedron $T^{(n)}$, defined in $E^{\binom{n}{2}}$ by $d_{ij} \geq 0$ and the triangle inequalities $d_{ij} + d_{jk} \geq d_{ik}$, has at most C^{n^2} faces of all dimensions, thus resolving an open question in a similar information bound approach to the shortest distance problem.

KEY WORDS AND PHRASES: decision tree model, Farkas Lemma, information bound, lower bound, maximum flow, polyhedron, shortest distance, shortest path

CR CATEGORIES: 5.3, 8.3

1. *Introduction*

Let G be a directed complete graph on n vertices $v_1, v_2, \ldots, v_n$, with a nonnegative distance d_{ij} associated with each edge (v_i, v_j). In the *all-pair shortest distance problem*, one wishes to compute the $n \times n$ *shortest distance matrix* $D^* = (d_{ij}^*)$, where d_{ij}^* is the minimum total length of any path from v_i to v_j (see, e.g., [1]). Efficient algorithms for this problem were devised by Dantzig [2], Dijkstra [3], and Floyd [5]. All these methods require at least Cn^3 time in the worst case. More recently, Fredman [6] gave an algorithm with running time $O(n^3(\log \log n/\log n)^{1/3})$, which is slightly better than $O(n^3)$. Substantial improvements over $O(n^3)$, however, are yet to be found. On the other hand, no lower bound better than Cn^2 is known to the all-pair shortest paths problem for programs with branching instructions. (Kerr [9] proved that Cn^3 steps are necessary for straightline programs with operations $\{\min, +\}$.)

A natural model incorporating branching instructions is the *decision tree model*, which is used, for example, in the study of many sorting type problems [10]. Indeed, all the

This research was supported in part by the National Science Foundation under Grants MCS-72-03752 A03 and MCS-77-05313.

Authors' addresses: R.L. Graham, Bell Laboratories, Murray Hill, NJ 07974; A.C. Yao and F.F. Yao, Computer Science Department, Stanford University, Stanford, CA 94305.

existing shortest paths algorithms mentioned above can be properly modeled by *linear decision trees*, where the primitives are ternary comparisons "$f(\{d_{ij}\}) \gtreqless 0$" with linear functions f. An apparently promising approach to obtaining lower bounds for linear decision trees was suggested by Yao, Avis, and Rivest [13]. It was shown that, in this model, $Cn^2\log n$ comparisons are necessary to compute the shortest distance matrix if a certain polyhedron $T^{(n)}$ in $\binom{n}{2}$-dimensional Euclidean space (see Section 2.3) has at least $\exp(Cn^2\log n)$ "edges," i.e., one-dimensional faces.[1] An interesting question is thus to determine if $T^{(n)}$ in fact has that many edges.

While counting the number of comparisons made in a decision tree tends to underestimate the "true" complexity of computing shortest distances (for example, Fredman [6] showed that for any given n, there exists a linear decision tree with $O(n^{2.5})$ comparisons), it seems to be at present the only hope for obtaining nontrivial lower bounds. In this paper we examine an approach based on information-theoretic arguments. As will become clear, a natural information lower bound is $\log_3 |P(n)| - n^2$, where $P(n)$ is defined as follows: For any $n \times n$ matrix $D = (d_{ij})$ with nonnegative entries, let *pattern*(D) denote the $n \times n$ matrix (p_{ij}), where p_{ij} is the set of all shortest paths from vertex v_i to v_j in the graph G associated with D. We define $P(n)$ to be the collection of all distinct patterns obtainable this way. As the cardinality of $P(n)$ is potentially large ($O(2^{n^2\log n})$, even if we require each p_{ij} to consist of a unique path), it appears hopeful that strong lower bounds could be established. However, we will show that in fact $\log|P(n)| = O(n^2)$; therefore no lower bounds better than Cn^2 can be derived from this approach. The enumeration of $P(n)$ is based on a study of "connection matrices," as described in the next paragraph.

Let $D = (d_{ij})$, $D' = (d'_{ij})$ be two $n \times n$ matrices with nonnegative entries. Then the *connection matrix* $C_{D,D'}$ for D and D' has as entries

$$C_{D,D'}[i, j] = \{\alpha \mid 1 \le \alpha \le n,\ d_{i\alpha} + d'_{\alpha j} = \min_k (d_{ik} + d'_{kj})\} \qquad \text{for} \quad 1 \le i, j \le n.$$

In Sections 2 through 5 we develop characterizations for $R(n)$, the set of all "realizable" connection matrices. As a result, $|R(n)|$ is shown to be of the order C^{n^2} (here again, rather short of its 2^{n^2} potential). In Section 6 we apply the scheme used in [1, p. 204] for reducing shortest distances computation to {min, +} matrix multiplication to establish a recurrence relation involving $|R(n)|$ and $|P(n)|$ and thereby show that $|P(n)| \le C^{n^2}$.

In another application of the concept of connection matrices, we show that, somewhat unexpectedly, each face of the polyhedron $T^{(n)}$ mentioned earlier corresponds naturally to a unique $n \times n$ connection matrix (see Section 2.3). Therefore $T^{(n)}$ has no more than C^{n^2} edges, which resolves the question in the polyhedron approach [13] as well.

2. *Connection Matrix, Information Bounds, and Triangular Polyhedron*

2.1 THE {MIN, +} MATRIX MULTIPLICATION. A *distance matrix* is a matrix of nonnegative real numbers. For two $n \times n$ distance matrices $D = (d_{ij})$ and $D' = (d'_{ij})$, define their *sum* $A = (a_{ij}) = D \oplus D'$ and *product* $B = (b_{ij}) = D \otimes D'$, respectively, by $a_{ij} = \min\{d_{ij}, d'_{ij}\}$ and $b_{ij} = \min\{d_{ik} + d'_{kj} \mid 1 \le k \le n\}$. The multiplicative operation $\otimes$ is also called the {*min*, +} *matrix multiplication*. It is well known [1, 4, 11] that the complexity of {min, +} matrix multiplication is closely related to that of finding all-pair shortest distances, i.e., computing the *transitive closure* $D^* = (d^*_{ij})$ of a matrix D, where $d^*_{ii} = 0$ and $d^*_{ij} = (D \oplus D^2 \oplus D^3 \oplus \cdots)_{ij}$ for $i \ne j$. ($D^i = D^{i-1} \otimes D$ by definition.) We first focus attention on the {min, +} matrix multiplication for its conceptual simplicity. The discussions are then extended to the computation of shortest distances in Section 6.

We consider the computation of {min, +}-product for two $n \times n$ matrices in the *decision tree model*. An algorithm in this model is a ternary tree. Each internal node contains a test "$f(D, D') : 0$" for some nonconstant rational function f of $2n^2$ arguments. Each leaf of the

[1] It was incorrectly claimed in [13] that $T^{(n)}$ could be shown to have $\exp(Cn^2\log n)$ edges, which would then imply the $\Omega(n^2\log n)$ lower bound. A revised version of [13] will appear as [14].

tree contains a set of rational functions $\{q_{ij}, i \le i, j \le n\}$ on the $2n^2$ variables $\{d_{ij}, d'_{ij}\}$. For any input (D, D'), the algorithm moves from the root down the tree, at each node testing and then branching according to whether $f(D, D')$ is >0, $=0$, or <0, until a leaf is reached. At that point the product $B = D \otimes D'$ is given by $b_{ij} = q_{ij}(D, D')$. The cost of the algorithm is defined to be the height of the tree. The complexity $L(n)$ in this model is the minimum cost over all such algorithms. When all the functions f, q_{ij} are restricted to be linear functions, the model is called the *linear decision tree model*, and the corresponding complexity is denoted by $L_0(n)$. Trivially, $L(n) \le L_0(n)$.

We shall be interested in a natural information-theoretic bound on $L(n)$ and $L_0(n)$.

2.2 CONNECTION MATRICES AND INFORMATION BOUNDS. The concept of a connection matrix has been defined in Section 1. We now give some illustrations and examine the relationship between connection matrices and {min, +}-multiplication.

Consider the following interpretation of the product $B = (b_{ij}) = D \otimes D'$ (see, e.g., [1]): Let $X = \{x_1, x_2, \ldots, x_n\}$, $Y = \{y_1, y_2, \ldots, y_n\}$, and $Z = \{z_1, z_2, \ldots, z_n\}$ be three disjoint sets of cities, with d_{ik} and d'_{kj} being the distances from x_i to y_k and from y_k to z_j, respectively. Then b_{ij} is the "shortest distance" from x_i to z_j via some intermediate city in Y. This suggests another way of representing the product $D \otimes D'$, namely, we can list for each pair $[i, j]$ the set of all connecting cities y_k for which $d_{ik} + d'_{kj}$ achieves the minimum b_{ij}. Such information can be tabulated into an $n \times n$ matrix $C_{D,D'}$, whose $[i, j]$-entry is the *set* of integers $\{\alpha \mid d_{i\alpha} + d'_{\alpha j} = \min_k(d_{ik} + d'_{kj})\}$. Clearly $C_{D,D'}$ is the *connection matrix* for D and D', as defined earlier.

Example 1. For the graph shown in Figure 1, we have $D = \left(\begin{smallmatrix} 20 & 30 \\ 10 & 15 \end{smallmatrix}\right)$ and $D' = \left(\begin{smallmatrix} 15 & 20 \\ 10 & 10 \end{smallmatrix}\right)$. The connection matrix $C_{D,D'}$ is $\left(\begin{smallmatrix} 1 & 1, 2 \\ 1, 2 & 2 \end{smallmatrix}\right)$.

