

Hypersphere Mapper: A Nonlinear Programming Approach to the Hypercube Embedding Problem*

JOHN K. ANTONIO

School of Electrical Engineering, Purdue University, 1285 Electrical Engineering Building, West Lafayette, Indiana 47907-1285

AND

RICHARD C. METZGER

Rome Laboratory/C3CB, Software Engineering Branch, 525 Brooks Road, Griffiss AFB, New York 13441-5700

A nonlinear programming approach is introduced for solving the hypercube embedding problem. The basic idea of the proposed approach is to approximate the discrete space of an n -dimensional hypercube, i.e., $\{z: z \in \{0, 1\}^n\}$, with the continuous space of an n -dimensional hypersphere, i.e., $\{x: x \in \mathbb{R}^n \text{ \& \; } \|x\|^2 = 1\}$. The mapping problem is initially solved in the continuous domain by employing the gradient projection technique to a continuously differentiable objective function. The optimal process "locations" from the solution of the continuous hypersphere mapping problem are then discretized onto the n -dimensional hypercube. The proposed approach can solve, directly, the problem of mapping P processes onto N nodes for the general case where $P > N$. In contrast, competing embedding heuristics from the literature can produce only one-to-one mappings and cannot, therefore, be directly applied when $P > N$. © 1993 Academic Press, Inc.

1. INTRODUCTION

A. Motivation

The hypercube structure has been a popular choice for interconnecting large numbers of processing elements in parallel processing systems. Examples of commercially available parallel machines that utilize a hypercube interconnection topology include nCUBE's nCUBE 2, Connection Machine's CM2, and Intel's iPSC2.

An n -dimensional hypercube has two connected nodes along each of n dimensions for a total of $N = 2^n$ nodes. Some of the attractive features of the hypercube are: a relatively low number of incident links at each node (node degree = $n = \log_2 N$), a small hop distance between nodes (network diameter = $n = \log_2 N$), and a large number of alternate paths between node pairs.

* This work was supported by Rome Laboratory under Contract F30602-92-C-0108.

The problem addressed in the present paper is how to map a given collection of processes onto the nodes (i.e., processing elements) of the hypercube topology so that the available communication resources are effectively utilized. The specific objective of interest is to map the processes so that the average Hamming distance (i.e., path length), between those pairs of processes that require communication, is minimized. For both circuit-switched and virtual cut-through routing schemes, minimizing the average Hamming distance between communicating process pairs reduces the total number of communication links needed to establish all required connections, which may reduce the latency caused by contention for common communication links.

Effectively mapping a collection of processes onto the nodes of a massively parallel computer is especially important when considering problem domains where the interprocess communication pattern is inherently irregular and/or data dependent. For example, in the general area of computational fluid dynamics, certain applications require that the solution of a partial differential equation be approximated over an *irregular* grid of discrete points [8].

In other applications, where the processes are *functionally independent*, the associated interprocess communication pattern may also be irregular and possibly data dependent. For example, in a large embedded information processing system (e.g., a data fusion system), the processing of the sensor data may be done using functional parallelism, i.e., one process (or perhaps a collection of subprocesses) may be performing FFTs, while other processes are involved in solving systems of linear equations, sorting integers, etc. In such a system, it is likely that the communication patterns between the functionally independent processes will be irregular and also depend on the values of the input data supplied by the sensors.

B. Related Work

The hypercube embedding problem has been studied extensively in the past and a variety of mapping heuristics and theoretical results have been reported in the literature. In Ref. [7], twelve hypercube mapping heuristics are evaluated in terms of mapping quality and running time (of the mapping algorithms themselves) for several classes of interprocess communication patterns. In other papers, fundamental theoretical results have been established for the case where the communication patterns are assumed to have regular structures such as grids [4], trees [11], binary trees [10], and hypercubes [3].

One potential drawback of previous embedding heuristics is that they produce one-to-one mappings. In practical applications, the requirement may arise to map P processes onto N nodes, where $P > N$; thus, many-to-one mappings are generally required. It appears that many of the known mapping heuristics (including the twelve evaluated in Ref. [7]) assume $P \leq N$. One way to overcome the one-to-one limitation of such mapping heuristics is to initially cluster the P processes into N (or fewer) clusters and then execute the mapping heuristic based on the *intercluster* communication pattern. However, because the clustering problem is generally NP-complete, an additional heuristic would be required to do the initial clustering before actually executing the mapping heuristic. One of the advantages of the algorithm proposed in this paper is that the case of $P > N$ is handled directly because the produced mapping solution is not limited to the class of one-to-one functions.

C. Organization of the Paper

In Section II, the general hypercube embedding problem is formulated in mathematical terms. A brief overview of the proposed algorithm, called hypersphere mapper, is given in Section III. Because hypersphere mapper utilizes an iterative gradient descent technique known as the gradient projection method, a general overview of the gradient projection method is provided in Section IV. Section V contains the detailed description of hypersphere mapper. In Section VI, simulation studies for hypersphere mapper are conducted for the case of randomly generated interprocess communication patterns.

II. THE HYPERCUBE EMBEDDING PROBLEM

Mapping a set of P processes, denoted $\mathcal{P} = \{0, 1, \dots, P - 1\}$, onto an $N = 2^n$ node (i.e., n -dimensional) hypercube can be viewed as a problem of determining a collection of P binary n -vectors, $z_i \in \{z: z \in \{0,1\}^n\}$, $i \in \mathcal{P}$, where the components of z_i comprise the binary address of the node that process i is to be mapped to. For in-

stance, for the case $n = 4$, $z_0 = [1 \ 1 \ 0 \ 0]^T$ indicates that process 0 is to be mapped onto node 3.

Let $d_H(z_i, z_j)$ denote the Hamming distance between nodes with addresses z_i and z_j , where the Hamming distance equals the number of differing binary components associated with the vectors z_i and z_j . For example, $d_H([1 \ 0 \ 1 \ 0]^T, [1 \ 1 \ 0 \ 0]^T) = 2$. Let \mathcal{W} denote the set of pairs of processes that require communication. Thus, $\mathcal{W} = \{(i, j): i, j \in \mathcal{P} \text{ \& process } i \text{ communicates with process } j\}$.

Therefore, for a given set \mathcal{W} (i.e., a given interprocess communication pattern), the hypersphere embedding problem involves finding P binary n -vectors, denoted by z_0, z_1, \dots, z_{P-1} , that minimize

$$f_H(z) = \frac{1}{|\mathcal{W}|} \sum_{(i,j) \in \mathcal{W}} d_H(z_i, z_j). \quad (II.1)$$

The standard hypercube embedding formulation assumes that $P \leq N$ and that the mapping is a one-to-one function, i.e., $z_i \neq z_j$, for all $i, j \in \mathcal{P}$, with $i \neq j$. It has been proven in Ref. [6] that the standard embedding problem is NP-hard.

III. OVERVIEW OF HYPERSPHERE MAPPER

The main premise of the hypersphere mapper approach is to approximate the discrete space of the hypercube, i.e., $\{z: z \in \{0, 1\}^n\}$, with a continuous n -dimensional (unit radius) hypersphere, i.e., $\{x: x \in \mathbb{R}^n \text{ \& } \|x\|^2 = 1\}$. The mapping problem in the continuous domain is to determine a collection of P real n -vectors, $x_i \in \{x: x \in \mathbb{R}^n \text{ \& } \|x\|^2 = 1\}$, $i \in \mathcal{P}$. The advantage of the continuous domain is that techniques from nonlinear programming can be employed to solve a constrained optimization problem. In particular, by defining a continuously differentiable objective function, the values of x_i are updated in an iterative fashion by using the gradient projection technique. Geometrically, the vectors x_i are points that reside on the surface of a hypersphere in the continuous domain of \mathbb{R}^n . The application of the iterative gradient projection technique acts to migrate these points around the surface of the hypersphere so as to minimize a desired objective function.

The objective function used by hypersphere mapper comprises two components: the first component is the average squared Euclidean distance between communicating pairs of processes, the second component is the average of the inverted squared Euclidean distance between every pair of processes. The first component acts to bring pairs of communicating processes close together; the second component acts to "spread-out" the process locations (and prevent any pair of processes from residing at the same point on the continuous hypersphere).

Once a solution is obtained on the continuous hypersphere, the solution vectors, denoted by x_i^* , are discretized onto the n -dimensional binary hypercube according to the sign of each component. For example, in three dimensions, the continuous vector $x_0^* = [0.4000 \ -0.5000 \ 0.7681]^T$ discretizes to $z_0^* = [1 \ 0 \ 1]^T$, which indicates that process 0 is to be mapped onto node 5.