Not all matrices can be realized as connection matrices for some D and D'. For example, there do not exist 2×2 distance matrices D and D' whose connection matrix $C_{D,D'}$ is $\left(\begin{smallmatrix} 1 & 2 \\ 2 & 1 \end{smallmatrix}\right)$. For if we assume that $C_{D,D'} = \left(\begin{smallmatrix} 1 & 2 \\ 2 & 1 \end{smallmatrix}\right)$ for some $D = (d_{ij})$ and $D' = (d'_{ij})$, we have then four inequalities:

$$d_{11} + d'_{11} < d_{12} + d'_{21},$$
$$d_{12} + d'_{22} < d_{11} + d'_{12},$$
$$d_{22} + d'_{21} < d_{21} + d'_{11},$$
$$d_{21} + d'_{12} < d_{22} + d'_{22}.$$

Adding the above four inequalities together, one obtains $0 < 0$, a contradiction.

Definition 1. An *n-ary matrix* M is a matrix where each entry $M[i, j]$ is a subset of $\{1, 2, \ldots, n\}$. An n-ary matrix is said to be *simple* if $|M[i, j]| = 1$ for all i, j.

A connection matrix $C_{D,D'}$ is an n-ary matrix of dimension $m \times p$ if D and D' have dimensions $m \times n$ and $n \times p$, respectively. For simplicity, we will only consider the case $m = p = n$, while noting that all discussions have immediate generalizations to rectangular matrices. Thus, when there is no danger of confusion, an $n \times n$ n-ary matrix will simply be called an n-ary matrix.

As illustrated in the discussion above, not all of the 2^{n^3} $n \times n$ n-ary matrices are connection matrices.

Definition 2. An n-ary matrix M is said to be *realizable* (as a connection matrix) if $M = C_{D,D'}$ for some distance matrices D, D'. Let $R(n)$ denote the family of all $n \times n$ realizable n-ary matrices M.

A subfamily of $R(n)$ deserves special attention.

Definition 3. Let $SR(n)$ be the subset of $R(n)$ consisting of all simple n-ary matrices.

We now give lower bounds to the complexity of {min, +}-multiplication in terms of $|R(n)|$ and $|SR(n)|$. It is plausible that to compute the shortest distance between x_i and z_j,

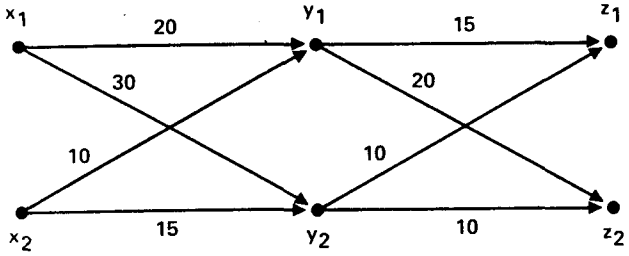

FIG. 1. An example of a connection matrix.

one has to find the best connecting cities y_k. Thus there must be as many leaves as $|R(n)|$ (or $|SR(n)|$) in a decision tree. The logarithm of the number of leaves then gives a lower bound to the height of a tree, which is usually referred to as the information-theoretic bound.

THEOREM 1. *$L(n) \geq \log_2|SR(n)|$ for all $n \geq 1$.*

PROOF. Let A be any decision tree algorithm computing the {min, +}-product of $n \times n$ matrices $D \otimes D'$. Let $\mathcal{D}$ be the set of input pairs (D, D') with all their entries strictly positive and for which the test result is never zero at any internal point, i.e., $\prod_{i \in A} f_i(D, D') \neq 0$, where f_i are the test functions at internal node i. Clearly $\mathcal{D}$ is an open set in the Euclidean space E^{2n^2}, and is dense in the positive quadrant (all coordinates ≥ 0). For each element $M \in SR(n)$, choose D_M, D'_M such that $C_{D_M, D'_M} = M$ and $(D_M, D_{M'}) \in \mathcal{D}$, which can be done since, for any distance-matrix pair (D, D') with $C_{D,D'} = M$, all $(D_M, D'_M) \in \mathcal{O} \cap \mathcal{D}$ satisfy $C_{D_M, D'_M} = M$, where $\mathcal{O}$ is a sufficiently small neighborhood of (D, D') in E^{2n^2}. For any such $(D_M, D_{M'})$, the computation will end at some leaf l_M without taking an equality branch at any internal node. Let $M[i, j] = \{k_{ij}\}$; then in some sufficiently small open set $\mathcal{O} \subseteq \mathcal{D}$ around (D_M, D'_M), the shortest distance from x_i to z_j $(1 \leq i, j \leq n)$ is through $y_{k_{ij}}$ uniquely for each $(D, D') \in \mathcal{O}$, and furthermore, every $(D, D') \in \mathcal{O}$ leads to the same leaf l_M. Since two rational functions agreeing in an open set must be identical, we know that the set of output functions $\{q_{ij}\}$ at l_M must be $q_{ij}(D, D') = d_{i,k_{ij}} + d'_{k_{ij},j}$. It follows that no two distinct $M \in SR(n)$ can have the same l_M. Now if we prune all the equality branches from the tree A, we have a binary tree with at least $|SR(n)|$ leaves. The height of A is therefore at least $\log_2|SR(n)|$, which implies $L(n) \geq \log_2|SR(n)|$. □

The above argument does not apply when $SR(n)$ is replaced by $R(n)$, since for $M \in R(n)$, the set of (D, D') satisfying $C_{D,D'} = M$ in general does not contain an open set. However, in the more restricted model of linear decision trees, $R(n)$ does provide a lower bound.

THEOREM 2. *$L_0(n) \geq \log_3|R(n)| - 2n^2$.*

PROOF. Let A be an optimal linear decision tree for computing the $n \times n$ matrix product $D \otimes D'$. Consider the algorithm A' which begins with a sequence of $2n^2$ tests $\{d_{ij}:0, d'_{ij}:0, 1 \leq i, j \leq n\}$, and then proceeds exactly as algorithm A, ignoring the outcomes of the first $2n^2$ tests. Represented as a linear decision tree, the algorithm A' has height $L_0(n) + 2n^2$. We will show that, for algorithm A', all input pairs of distance matrices (D, D') reaching the same leaf must have the same connection matrix $C_{D,D'}$. This will prove $L_0(n) + 2n^2 \geq \log_3|R(n)|$, hence the theorem.

Let l be any leaf with output functions $\{q_{ij}\}$. Let $\mathcal{L} = \{g_1 < 0, g_2 < 0, \ldots, g_s < 0, h_1 = 0, h_2 = 0, \ldots, h_t = 0\}$ be the system of linear inequalities and equalities obtained along the path from the root to l. Then for any $1 \leq i, j, k \leq n$, $q_{ij}(D, D') \leq d_{ik} + d'_{kj}$ must be a consequence of the system $\mathcal{L}$. Because of the Farkas Lemma (for inhomogeneous systems; see, e.g., [12, Theorem 1.4.4]), one can obtain $q_{ij}(D, D') \leq d_{ik} + d'_{kj}$ by taking convex linear combinations of formulas in the system $\mathcal{L} \cup \{0 < 1\}$. But this process

actually yields either "<" or "=" explicitly. Thus we actually know at leaf l whether $q_{ij}(D, D') < d_{ik} + d'_{kj}$ or $q_{ij}(D, D') = d_{ik} + d'_{kj}$ for all i, j, k. This proves that the connection matrix is determined at each leaf reachable by inputs, as was to be shown.[2] □

We regard the two preceding theorems as information bounds on $L(n)$ and $L_0(n)$, respectively. As there are n^{n^2} simple n-ary matrices, and 2^{n^3} n-ary matrices, of which $SR(n)$ and $R(n)$ are subsets, respectively, Theorems 1 and 2 could potentially give lower bounds of the order $n^2 \log n$ or higher. The characterization and enumeration of $SR(n)$ and $R(n)$ will be the subject of Sections 3 through 5. First we define the Triangular polyhedron $T^{(n)}$ and relate it to our present approach.

2.3 THE TRIANGULAR POLYHEDRON $T^{(n)}$. A set Z in E^N is a *polyhedron* if $Z = \{\vec{x} \mid \vec{x} \in E^N, l_i(\vec{x}) \le 0, i = 1, 2, \ldots, m\}$, where m is an integer, $\vec{x} = (x_1, x_2, \ldots, x_N)$, and $l_i(\vec{x}) = \sum_{1 \le j \le n} c_{ij}x_j - c'_i$ for real numbers c_{ij}, c'_i. To each subset $J \subseteq \{1, 2, \ldots, m\}$ (possibly empty), let $F_J(Z) = \{\vec{x} \mid l_i(\vec{x}) < 0$ for each $i \in J$; $l_i(\vec{x}) = 0$ for each $i \notin J\}$. We call $F_J(Z)$ a *face of dimension* t of Z if $F_J(Z) \neq \emptyset$ and the smallest subspace of E^N containing $F_J(Z)$ has dimension t. Let $\mathcal{F}_t(Z)$ be the set of faces of dimension t of Z for $1 \le t \le N$. (For more information on polyhedra, faces, etc., see [7, 12].)

The *Triangular polyhedron* $T^{(n)}$ is a polyhedron in E^N for $N = \binom{n}{2}$. Let $\Pi = \{(i, j) \mid 1 \le i < j \le n\}$ and $\Sigma = \{(i, j, k) \mid (i, j) \in \Pi,\ 1 \le k \le n$ and $k \neq i, k \neq j\}$. Write a vector in E^N as $\vec{x} = (x_{ij}, (i, j) \in \Pi)$. Then $T^{(n)}$ is defined by

$$T^{(n)} = \{\vec{x} \mid x_{ij} \ge 0 \text{ for } (i, j) \in \Pi,\ x_{ij} \le x_{ik} + x_{kj} \text{ for } (i, j, k) \in \Sigma\},$$

where[3] we interpret x_{ik} to be x_{ki} if $i > k$.

THEOREM 3. $|\bigcup_{t=0}^{N} \mathcal{F}_t(T^{(n)})| \le |R(n)|$, *where* $N = \binom{n}{2}$.

COROLLARY. $|\mathcal{F}_1(T^{(n)})| \le |R(n)|$.

PROOF. It suffices to establish a one-to-one mapping φ from $\bigcup_{t=0}^{N} \mathcal{F}_t(T^{(n)})$, i.e., the set of all faces of $T^{(n)}$, into $R(n)$. Write $l_{ijk}(\vec{x}) = x_{ij} - x_{ik} - x_{kj}$ for $(i, j, k) \in \Sigma$. Let F be a face of $T^{(n)}$, specified by a partition of Π into $\Pi_1 \cup \Pi_2$, Σ into $\Sigma_1 \cup \Sigma_2$, such that

$$F = \{\vec{x} \mid x_{ij} > 0 \text{ if } (i, j) \in \Pi_1,\ l_{ijk} < 0 \text{ if } (i, j, k) \in \Sigma_1,$$
$$\text{and } x_{ij} = 0 \text{ if } (i, j) \in \Pi_2,\ l_{ijk} = 0 \text{ if } (i, j, k) \in \Sigma_2\}.$$

We now define $\varphi(F)$ to be the $n \times n$ n-ary matrix M, given by

$$M[i, j] = M[j, i] = \{k \mid (i, j, k) \in \Sigma_2\} \cup \{i, j\} \quad \text{if } i < j,$$

and

$$M[i, i] = \{k \mid \{(i, k), (k, i)\} \cap \Pi_2 \neq \emptyset\} \cup \{i\}.$$

The mapping φ is one-to-one, as Σ_2 and Π_2 can be reconstructed from $\varphi(F)$.

To complete the proof of the theorem, it remains to show that $\varphi(F)$ defines a realizable matrix M. Choose $\vec{x} = (x_{ij}, 1 \le i < j \le n)$ to be any point on F. Define a distance matrix $D = (d_{ij})$ from $\vec{x}$ by letting

$$d_{ij} = d_{ji} = x_{ij} \quad \text{for } 1 \le i < j \le n,$$

and

$$d_{ii} = 0 \quad \text{for } 1 \le i \le n.$$

It is easy to check that $D \otimes D = D$. It follows that the connection matrix $C_{D,D}$ is given by

$$C_{D,D}[i, j] = C_{D,D}[j, i] = \{k \mid l_{ijk}(\vec{x}) = 0,\ 1 \le k \le n\} \cup \{i, j\} \quad \text{if } i < j,$$

[2] We introduced A' in the proof for the following reason: The system of constraints $\mathcal{L}$ at a leaf of A contains the $2n^2$ "mixed" constraints $d_{ij} \ge 0$, $d'_{ij} \ge 0$ which are neither equalities nor strict inequalities.

[3] David Avis pointed out that the conditions $x_{ij} \ge 0$ are implied by the other conditions in the definition of $T^{(n)}$.

and

$$C_{D,D}[i, i] = \{k \mid x_{ik} = 0 \text{ or } x_{ki} = 0,\ 1 \le k \le n\} \cup \{i\}.$$

This proves that $\varphi(F) = M = C_{D,D}$. The proof of the theorem is complete. □

3. *A Characterization of Simple Connection Matrices*

We will give a necessary and sufficient condition for a simple n-ary matrix to be a connection matrix. We first define some useful concepts.

Definition 4. The *weight distribution* $W(M)$ of an n-ary matrix M is the integer matrix defined by $W(M)_{i,j} = |M[i, j]|$. The sum $\sum_{i,j} |M[i, j]|$ is called the *total weight* of M, denoted by $w(M)$.

Example 2. Let

$$M = \begin{pmatrix} 3 & 1, 2 & 2, 3 \\ 1 & 1 & 2 \\ 1, 2, 3 & 3 & 2 \end{pmatrix}.$$

The weight distribution of M is

$$W(M) = \begin{pmatrix} 1 & 2 & 2 \\ 1 & 1 & 1 \\ 3 & 1 & 1 \end{pmatrix},$$

with total weight $w(M) = 13$.

Definition 5. Let M be an n-ary matrix of dimension $m \times p$. For $1 \le i \le m$, the *ith row signature* of M is the vector $\vec{r}^{(i)} = (r_1^{(i)}, r_2^{(i)}, \ldots, r_n^{(i)})$, where $r_l^{(i)}$ is the number of times integer l appears in the ith row. For $1 \le j \le p$, the *jth column signature* $\vec{c}^{(j)} = (c_1^{(j)}, c_2^{(j)}, \ldots, c_n^{(j)})$ of M is defined in a similar way, i.e., $c_l^{(j)}$ is the number of occurrences of l in the jth column. The sequence of $m + p$ vectors $\vec{r}^{(1)}, \vec{r}^{(2)}, \ldots, \vec{r}^{(m)}, \vec{c}^{(1)}, \vec{c}^{(2)}, \ldots, \vec{c}^{(p)}$) is then called the *signature of* M, denoted by $s(M)$.

In Example 2 above, the row signatures of M are $\vec{r}^{(1)} = (1, 2, 2)$, $\vec{r}^{(2)} = (2, 1, 0)$, and $\vec{r}^{(3)} = (1, 2, 2)$; the column signatures are $\vec{c}^{(1)} = (2, 1, 2)$, $\vec{c}^{(2)} = (2, 1, 1)$, and $\vec{c}^{(3)} = (0, 3, 1)$.

Definition 6. An n-ary simple matrix M is said to be *s-unique* if no other n-ary simple matrix M' can have the same signature as M.

We will show that, for a simple n-ary matrix M, the property of s-uniqueness is the answer to the question of whether M is realizable as a connection matrix.

THEOREM 4. *Let M be an $n \times n$ simple n-ary matrix. Then $M \in SR(n)$ if and only if M is s-unique.*

PROOF

Necessity. Let M be a simple n-ary matrix such that $M = C_{D,D'}$ for distance matrices $D = (d_{ij})$ and $D' = (d'_{ij})$. Assume that there exists another simple n-ary matrix $M' \ne M$ with $s(M') = s(M)$. We will show that this leads to a contradiction.

Write $M = (m_{ij})$ and $M' = (m'_{ij})$. We have

$$d_{i,m_{ij}} + d'_{m_{ij},j} \le d_{i,m'_{ij}} + d'_{m'_{ij},j} \qquad \text{for} \quad 1 \le i, j \le n, \tag{1}$$

by the definition of the connection matrix $C_{D,D'}$. Furthermore, the inequality (1) is strict if $m_{ij} \ne m'_{ij}$. Adding up the n^2 inequalities in (1) we obtain

$$\sum_i \sum_j d_{i,m_{ij}} + \sum_j \sum_i d'_{m_{ij},j} < \sum_i \sum_j d_{i,m'_{ij}} + \sum_j \sum_i d'_{m'_{ij},j}, \tag{2}$$

where the inequality is strict since $m_{ij} \ne m'_{ij}$ for some i, j. Now, by the definition of the row

and column signatures $\vec{r}^{(i)}$, $\vec{c}^{(j)}$ of M and $\vec{r}'^{(i)}$, $\vec{c}'^{(j)}$ of M', respectively, (2) is equivalent to

$$\sum_i \sum_l r_l^{(i)} d_{il} + \sum_j \sum_l c_l^{(j)} d'_{lj} < \sum_i \sum_l r_l'^{(i)} d_{il} + \sum_j \sum_l c_l'^{(j)} d'_{lj}. \tag{3}$$

But by assumption M and M' have the same signature, so the left-hand side of (3) is equal to the right-hand side, which is a contradiction. This proves the necessity of s-uniqueness for a simple connection matrix.