IV. THE GRADIENT PROJECTION METHOD

The gradient projection method is an iterative gradient descent technique for solving constrained optimization problems. Each iteration of the gradient projection method comprises two parts: (1) updating the values of the variables by moving in the direction of the negative gradient of the objective function and (2) orthogonally projecting the values of the variables onto the constraint space.

Consider the following general constrained optimization problem:

$$\text{minimize } f(x) \quad (\text{IV.1})$$

$$\text{subject to } g(x) = 0. \quad (\text{IV.2})$$

The first part of the gradient projection method is to update x by moving in the direction of the negative gradient of f . Thus, letting $x^{(k)}$ denote the value of the vector x at iteration k , the first part of the gradient projection method is given by

$$x^{(k+1)} \leftarrow x^{(k)} - \alpha \nabla f(x)|_{x=x^{(k)}}, \quad (\text{IV.3})$$

where α is a positive scalar called the stepsize.

The second part of the gradient projection method is to orthogonally project $x^{(k+1)}$ onto the constraint space $g(x) = 0$. This projection operation is denoted by

$$x^{(k+1)} \leftarrow [x^{(k+1)}]_{g(x)=0}. \quad (\text{IV.4})$$

The mathematical description of precisely how to do the projection for general constraint surfaces shall not be included here. The interested reader is referred to [2]. A geometric view of an iteration from the gradient projection method is depicted in Fig. 1.

V. DETAILED DESCRIPTION OF HYPERSPHERE MAPPER

The detailed description of hypersphere mapper is divided into three main parts. First, a constrained optimization problem is formulated for the problem of mapping processes onto the surface of a continuous n -dimensional hypersphere. Second, the gradient projection technique is applied as a means of solving the proposed constrained




FIG. 1. A geometric view of an iteration of the gradient projection method being applied to a general constrained optimization problem. The dashed lines represent constant value contours for the objective function $f(x)$, i.e., $f(x) = C_1$ and $f(x) = C_2$. The value of $x^{(k)}$ is initially updated by moving in the direction of the negative gradient, i.e., $x^{(k)}$ is updated to the point $x^{(k)} - \alpha \nabla f(x)|_{x=x^{(k)}}$. Then, the updated value is projected back to the constraint surface $g(x) = 0$.

optimization problem. Finally, the method of discretizing the continuous process "locations" (from the solution of the constrained optimization problem) is described. After presenting the detailed description of hypersphere mapper, an illustrative example application is given. The final subsection discusses a technique for incrementally spreading-out the hypersphere mapper solution so as to obtain more uniform mappings.

A. Formulating the Constrained Optimization Problem

In the domain of the continuous hypersphere, the underlying objective is to minimize the average squared Euclidean distance between pairs of communicating processes, where the processes are located on the surface of an n -dimensional hypersphere in \mathcal{R}^n . Thus, given \mathcal{W} (i.e., the set of pairs of processes that require communication), the problem is to find P real n -vectors, $x_i \in \{x: x \in \mathcal{R}^n \ \& \ \|x\|^2 = 1\}$, $i \in \mathcal{P}$, that minimize

$$\frac{1}{|\mathcal{W}|} \sum_{(i,j) \in \mathcal{W}} \|x_i - x_j\|^2, \quad (\text{V.1})$$

where $x_i = (x_{i,0} \ x_{i,1} \ \dots \ x_{i,n-1})^T$ and $\|x_i - x_j\|^2 = \sum_{l=0}^{n-1} (x_{i,l} - x_{j,l})^2$.

It is also desirable to enforce that $x_i \neq x_j$, for all $i, j \in \mathcal{P}$, with $i \neq j$. To enforce this constraint, the following penalty term is proposed:

$$\frac{2}{P(P-1)} \sum_{i=0}^{P-2} \sum_{j=i+1}^{P-1} \frac{1}{\|x_i - x_j\|^2}, \quad (\text{V.2})$$

which is the average of the inverted squared Euclidean distance between all pairs of processes. The desired objective function comes by adding Eqs. (V.1) and (V.2):

$$\frac{1}{|W|} \sum_{(i,j) \in W} \|x_i - x_j\|^2 + \frac{2}{P(P-1)} \sum_{i=0}^{P-2} \sum_{j=i+1}^{P-1} \frac{1}{\|x_i - x_j\|^2}. \quad (V.3)$$

The first term in Eq. (V.3) acts to bring the communicating pairs of processes close together while the second term ensures that no two processes reside at the same position.