Sufficiency. We next show that if a simple n-ary matrix M is s-unique, then there exist distance matrices D and D' such that $M = C_{D,D'}$. What we look for are $D = (d_{ij})$ and $D' = (d'_{ij})$ that satisfy the following system of inequalities:

$$(\mathscr{S}) \begin{cases} g_{i,j,\alpha,\beta}(D, D') = (d_{i\alpha} + d'_{\alpha j}) - (d_{i\beta} + d'_{\beta j}) < 0 \\ \qquad\qquad \text{for} \quad \alpha = m_{ij}, \quad \beta \neq \alpha, \quad 1 \le i, j \le n, \\ h_{i,j,\alpha,\alpha}(D, D') = (d_{i\alpha} + d'_{\alpha j}) - (d_{i\alpha} + d'_{\alpha j}) = 0 \\ \qquad\qquad \text{for} \quad \alpha = m_{ij}, \quad 1 \le i, j \le n. \end{cases}$$

Assume that the system $(\mathscr{S})$ has no solution. We will show that this implies M is not s-unique. First note that $(\mathscr{S})$ contains at least one strict inequality $g_{i,j,\alpha,\beta} < 0$, for $n \ge 2$. By the theorem of Kuhn–Fourier (see [12, Theorem 1.1.9]), $(\mathscr{S})$ is not solvable only if there exist nonnegative numbers $\lambda_{i,j,\alpha,\beta}$ such that

$$\sum_{\substack{1\le i,j\le n \\ \alpha = m_{ij} \\ \beta \neq \alpha}} \lambda_{i,j,\alpha,\beta}\, g_{i,j,\alpha,\beta} + \sum_{\substack{1\le i,j\le n \\ \alpha = m_{ij}}} \lambda_{i,j,\alpha,\alpha}\, h_{i,j,\alpha,\alpha}$$

$$= (0 \cdot d_{11} + \cdots + 0 \cdot d_{ij} + \cdots + 0 \cdot d_{nn}) + (0 \cdot d'_{11} + \cdots + 0 \cdot d'_{ij} + \cdots + 0 \cdot d'_{nn}), \tag{4}$$

where $\lambda_{i,j,\alpha,\beta} > 0$ for the coefficient of some $g_{i,j,\alpha,\beta}$. We can scale the coefficients in (4) so that every λ is $\le 1/n$, except for $\lambda_{i,j,\alpha,\alpha}$. The values of $\lambda_{i,j,\alpha,\alpha}$ ($1 \le i, j \le n$, $\alpha = m_{ij}$) can be chosen freely in (4) since $h_{i,j,\alpha,\alpha} = 0$, and we shall choose them so that for any fixed i, j, and $\alpha = m_{ij}$,

$$\sum_{1\le\beta\le n} \lambda_{i,j,\alpha,\beta} = 1. \tag{5}$$

Let us rewrite (4) as

$$\sum_{\substack{1\le i,j\le n \\ \alpha = m_{ij}}} \sum_{1\le\beta\le n} \lambda_{i,j,\alpha,\beta}(d_{i\alpha} + d'_{\alpha j}) = \sum_{\substack{1\le i,j\le n \\ \alpha = m_{ij}}} \sum_{1\le\beta\le n} \lambda_{i,j,\alpha,\beta}(d_{i\beta} + d'_{\beta j}). \tag{6}$$

By eq. (5), the left-hand side of (6) is

$$\sum_{\substack{1\le i,j\le n \\ \alpha = m_{ij}}} (d_{i\alpha} + d'_{\alpha j}),$$

or, equivalently,

$$\sum_{1\le i\le n} \sum_{1\le l\le n} r_l^{(i)} d_{il} + \sum_{1\le j\le n} \sum_{1\le l\le n} c_l^{(j)} d'_{lj}, \tag{7}$$

where $r_l^{(i)}$, $c_l^{(j)}$ are the row and column signatures of M. By comparing the coefficient of each variable d_{il}, d'_{jl} in (7) with that in the right-hand side of (6), we obtain

$$\sum_{\substack{1\le j\le n \\ \alpha = m_{ij}}} \lambda_{i,j,\alpha,l} = r_l^{(i)} \qquad \text{for} \quad 1 \le i \le n, \quad 1 \le l \le n, \tag{8}$$

$$\sum_{\substack{1\le i\le n \\ \alpha = m_{ij}}} \lambda_{i,j,\alpha,l} = c_l^{(j)} \qquad \text{for} \quad 1 \le j \le n, \quad 1 \le l \le n. \tag{9}$$

The equalities in (8) and (9) are best represented in terms of a network flow problem. Let $\mathscr{N}(M)$ be a network with source S, sink T, and, in between, three levels of nodes, with n^2 nodes on each level (Figure 2). The nodes on the first level are $R_l^{(i)}$ ($1 \le i, l \le n$), on the

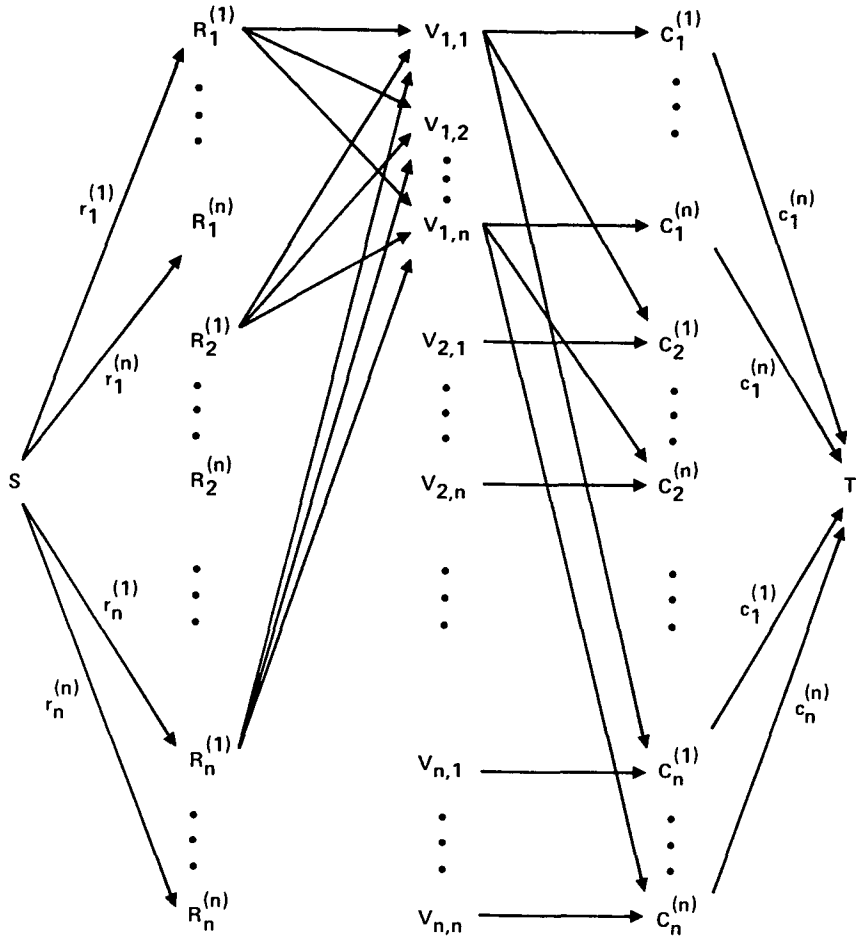

FIG. 2. Network $\mathcal{N}(M)$.

second level V_{ij} $(1 \le i, j \le n)$, and on the third level $C_l^{(j)}$ $(1 \le j, l \le n)$. Each $R_l^{(i)}$ is connected with the source and the n nodes V_{ij} $(1 \le j \le n)$; each $C_l^{(j)}$ is connected with the sink and the n nodes V_{ij} $(1 \le i \le n)$. We shall consider maximum flows in $\mathcal{N}(M)$, subject to the following capacity constraints on the nodes (cf. [7]): Node $R_l^{(i)}$ has capacity $r_l^{(i)}$, node $C_l^{(j)}$ has capacity $c_l^{(j)}$, and node V_{ij} has capacity 1.

The value of a maximum flow in $\mathcal{N}(M)$ is clearly at most $\sum_i \sum_l r_l^{(i)} = \sum_j \sum_l c_l^{(j)} = n^2$, if all nodes are saturated to their capacities. We will demonstrate two flow functions y^* and $\bar{y}$ that can achieve this maximum. Each function assigns the same value to both arcs $(R_l^{(i)}, V_{ij})$ and $(V_{ij}, C_l^{(j)})$. We denote this value by $y^*(i, j, l)$ and $\bar{y}(i, j, l)$, respectively.

In the first maximum flow y^*, we let

$$y^*(i, j, l) = \begin{cases} 1 & \text{if } l = m_{ij}, \\ 0 & \text{otherwise.} \end{cases} \tag{10}$$

There is one unit of flow through every node V_{ij}. Furthermore, each node $R_l^{(i)}$, $C_l^{(j)}$ is balanced and saturated by definition of the capacities $r_l^{(i)}$, $c_l^{(j)}$.

The other flow function $\bar{y}$ makes the assignment

$$\bar{y}(i, j, l) = \lambda_{i,j,\alpha,l}, \tag{11}$$

where $\alpha = m_{ij}$. The amount of flow through V_{ij} is

$$\sum_{1 \le l \le n} \bar{y}(i, j, l) = 1$$

by eq. (5). The total flow out of node $R_l^{(i)}$ is

$$\sum_{1\le j\le n} \bar{y}(i, j, l) = \sum_{\substack{1\le j\le n \\ \alpha = m_{ij}}} \lambda_{i,j,\alpha,l} = r_l^{(i)}$$

by eq. (8); similarly the total flow into node $C_l^{(j)}$ is

$$\sum_{1\le i\le n} \bar{y}(i, j, l) = \sum_{\substack{1\le i\le n \\ \alpha = m_{ij}}} \lambda_{i,j,\alpha,l} = c_l^{(j)}$$

by eq. (9). Therefore $\bar{y}$ also defines a maximum flow in $\mathcal{N}(M)$. Note that y^* and $\bar{y}$ are in fact two distinct flow functions. This is so because $\lambda_{i,j,\alpha,\beta} > 0$ for some $i, j, \alpha = m_{ij}$, and $\beta \neq \alpha$ when we formed eq. (4); it then follows from definitions of y^* and $\bar{y}$ in (10) and (11) that to the particular arc $(R_l^{(i)}, V_{ij})$, with $l = \beta$, we have

$$y^*(i, j, l) = 0, \qquad \bar{y}(i, j, l) > 0. \tag{12}$$

We are now ready to derive a contradiction that M could not be s-unique. Formulate the maximum flow problem for $\mathcal{N}(M)$ as a linear program in the standard way (e.g., [8, Chapter 8]):

$$\begin{array}{ll} \text{maximize} & z = c\cdot y, \\ \text{subject to} & A\cdot y = b, \qquad y \ge 0, \end{array}$$

with suitable vectors b, c, and matrix A. It is known [8, Theorem 8.8] that in the present case, when A is unimodular and b is an integer vector (representing the capacity constraints in $\mathcal{N}(M)$), the bounded polyhedron Y defined by $Ay = b$, $y \ge 0$ has the property that all of its extreme points have integer components. Let us write $\bar{y}$ as a convex linear combination of the extreme points of Y (this is always possible; see [12, Theorem 2.12.2]),

$$\bar{y} = \sum a_k y_k \qquad \text{where} \quad a_k \ge 0, \quad \sum a_k = 1.$$

Since $\bar{y} \neq y^*$, we must have $a_k > 0$ for some extreme point y_k with $y_k \neq y^*$. Denote this y_k by y'. Because of (12), we can further assume that y' is chosen such that

$$y'(i, j, l) > 0 \tag{13}$$

for the particular triple (i, j, l) in (12). By the theorem quoted above, y' has integer components. Furthermore, since z is a concave function of y, that is,

$$\begin{aligned} c\cdot\bar{y} &= c\cdot\left(\sum_{a_k>0} a_k y_k\right) \\ &= \sum_{a_k>0} a_k (c\cdot y_k) \\ &\le \max_{a_k>0} c\cdot y_k, \end{aligned}$$

the fact that z is maximized at $\bar{y}$ implies that it must be maximized at all y_k with $a_k > 0$. In summary, we know that (i) y' is a maximum flow for $\mathcal{N}(M)$, distinct from y^* and satisfying (13); (ii) y' has integer assignments to all arcs in $\mathcal{N}(M)$—in fact the assignments are 0–1 valued since the total flow through any V_{ij} is 1.

We now define a simple n-ary matrix $M' = (m'_{ij})$ corresponding to y' by letting $m'_{ij} = l$, where l is the unique integer with $y'(i, j, l) = 1$. The fact that all nodes $R_l^{(i)}$ and $C_l^{(j)}$ are saturated under y' implies that M' has row and column signatures as given by $r_l^{(i)}$ and $c_l^{(j)}$. Note that $M' \neq M$ since $m'_{ij} = l$ by (13), while $m_{ij} \neq l$ by (10) and (12), for some triple (i, j, l). But this contradicts the assumption that M is s-unique. We therefore conclude that the system $(\mathcal{S})$ can be solved to find D, D' such that $M = C_{D,D'}$. The proof of Theorem 4 is thus complete. □

4. *Bounds on the Number of Simple Connection Matrices*

On the basis of the characterization derived in the previous section, we shall find bounds on the number of $n \times n$ n-ary simple matrices that are realizable.

THEOREM 5. $(C/n)^{n/2}4^{n^2} \leq |SR(n)| \leq 4^{2n^2}$, *for some constant* $C > 0$.

PROOF. We first show the upper bound. By Theorem 4, an $n \times n$ n-ary simple matrix M is in $SR(n)$ only if M has a unique signature among simple matrices. Therefore $|SR(n)|$ cannot be greater than the total number of such distinct signatures. In a signature $(\vec{r}^{(1)}, \vec{r}^{(2)}, \ldots, \vec{r}^{(n)}, \vec{c}^{(1)}, \vec{c}^{(2)}, \ldots, \vec{c}^{(n)})$, each component $\vec{r}^{(i)} = (r_1^{(i)}, r_2^{(i)}, \ldots, r_n^{(i)})$ can be viewed as a partition of integer n into n labeled parts. Thus each $\vec{r}^{(i)}$ can take at most $\binom{n+n-1}{n-1} \leq 4^n$ different values. It follows that the total number of distinct signatures (for simple matrices) is at most $(4^n)^{2n} = 4^{2n^2}$. This proves $|SR(n)| \leq 4^{2n^2}$.

The rest of this section is devoted to the proof of $|SR(n)| \geq (C/n)^{n/2}4^{n^2}$. We define a class of matrices, called row-ordered matrices, and show that they have the property of being s-unique. It follows from Theorem 4 that they are all in $SR(n)$. A demonstration that there are at least $(C/n)^{n/2}4^{n^2}$ such row-ordered matrices then complete the proof.

Definition 7. A simple n-ary matrix is *row-ordered* if the entries are nondecreasing along each row. For example, the following matrix is row-ordered:

$$\begin{pmatrix} 1 & 1 & 2 & 3 \\ 1 & 3 & 4 & 4 \\ 2 & 2 & 2 & 3 \\ 1 & 2 & 3 & 3 \end{pmatrix}$$

THEOREM 6. *A row-ordered matrix is s-unique.*

PROOF. Let M be a row-ordered matrix, and let $(\vec{r}^{(i)})$, $(\vec{c}^{(j)})$ be its row and column signatures. We shall show that M is the only simple n-ary matrix whose signatures are $(\vec{r}^{(i)})$ and $(\vec{c}^{(j)})$.

Let $\bar{M}$ be any simple n-ary matrix with signatures $(\vec{r}^{(i)})$ and $(\vec{c}^{(j)})$. Clearly $\bar{M}$ must have the same dimensions as M. We shall now prove that the signatures determine which entries of $\bar{M}$ contain a 1, which entries contain a 2, . . . , etc.

Let a be the smallest integer that appears in $\bar{M}$. Note that a is uniquely determined by the signatures. We first show that the positions (i, j) in $\bar{M}$ where a occurs are determined by the signatures.

LEMMA 1. $\bar{M}[i, j] = \{a\}$ *if and only if* $r_a^{(i)} \geq j$.

PROOF. As $(\vec{r}^{(i)})$, $(\vec{c}^{(j)})$ are signatures arising from the row-ordered matrix M, we have

$$c_a^{(1)} = |\{i \mid r_a^{(i)} \geq 1\}|, \tag{14}$$

and, in general,

$$c_a^{(j)} = |\{i \mid r_a^{(i)} \geq j\}|. \tag{15}$$

We can now prove the lemma by induction on j.

$j = 1$. The only positions $(i, 1)$ in the first column of $\bar{M}$ where a may appear are those with $r_a^{(i)} \geq 1$. But, by (14), we must actually place a's in all such positions in order to satisfy the requirement of having $c_a^{(1)}$ a's in the first column.

Induction step. Suppose the lemma is true for all $j \leq j_0$. We will prove it for $j = j_0 + 1$. Consider the $(j_0 + 1)$st column of $\bar{M}$. By the induction hypothesis, each row i has had exactly $\min\{r_a^{(i)}, j_0\}$ a's appearing in column 1 through column j_0. Therefore, only those rows i with $r_a^{(i)} \geq j_0 + 1$ could have a's appearing in the $(j_0 + 1)$st column. By (15), all such rows must actually have a's in the $(j_0 + 1)$st column in order to satisfy (15). This completes the induction step of the lemma. □

PROOF OF THEOREM 6 (CONTD.). Now we complete the proof of Theorem 6 by induction on a, the smallest integer that occurs in $\bar{M}$, for $a = n, n-1, \ldots, 1$. When $a = n$, $\bar{M}$ has integer n in every entry, and this is obviously uniquely determined from the signature. Suppose it is true that $\bar{M} = M$ whenever $a \geq a_0 + 1$; we will prove it for $a = a_0$. By the preceding lemma, the positions in $\bar{M}$ where a_0 occurs are only dependent on the signature. Therefore M and $\bar{M}$ have a_0 at exactly the same positions. Now replace the a_0's in both M and $\bar{M}$ by $a_0 + 1$ and call the new matrices M' and $\bar{M}'$, respectively. Clearly this transformation still leaves M' and $\bar{M}'$ with the same signature, and M' is again a row-ordered matrix. By the induction hypothesis, since the smallest integer in $\bar{M}'$ is $a_0 + 1$, we must have $\bar{M}' = M'$. But this implies that, before replacing a_0 by $a_0 + 1$, it must be true that $\bar{M} = M$. This proves Theorem 6. □

It is easy to see that any matrix which can be transformed into a row-ordered matrix through row and column permutations is also s-unique.

PROOF OF THEOREM 5 (CONTD.). We now count the number of row-ordered matrices. As demonstrated earlier, the number of choices of $\vec{r}^{(i)}$ is $\binom{2n-1}{n-1} = \frac{1}{2}\binom{2n}{n} = (1/2\sqrt{\pi n})4^n \cdot (1 + O(1/n)) \geq (C/n)^{1/2}4^n$ for some $C > 0$. Therefore the number of possible signatures $(\vec{r}^{(1)}, \vec{r}^{(2)}, \ldots, \vec{r}^{(n)})$ is at least $(C/n)^{n/2}4^{n^2}$. Since every such signature can be achieved by some row-ordered matrix, we have established that there are at least $(C/n)^{n/2}4^{n^2}$ row-ordered matrices and hence $|SR(n)| \geq (C/n)^{n/2}4^{n^2}$. This completes the proof of Theorem 5. □

5. *Enumeration and Characterization of General Connection Matrices*

We extend the preceding results about $SR(n)$ to $R(n)$, the set of all connection matrices. In Section 5.1 we introduce the notion of "spanning matrices" and discuss their properties. The results are used in Section 5.2 to derive an upper bound of C^{n^2} on $|R(n)|$, which by Theorem 3 is also an upper bound on the number of edges of the Triangular polyhedron $T^{(n)}$. Finally a characterization of $R(n)$ similar to Theorem 4 is given in Section 5.3.