Combining the objective function of Eq. (V.3) with the constraint that the processes must be located on the surface of an n -dimensional hypersphere results in the following constrained optimization problem:

$$\begin{aligned} \text{minimize} \quad & \left\{ \frac{1}{|W|} \sum_{(i,j) \in W} \|x_i - x_j\|^2 \right. \\ & \left. + \frac{2}{P(P-1)} \sum_{i=0}^{P-2} \sum_{j=i+1}^{P-1} \frac{1}{\|x_i - x_j\|^2} \right\} \end{aligned} \quad (V.4)$$

$$\text{subject to} \quad \|x_i\|^2 = 1, \quad \text{for all } i \in \mathcal{P}. \quad (V.5)$$

B. Application of the Gradient Projection Method

Note that the constrained optimization problem of Eqs. (V.4) and (V.5) can be expressed as:

$$\text{minimize} \quad f(x) \quad (V.6)$$

$$\text{subject to} \quad g(x) = 0, \quad (V.7)$$

where

$$f(x) = \left\{ \frac{1}{|W|} \sum_{(i,j) \in W} \|x_i - x_j\|^2 + \frac{2}{P(P-1)} \sum_{i=0}^{P-2} \sum_{j=i+1}^{P-1} \frac{1}{\|x_i - x_j\|^2} \right\}, \quad (V.8)$$

$$g(x) = \begin{pmatrix} \|x_0\|^2 - 1 \\ \|x_1\|^2 - 1 \\ \vdots \\ \|x_{P-1}\|^2 - 1 \end{pmatrix}, \quad (V.9)$$

and

$$x = \begin{pmatrix} x_0 \\ x_1 \\ \vdots \\ x_{P-1} \end{pmatrix}, \quad (V.10)$$

Derivation of the gradient. Recall from Eq. (IV.3) that the first part of each iteration of the gradient projection method requires the gradient of $f(x)$, i.e., $\nabla f(x)$. By differentiating terms from $f(x)$, it can be verified that

$$\frac{\partial}{\partial x_i} \{\|x_i - x_j\|^2\} = 2(x_i - x_j),$$

$$\frac{\partial}{\partial x_j} \{\|x_i - x_j\|^2\} = 2(x_j - x_i),$$

$$\frac{\partial}{\partial x_i} \left\{ \frac{1}{\|x_i - x_j\|^2} \right\} = \frac{2}{\|x_i - x_j\|^4} (x_j - x_i),$$

and

$$\frac{\partial}{\partial x_j} \left\{ \frac{1}{\|x_i - x_j\|^2} \right\} = \frac{2}{\|x_i - x_j\|^4} (x_i - x_j).$$

Define $\mathcal{W}(i) = \{j: j \in \mathcal{P} \text{ \& process } i \text{ and process } j \text{ communicate}\}$. Also, define $\mathcal{P}(i) = \mathcal{P} - \{i\}$. With these definitions, the gradient of $f(x)$ is expressed as

$$\begin{aligned} \nabla f(x) = & \begin{pmatrix} \frac{2}{|W|} \sum_{j \in \mathcal{W}(0)} (x_0 - x_j) \\ \frac{2}{|W|} \sum_{j \in \mathcal{W}(1)} (x_1 - x_j) \\ \vdots \\ \frac{2}{|W|} \sum_{j \in \mathcal{W}(P-1)} (x_{P-1} - x_j) \end{pmatrix} \\ & - \begin{pmatrix} \frac{4}{P(P-1)} \sum_{j \in \mathcal{P}(0)} \frac{(x_0 - x_j)}{\|x_0 - x_j\|^4} \\ \frac{4}{P(P-1)} \sum_{j \in \mathcal{P}(1)} \frac{(x_1 - x_j)}{\|x_1 - x_j\|^4} \\ \vdots \\ \frac{4}{P(P-1)} \sum_{j \in \mathcal{P}(P-1)} \frac{(x_{P-1} - x_j)}{\|x_{P-1} - x_j\|^4} \end{pmatrix} \end{aligned} \quad (V.11)$$

Thus, the first part of the iteration for the gradient projection method (i.e., moving in the direction of the negative gradient) is given by

$$x^{(k+1)} \leftarrow x^{(k)} - \alpha \nabla f(x)|_{x=x^{(k)}}, \quad (V.12)$$

where $\nabla f(x)|_{x=x^{(k)}}$ is computed according to Eq. (V.11).