5.1 SPANNING MATRICES. Let M be any $n \times n$ n-ary matrix. Define $\mathcal{S}_M$ to be the following induced system of linear equations:

$$\mathcal{S}_M : h_{i,j,\alpha,\beta} = (d_{i\alpha} + d'_{\alpha j}) - (d_{i\beta} + d'_{\beta j}) = 0 \quad \text{for} \quad \alpha, \beta \in M(i,j), \quad \alpha \neq \beta, \quad 1 \leq i,j \leq n. \tag{16}$$

As there are only $2n^2$ variables d_{ij} and d'_{ij}, at most $2n^2$ of these equations can be linearly independent. Choose any fixed maximal independent subset $\mathcal{L}$ of $\mathcal{S}_M$ (clearly $|\mathcal{L}| \leq 2n^2$). We define an n-ary matrix H by

$$H[i,j] = \begin{cases} M[i,j] & \text{if} \quad |M[i,j]| = 1, \qquad (17) \\ \{\alpha \mid h_{i,j,\alpha,\beta} = 0 \text{ is in } \mathcal{L} \text{ for some } \beta\} \cup \{\beta \mid h_{i,j,\alpha,\beta} = 0 \text{ is in } \mathcal{L} \text{ for some } \alpha\} & \text{if} \quad |M[i,j]| > 1. \qquad (18) \end{cases}$$

An n-ary matrix H obtained this way is called a *spanning matrix* for M. The total weight of H clearly satisfies $w(H) \leq n^2 + 2|\mathcal{L}| \leq 5n^2$. A basic property of H is the following: For a pair of distance matrices D and D', if it is known that $\min\{d_{ik} + d'_{kj} \mid 1 \leq k \leq n\}$ is achieved by every $\alpha \in H[i,j]$ (for all $1 \leq i,j \leq n$), then it is also achieved by every $\alpha \in M[i,j]$. Formally, we have the lemma following Definition 8.

Definition 8. For two n-ary matrices M and M', we say $M' \subseteq M$ if $M'[i,j] \subseteq M[i,j]$ for all i,j.

LEMMA 2. *Let H be a spanning matrix of an $n \times n$ n-ary matrix M. If $M' \in R(n)$ is a connection matrix and $H \subseteq M'$, then $M \subseteq M'$.*

PROOF. Let $M' = C_{\bar{D},\bar{D}'}$. By the assumption that $H \subseteq M'$, we have for any i,j,

$$\bar{d}_{i\alpha} + \bar{d}'_{\alpha j} \leq \bar{d}_{ik} + \bar{d}'_{kj}, \qquad 1 \leq k \leq n, \quad \alpha \in H[i,j]. \tag{19}$$

This implies $h_{i,j,\alpha,\beta}(\bar{D}, \bar{D}') = 0$, $1 \le i, j \le n$, $\alpha, \beta \in H[i, j]$, $\alpha \ne \beta$. As H is derived from a maximal independent subset of $\mathcal{S}_M$ in (16), we have

$$h_{i,j,\alpha,\beta}(\bar{D}, \bar{D}') = 0, \qquad 1 \le i, j \le n, \quad \alpha, \beta \in M[i, j], \quad \alpha \ne \beta. \tag{20}$$

Formulas (19) and (20) imply that, if $|M[i, j]| > 1$, then

$$\bar{d}_{i\alpha} + \bar{d}'_{\alpha j} \le \bar{d}_{ik} + \bar{d}'_{kj}, \qquad 1 \le k \le n, \quad \alpha \in M[i, j],$$

and therefore $M[i, j] \subseteq M'[i, j]$.

If $|M[i, j]| = 1$, then $M[i, j] = H[i, j] \subseteq M'[i, j]$. □

THEOREM 7. *Let H and H' be spanning matrices for connection matrices M and M', respectively. If H and H' have the same weight distribution and the same signature, then $M = M'$.*

If a connection matrix M is simple, the only spanning matrix for M is itself. In this case the above theorem becomes a weaker form of the s-uniqueness condition for M in Theorem 4 (weaker because M' is assumed to be a connection matrix).

PROOF. Since H and H' have the same weight distribution, $|H[i, j]| = |H'[i, j]|$ for all i, j. Let us match the elements of $H[i, j]$ and $H'[i, j]$ in disjoint pairs as $Q_{ij} = \{(\alpha, \beta)\}$, where $\alpha \in H[i, j]$, $\beta \in H'[i, j]$, and $|Q_{ij}| = |H[i, j]|$.

Let $M = C_{D,D'}$ for $D = (d_{ij})$ and $D' = (d'_{ij})$; we can write down the following set of inequalities:

$$\mathcal{H}: d_{i\alpha} + d'_{\alpha j} \le d_{i\beta} + d'_{\beta j} \qquad \text{for} \quad (\alpha, \beta) \in Q_{ij}, \quad 1 \le i, j \le n, \quad \text{with equality only if} \quad \beta \in M[i, j].$$

When we add up the $w(H)$ inequalities in $\mathcal{H}$, we obtain

$$\sum_i \sum_l r_l^{(i)} d_{il} + \sum_j \sum_l c_l^{(j)} d'_{lj} \le \sum_i \sum_l r_l'^{(i)} d_{il} + \sum_j \sum_l c_l'^{(j)} d'_{lj}, \tag{21}$$

with equality holding only if $H' \subseteq M$, where $(r_l^{(i)}, c_l^{(j)})$ and $(r_l'^{(i)}, c_l'^{(j)})$ are the signatures of H and H', respectively. Since by assumption H and H' have the same signature, the two sides in eq. (21) are equal. Therefore $H' \subseteq M$. By Lemma 2, this implies $M' \subseteq M$.

A similar argument shows $M \subseteq M'$. Hence $M = M'$. □

5.2 A C^{n^2} BOUND FOR $|R(n)|$. We will show that there are at most C^{n^2} connection matrices (out of the 2^{n^2} $n \times n$ n-ary matrices).

THEOREM 8. $|R(n)| \le C^{n^2}$ *for some constant* C.

COROLLARY

$$\Big| \bigcup_{0 \le s \le \binom{n}{2}} \mathcal{F}_s(T^{(n)}) \Big| \le C^{n^2}.$$

PROOF. For each $M \in R(n)$, choose a spanning matrix H_M. By Theorem 7, all the weight distribution-signature pairs of H_M, i.e., $(W(H_M), s(H_M))$, are distinct. Furthermore, the total weight of H_M satisfies $n^2 \le w(H_M) \le 5n^2$. Therefore $|R(n)|$ is bounded by the product $u \cdot v$, where u is the number of ways for distributing a total weight A, $n^2 \le A \le 5n^2$, to the n^2 entries in the $n \times n$ matrix and v is an upper bound on the maximum number of distinct signatures under any fixed weight distribution (with total weight $n^2 \le A \le 5n^2$). We will show that $u \le (64)^{n^2}$ and $v \le c^{n^2}$ for some constant c, which then implies the theorem.

The number u is bounded by the number of ways of partitioning the integer $5n^2$ into $n^2 + 1$ labeled parts, where the last part specifies $5n^2 - A$. Therefore

$$u \le \binom{5n^2 + n^2}{n^2} \le 2^{6n^2} = (64)^{n^2}.$$

To estimate v, let b_W be the total number of distinct n-tuples of row signatures

$(\vec{r}^{(1)}, \vec{r}^{(2)}, \ldots, \vec{r}^{(n)})$ subject to a fixed weight distribution W. It then follows that $v \leq \max_W (b_W)^2$, where we have restricted W to those with total weight $n^2 \leq A \leq 5n^2$. For any such W, let us regard every $(\vec{r}^{(1)}, \vec{r}^{(2)}, \ldots, \vec{r}^{(n)})$ as an $n \times n$ matrix whose (i, j)th entries are $r_j^{(i)}$. It follows that b_W is bounded by the number of ways of partitioning the integer A into n^2 labeled parts. Thus

$$b_W \leq \binom{A + n^2 - 1}{n^2 - 1} \leq 2^{A+n^2-1} \leq 2^{6n^2}.$$

It follows that

$$v \leq (2^{6n^2})^2 = 2^{12n^2}.$$

This completes the proof of the theorem. The corollary follows immediately from Theorem 3. □

We have not tried to obtain the best constants C in the above proof. David Avis reported (private communication) that, by sharpening the above arguments, the constant C in the Corollary to Theorem 8 can be taken to be $2^{2.72}$.

5.3 CHARACTERIZATION OF CONNECTION MATRICES. We will state a necessary and sufficient condition for an n-ary matrix to be a member of $R(n)$. The proof is a slight extension of that given for Theorem 4 and hence will not be repeated.

Definition 9. A *multiset* U is analogous to a set except that an element may appear more than once in U. We use $|U|$ to denote the total number of elements appearing in U. Thus $|U| = 6$ for $U = \{1, 2, 2, 2, 3, 3\}$.

Definition 10. An *n-ary multimatrix* M is a matrix where each entry $M[i, j]$ is a multiset whose elements are drawn from $[1, 2, \ldots, n]$, with $|M[i, j]| \leq n$.

The concepts of *weight distribution* and *signature* defined in Section 3 can also be generalized to an n-ary multimatrix in the obvious way.

Definition 11. For two n-ary multimatrices M and M', we say $M' \subseteq M$ if every element that appears in the multiset $M'[i, j]$ also occurs at least once in $M[i, j]$, for $1 \leq i, j \leq n$.

We generalize the definition of s-uniqueness to n-ary matrices as follows.

Definition 12. An n-ary matrix M is said to be *s-unique* if for any n-ary multimatrix M' with the same weight distribution, $s(M') = s(M)$ implies that $M' \subseteq M$.

THEOREM 9. *Let M be an $n \times n$ n-ary matrix. Then $M \in R(n)$ if and only if M is s-unique.*

6. *Enumeration of the Patterns of Shortest Paths*

In this section we examine an information bound based directly on the solution space of computing shortest distances. Let G be a directed complete graph on n vertices $\{v_1, v_2, \ldots, v_n\}$, with a nonnegative distance d_{ij} assigned to each edge (v_i, v_j). A *path* from v_i to v_j is a finite sequence of vertices $(i = k_0, k_1, k_2, \ldots, k_{m-1}, k_m = j)$, not necessarily all distinct. The *length* of such a path is $\sum_{1 \leq l \leq m} d_{k_{l-1}, k_l}$. We shall also consider the sequence of a single point (i) to be a path from i to i, called a *null path*, with length 0. The entry d_{ij}^* in the transitive closure D^* is then the minimum length of any path from i to j. For any i, j, let p_{ij} be the set of all shortest paths in G from v_i to v_j. (The set p_{ij} may be infinite.) We denote by *pattern*(D) the $n \times n$ matrix (p_{ij}) associated with the distance matrix $D = (d_{ij})$. Let $P(n)$ be the collection of all distinct patterns induced by $n \times n$ distance matrices. By an argument similar to that used in Theorem 2, one can show that any linear decision tree for computing the shortest distance matrix D^*, given D, requires at least $\log_3 |P(n)| - n^2$ comparisons in the worst case. This, intuitively, is probably the best information lower bound one can hope for; the previous approach using connection matrices can be regarded

as a special case with the vertices divided into three disjoint sets X_0, X_1, X_2, such that all edges except those from X_0 to X_1 and from X_1 to X_2 are effectively ∞.

The rest of this section is devoted to proving the following theorem, which states that no nontrivial lower bound can be obtained even in the present version of the information-theoretic approach.

THEOREM 10. *$|P(n)| \leq C^{n^2}$ for some constant $C > 0$.*

We first generalize the notion of a connection matrix to that for $m + 1$ consecutive sets of "cities" $X_0, X_1, \ldots, X_m$. Assuming that $D^{(l)} = (d_{ij}^{(l)})$ defines the distances between any pairs of cities in $X_{l-1} \times X_l$, then $C_{D^{(1)},\ldots,D^{(m)}}[i, j]$ is to be the set of best connecting paths from city $i \in X_0$ to city $j \in X_m$. Formally, if $(D^{(1)}, D^{(2)}, \ldots, D^{(m)})$ is a sequence of m $n \times n$ matrices, then their *m-connection matrix* $C_{D^{(1)},\ldots,D^{(m)}}$ is defined by

$$\begin{aligned} C_{D^{(1)},\ldots,D^{(m)}}[i, j] = \{(\alpha_1, \alpha_2, \ldots, \alpha_{m-1}) \mid & 1 \leq \alpha_l \leq n \text{ for all } l, \\ & \text{and } d_{i\alpha_1}^{(1)} + d_{\alpha_1\alpha_2}^{(2)} + \cdots + d_{\alpha_{m-1}j}^{(m)} \\ & = \min_{k_1,\ldots,k_{m-1}} (d_{ik_1}^{(1)} + d_{k_1k_2}^{(2)} + \cdots + d_{k_{m-1}j}^{(m)})\}. \end{aligned}$$

This definition reduces to the connection matrix defined previously when $m = 2$.

Let $R_m(n)$ denote the set of all possible $n \times n$ m-connection matrices.

LEMMA 3. *$|R_m(n)| \leq |R(n)|^{m-1}$ for $m \geq 2$.*

PROOF. We will show that, for $m > 2$, $C_{D^{(1)},\ldots,D^{(m)}}$ is determined by $C_{D^{(1)},\ldots,D^{(m-1)}}$ and $C_{A,D^{(m)}}$, where $A = D^{(1)} \otimes D^{(2)} \otimes \cdots \otimes D^{(m-1)}$. This will imply that $|R_m(n)| \leq |R_{m-1}(n)| \cdot |R_2(n)|$. The lemma then follows by induction, observing that $|R_2(n)| = |R(n)|$.

Let $A = D^{(1)} \otimes D^{(2)} \otimes \cdots \otimes D^{(m-1)}$. Since $\min_{k_1,\ldots,k_{m-1}}(d_{ik_1}^{(1)} + \cdots + d_{k_{m-1}j}^{(m)}) = \min_{k_{m-1}}(\min_{k_1,\ldots,k_{m-2}}(d_{ik_1}^{(1)} + \cdots + d_{k_{m-2}k_{m-1}}^{(m-1)}) + d_{k_{m-1}j}^{(m)})$, an alternative description of $C_{D^{(1)},\ldots,D^{(m)}}[i, j]$ is the set of $(\alpha_1, \ldots, \alpha_{m-2}, \alpha_{m-1})$ such that $\alpha_{m-1} \in C_{A,D^{(m)}}[i, j]$, and $(\alpha_1, \ldots, \alpha_{m-2}) \in C_{D^{(1)},\ldots,D^{(m-1)}}[i, \alpha_{m-1}]$. This proves that $C_{D^{(1)},\ldots,D^{(m)}}$ is determined by $C_{D^{(1)},\ldots,D^{(m-1)}}$ and $C_{A,D^{(m)}}$. □

PROOF OF THEOREM 10. We shall derive a recurrence relation on $|P(n)|$. We use the idea employed in [1] for reducing the shortest paths problem to {min, +} multiplication. Let X be any $2n \times 2n$ distance matrix on vertices $\{1, 2, \ldots, 2n\}$. We write it in the form of four $n \times n$ blocks

$$X = \begin{pmatrix} A & B \\ Y & D \end{pmatrix} \tag{22}$$

The shortest distances matrix X^* then satisfies the following recurrence formula [1, p. 204]:

$$X^* = \begin{pmatrix} E^* & E^* \otimes B \otimes D^* \\ D^* \otimes Y \otimes E^* & D^* \oplus (D^* \otimes Y \otimes E^* \otimes B \otimes D^*) \end{pmatrix}, \tag{23}$$

where $E = (A \oplus (B \otimes D^* \otimes Y))$. Actually, implicit in the derivation of (23) is an enumeration of all possible shortest paths between any two of the $2n$ vertices in terms of quantities involving only $n \times n$ matrices. We now make this statement precise in a lemma.

Definition 13. Let $\mathscr{E}$ and $\mathscr{E}'$ be the $n \times n$ matrices of 0's and ± 1's defined below:

$$\mathscr{E}_{ij} = \begin{Bmatrix} -1 \\ 0 \\ 1 \end{Bmatrix} \quad \text{if} \quad (A)_{ij} \begin{Bmatrix} < \\ = \\ > \end{Bmatrix} (B \otimes D^* \otimes Y)_{ij},$$

$$\mathscr{E}'_{ij} = \begin{Bmatrix} -1 \\ 0 \\ 1 \end{Bmatrix} \quad \text{if} \quad (D^*)_{ij} \begin{Bmatrix} < \\ = \\ > \end{Bmatrix} (D^* \otimes Y \otimes E \otimes B \otimes D^*)_{ij}.$$

Define the *counting vector* $\mu(X)$, for X as in (26), to be $\mu(X) = (pattern(D), pattern(E), C_{E^*,B,D^*}, C_{D^*,Y,E^*}, C_{D^*,Y,E^*,B,D^*}, C_{B,D^*,Y}, \mathscr{E}, \mathscr{E}')$.

LEMMA 4. *The matrix pattern(X) is determined by the counting vector* $\mu(X)$.

PROOF. We shall show that the (i, j)th entry of *pattern*(X) is determined by $\mu(X)$ for all i, j.

First we assume $1 \le i, j \le n$. Following the original argument [1, p. 204] leading to (23), any path from vertex i to vertex j can be written uniquely as

$$(i = k_0, \sigma_1, k_1, \sigma_2, k_2, \ldots, k_{l-1}, \sigma_l, k_l, \ldots, \sigma_m, k_m = j),$$

where each $k_l \in \{1, 2, \ldots, n\}$ and each σ_l is a sequence of vertices (possibly empty) in $\{n+1, n+2, \ldots, 2n\}$ (m may be 0 when $i = j$). A shortest path from i to j is characterized by the following conditions:

(a) Each $k_{l-1}\sigma_l k_l$ is among the shortest such paths from k_{l-1} to k_l; denote this length by $\text{leng}(k_{l-1}, k_l)$.
(b) The k's satisfy the condition that $\sum_l \text{leng}(k_{l-1}, k_l)$ is minimum for all possible choices of the k's.

We can restate the conditions as follows. Let

$$Q = pattern(E),$$
$$\Delta_{st} = \bigcup_{(h,h') \in C_{B,D^*,Y}[s,t]} [pattern(D)]_{h,h'},$$

and let Γ be the $n \times n$ matrix defined by

$$\Gamma_{st} = \begin{cases} \{\lambda\} & \text{if } \mathscr{E}_{st} = -1, \\ \Delta_{st} & \text{if } \mathscr{E}_{st} = 1, \\ \{\lambda\} \cup \Delta_{st} & \text{if } \mathscr{E}_{st} = 0, \end{cases}$$

where we use λ for the null sequence. Then condition (b) is equivalent to $(k_0, k_1, \ldots, k_m) \in Q_{ij}$ and condition (a) is equivalent to $\sigma_l \in \Gamma_{k_{l-1}k_l}$ for $1 \le l \le m$. But this implies that the (i, j)th entry of *pattern*(X), i.e., the set of all shortest paths from i to j, is determined by Q and Γ and hence by *pattern*(E), *pattern*(D), $C_{B,D^*,Y}$, and $\mathscr{E}$. This proves the lemma for the case $1 \le i, j \le n$.