Derivation of the projection. After the update of Eq. (V.12) is computed, the second part of the gradient projection method requires that the resultant $x^{(k+1)}$ be pro-

jected back to the surface $g(x) = 0$. This amounts to simply normalizing each x_i to be unit length (in the Euclidean sense). Thus, the projection operator is defined by

$$[x]_{g(x)=0} = \begin{pmatrix} \frac{x_0}{\|x_0\|} \\ \frac{x_1}{\|x_1\|} \\ \vdots \\ \frac{x_{p-1}}{\|x_{p-1}\|} \end{pmatrix}. \quad (V.13)$$

So, the second part of the iteration for the gradient projection method is given by

$$x^{(k+1)} \leftarrow [x^{(k+1)}]_{g(x)=0}, \quad (V.14)$$

where $[x^{(k+1)}]_{g(x)=0}$ is computed according to Eq. (V.13).

After a sufficient number of iterations, the gradient projection method converges to a fixed-point, i.e., the vector $x^{(k)}$ converges to a fixed-point x^* as k increases.

C. Discretizing onto the Hypercube

The final part of hypersphere mapper is to convert the solution vector x^* , which has continuous real components, to z^* , which has discrete components in the set $\{0, 1\}$. The conversion is done by simply discretizing each component of x (to either zero or one) according to their signs. Formally, for any $t \in \mathcal{R}$, the unit step function is defined as

$$u(t) = \begin{cases} 1 & \text{if } t \geq 0 \\ 0 & \text{if } t < 0. \end{cases}$$

Therefore, the discretization is defined by

$$z^* = u(x^*),$$

where it is understood that $u(\cdot)$ is applied to each component of x^* .

D. An Illustrative Example

The goal here is to illustrate how the process locations move around the surface of the hypersphere from one iteration to the next. Of particular interest is to show how the relative locations of the processes (for a given communication pattern, \mathcal{W}) affect the iterative migration to the fixed-point.

Consider the problem of mapping eight processes, $\mathcal{P} = \{0, 1, \dots, 7\}$, onto the nodes of a three-dimensional hyper-




FIG. 2. Graphical view of the example process communication pattern $\mathcal{W} = \{(0,4), (0,7), (1,7), (1,6), (2,4), (2,5), (3,5), (3,6)\}$.

cube. Assume the communication pattern is $\mathcal{W} = \{(0,4), (0,7), (1,7), (1,6), (2,4), (2,5), (3,5), (3,6)\}$. A graphical view of the assumed process communication pattern, i.e., \mathcal{W} , is shown in Fig. 2. The resulting iteration data from hypersphere mapper is given in Table I.

From Table I, note that the initial values of the continuous location vectors (i.e., the values of $x_i^{(0)}$, $i = 0, 1, \dots, 7$) were selected so that process i is (initially) mapped to node i , $i = 0, 1, \dots, 7$. After each iteration k , the resulting values of the continuous location vectors, $x_i^{(k)}$, and the associated discretized location vectors, $z_i^{(k)}$, are given. Also, the corresponding values of the continuous objective function, $f(x^{(k)})$ (refer to Eq. (V.8)), and the associated average Hamming distance, $f_H(z^{(k)})$ (refer to Eq. (II.1)), are given. The fact that the value of $f_H(x^{(k)})$ decreases as k increases indicates that the continuous objective function does indeed capture the essence of the discrete mapping problem.

Note that the process to node mapping produced by hypersphere mapper (taken from iteration $k = 10$ of Table I) is

- process 0 \rightarrow node 4
- process 1 \rightarrow node 7
- process 2 \rightarrow node 0
- process 3 \rightarrow node 3
- process 4 \rightarrow node 4
- process 5 \rightarrow node 1
- process 6 \rightarrow node 7
- process 7 \rightarrow node 5,

which is *not* a one-to-one mapping because two processes are mapped to node 4 and two processes are mapped to node 7 (and no processes are mapped to node 2 nor node 6). The average Hamming distance between communicating processes for this mapping is 0.75.