Similarly, one can show that the set of shortest paths from i to j is determined by *pattern*(E), $C_{B,D^*,Y}$, $\mathscr{E}$, and, in addition,

C_{E^*,B,D^*} and *pattern*(D)	if $1 \le i \le n$,	$n+1 \le j \le 2n$,
C_{D^*,Y,E^*} and *pattern* (D)	if $n+1 \le i \le 2n$,	$1 \le j \le n$,
C_{D^*,Y,E^*,B,D^*}, *pattern* (D), and $\mathscr{E}'$	if $n+1 \le i, j \le 2n$.	

We omit the details. □

PROOF OF THEOREM 10 (CONTD.). To complete the proof of Theorem 10, we note that, by Lemma 4, the number of distinct patterns is bounded by the number of distinct counting vectors. This leads to

$$|P(2n)| \le |P(n)|^2 \cdot |R(n)|^2 \cdot |R(n)|^2 \cdot |R(n)|^4 \cdot |R(n)|^2 \cdot 3^{2n^2}$$

by Definition 13 and Lemma 3.

Writing $f(n)$ for $|P(n)|$ and using Theorem 8, we obtain

$$f(2n) \le (f(n))^2 C^{n^2} \qquad \text{for some constant } C.$$

Taking logarithms,

$$\ln f(2n) \le 2 \ln f(n) + n^2 \ln C.$$

For $n = 2^k$ this leads to (noting that $f(1) = 2$)[4]

$$\ln f(2n) \le (\ln C)\left(n^2 + 2\left(\frac{n}{2}\right)^2 + 2^2\left(\frac{n}{2^2}\right)^2 + \cdots + 2^k\left(\frac{n}{2^k}\right)^2 + 2^{k+1}\ln f(1)\right)$$
$$\le 4n^2 \ln C.$$

This proves $f(n) \le C^{n^2}$ if n is a power of 2.

For general n one can easily show $f(n) \le f(2^{\lceil \log n \rceil})$ by adding extra points with effectively ∞ distances between these points and the other vertices. This leads to $f(n) \le C^{4n^2}$ immediately. The proof of Theorem 10 is thus complete. □

7. *Other Applications of s-Uniqueness*

The s-uniqueness characterization technique used in this paper might be useful in the study of other complexity problems. As an example we consider below a particular sorting-type problem.

Let $X = \{x_1, x_2, \ldots, x_n\}$ be a linearly ordered set of n distinct elements, and let S_1, $S_2, \ldots, S_m$ be specified subsets of X. Consider the problem of computing the minimum element of each S_i by pairwise comparisons between the elements. For $m = 2$ this can be done in $n - 1$ comparisons by first computing y_1, y_2, y_3 (the minimum elements in $S_1 \cap \bar{S}_2$, $S_1 \cap S_2$, $\bar{S}_1 \cap S_2$, respectively), followed by the computation of $\min\{y_1, y_2\}$ and $\min\{y_2, y_3\}$. This scheme can be extended to solve the problem, for any fixed m, in $n + O(1)$ comparisons. However, for arbitrary m and n, the problem does not seem to be solvable in linear time, i.e., using $O(m + n)$ comparisons. It is thus of interest to study the quantity $I(S_1, S_2, \ldots, S_m) \equiv \log_2(\#$ of possible answers), which is the information-theoretic lower bound in the decision tree model. We shall now demonstrate that $I(S_1, S_2, \ldots, S_m) \le m + n$, i.e., the information bound is weak for this problem. In the following, the definitions for the terms "realizable," "signatures," and "s-uniqueness" are for this problem only and are different from their usage in earlier sections.

Definition 14. Let $L = (x_{j_1}, x_{j_2}, \ldots, x_{j_m})$ be a *list of representatives*, i.e., $x_{j_i} \in S_i$ for $1 \le i \le m$. We call this list *realizable* provided that there exist a linear ordering of X relative to which $x_{j_i} = \min\{x_l \mid x_l \in S_i\}$ for $1 \le i \le m$.

Let $\mathcal{L}$ denote the set of all realizable lists of representatives. ($\mathcal{L}$ depends on S_1, $S_2, \ldots, S_m$.) Obviously $I(S_1, S_2, \ldots, S_m) = \log_2|\mathcal{L}|$. We shall prove that $|\mathcal{L}| \le 2^{m+n}$. This immediately implies $I(S_1, S_2, \ldots, S_m) \le m + n$, as was to be shown.

Definition 15. Let $L = (x_{j_1}, x_{j_2}, \ldots, x_{j_m})$ be a list of representatives and let r_k, $1 \le k \le n$, denote the number of i such that $j_i = k$. We refer to the vector $(r_1, r_2, \ldots, r_n)$ as the *signature* of L. The list L is *s-unique* if no other list L' has the same signature as L.

THEOREM 11. *A list of representatives $L = (x_{j_1}, x_{j_2}, \ldots, x_{j_m})$ is realizable iff it is s-unique.*

COROLLARY. $|\mathcal{L}| \le 2^{m+n}$.

PROOF

Necessity. Let $L' = (x_{j'_1}, x_{j'_2}, \ldots, x_{j'_m})$ be another list of representatives having the same signature as L, and assume that X is ordered so as to realize L. Then the mapping $j_i \to j'_i$ induces an automorphism with the multiset $\{x_{j_1}, x_{j_2}, \ldots, x_{j_m}\}$ which is nondecreasing, and strictly increasing at least once—an impossibility.

Sufficiency. Suppose that L is not realizable. Construct the directed graph G whose vertex set is X and which includes precisely the edges (x_{j_i}, x_k) such that $x_k \in S_i$ and $k \ne j_i$. The fact that L is not realizable implies that G contains a directed cycle, say,

$$(x_{t_1}, x_{t_2}), \ldots, (x_{t_{d-1}}, x_{t_d}), (x_{t_d}, x_{t_1}).$$

[4] When $n = 1$, $pattern(D) = (p_{11})$, where $p_{11} = \{(1)\}$ if $d_{11} > 0$ and $p_{11} = \{(1), (1, 1), (1, 1, 1), (1, 1, 1, 1), \ldots\}$ if $d_{11} = 0$.

By definition of G, there exist distinct $h_1, h_2, \ldots, h_d$ such that $x_{l_u}, x_{l_{u+1}} \in S_{h_u} (1 \leq u < d)$, $x_{l_d}, x_{l_1} \in S_{h_d}$, and x_{l_u} is the representative chosen from S_{h_u} for $1 \leq u \leq d$. However, from the sets S_{h_u} we can pick instead the representatives $x_{l_{u+1}}$ if $u < d$, x_{l_1} if $u = d$. Picking the same representatives from the remaining S_k as before, we obtain a different list of representatives with the same signature. Thus L is not s-unique. □

COROLLARY. *As the number of distinct signatures is at most* $\binom{m+n-1}{n-1} < 2^{m+n}$, *the corollary follows.* □

ACKNOWLEDGMENTS. We wish to thank David Avis and Michael Fredman for many helpful comments. The results in Section 7 are due to Michael Fredman.

REFERENCES

1. AHO, A.V., HOPCROFT, J.E., AND ULLMAN, J.D. *The Design and Analysis of Computer Algorithms.* Addison-Wesley, Reading, Mass., 1974.
2. DANTZIG, G.B. On the shortest route through a network. *Management Sci. 6* (1960), 187–190.
3. DIJKSTRA, E.W. A note on two problems in connection with graphs. *Numer. Math. 1* (1959), 269–271.
4. FISCHER, M.J., AND MEYER, A. Boolean matrix multiplication and transitive closure. Proc. 12th Symp. on Switching and Automata Theory, East Lansing, Mich., 1971, pp. 129–131.
5. FLOYD, R.W. Algorithm 97: shortest path. *Comm. ACM 5*, 6 (June 1962), p. 345.
6. FREDMAN, M.L. New bounds on the complexity of the shortest path problem. *SIAM J. Comptg. 5* (1976), 87–89.
7. GRUNBAUM, B. *Convex Polytopes.* Interscience, New York, 1967.
8. HU, T.C. *Integer Programming and Network Flows.* Addison-Wesley, Reading, Mass., 1969.
9. KERR, L.R. The effect of algebraic structure on the computational complexity of matrix multiplication. Ph.D. Th., Cornell U., Ithaca, N.Y., 1970.
10. KNUTH, D.E. *The Art of Computer Programming, Vol. 3, Sorting and Searching.* Addison-Wesley, Reading, Mass., 1973.
11. MUNRO, I. Efficient determination of the transitive closure of a directed graph. *Inform. Proc. Letters 1* (1971), 56–58.
12. STOER, J., AND WITZGALL, C. *Convexity and Optimization on Finite Dimensions I.* Springer-Verlag, Berlin-Heidelberg, 1970.
13. YAO, A.C., AVIS, D.M., AND RIVEST, R.L. An $\Omega(n^2 \log n)$ lower bound to the shortest paths problem. Proc. 9th Symp. on Theory of Comptg., Boulder, Colo. 1977, pp. 11–17.
14. YAO, A.C., AND RIVEST, R.L. On the polyhedral decision problem. *SIAM J. Comptg.* (May 1980).

RECEIVED APRIL 1979; REVISED MAY 1979; ACCEPTED JUNE 1979