It can be verified (through an exhaustive search, for

TABLE I
Iterations from Hypersphere Mapper

k	x_0^k z_0^k	x_1^k z_1^k	x_2^k z_2^k	x_3^k z_3^k	x_4^k z_4^k	x_5^k z_5^k	x_6^k z_6^k	x_7^k z_7^k	$f(x^k)$ $f_H(z^k)$
0	-0.66	0.61	-0.34	0.66	-0.03	0.09	-0.45	0.61	4.12
	-0.56	-0.75	0.07	0.55	-0.80	-0.91	0.61	0.55	
	-0.50	-0.26	-0.94	-0.50	0.59	0.41	0.65	0.57	
	0	1	0	1	0	1	0	1	2.25
1	-0.66	0.78	-0.36	0.69	-0.51	0.46	-0.33	0.73	3.06
	-0.68	-0.61	-0.16	0.59	0.15	-0.62	0.70	0.41	
	-0.32	-0.14	-0.92	-0.40	0.84	-0.63	0.63	0.54	
	0	1	0	1	0	1	0	1	2.00
2	-0.73	0.88	-0.38	0.68	-0.71	0.35	-0.07	0.81	2.31
	-0.68	-0.42	-0.17	0.61	-0.15	-0.47	0.87	0.21	
	-0.05	0.22	-0.91	-0.42	0.69	-0.81	0.48	0.54	
	0	1	0	1	0	1	0	1	1.50
9	-0.39	0.91	-0.62	0.57	-0.94	0.19	0.66	0.33	1.38
	-0.70	-0.01	-0.27	0.59	-0.33	-0.05	0.72	-0.34	
	0.60	0.42	-0.74	-0.58	0.07	-0.98	0.20	0.88	
	0	1	0	1	0	1	1	1	0.75
10	-0.39	0.89	-0.62	0.56	-0.93	0.16	0.68	0.32	1.37
	-0.68	0.01	-0.27	0.58	-0.33	-0.03	0.71	-0.36	
	0.61	0.45	-0.74	-0.59	0.05	-0.99	0.18	0.88	
	0	1	0	1	0	1	1	1	0.75

example) that an optimal *one-to-one* mapping is given by (i.e., 0.75 versus 1.00) at the “expense” of not being one-to-one.

- process 0 → node 5
- process 1 → node 6
- process 2 → node 0
- process 3 → node 3
- process 4 → node 4
- process 5 → node 1
- process 6 → node 2
- process 7 → node 7.

The above optimal one-to-one mapping has an associated average Hamming distance of 1.00. Thus, for this example, the solution from hypersphere mapper produces a mapping that is superior, in terms of the average Hamming distance, to the best possible one-to-one mapping

E. Converting Hypersphere Mapper Solutions into Uniform Mappings

An intuitive approach was developed for modifying the continuous solution produced by hypersphere mapper so that after discretization, the resulting mapping will be uniform (or to within a specified degree of closeness to being uniform). The basic idea of the approach is the incrementally spread (in n phases) the continuous process locations out of “overpopulated” sectors and into nearby “underpopulated” sectors. The detailed description of the spreading approach can be found in [1].

The spreading procedure was applied to the illustrative example of the previous subsection (i.e., the problem of

mapping eight processes onto eight nodes, assuming the communication pattern \mathcal{W} of Fig. 2). From the hypersphere mapper solution (without application of the spreading procedure), processes 0 and 4 are both mapped to node 4, processes 1 and 6 are both mapped to node 7, and no processes are mapped to either node 2 nor 6. After one phase of the spreading procedure, process 4 is moved from the overpopulated node 4 onto the underpopulated node 6. After the second spreading phase, process 6 is moved from the overpopulated node 7 onto the underpopulated node 2.

VI. EVALUATION OF HYPERSPHERE MAPPER

A. Comparison with Other Hypercube Embedding Heuristics

In Ref. [7], evaluations and comparisons (based on simulation studies) are reported for twelve hypercube embedding heuristics. Several classes of process communication patterns are assumed, including random patterns, permuted hypercube patterns, and tree patterns. Of the heuristics evaluated (and for all of the assumed process communication patterns), the simulated annealing heuristic developed in [7] produced mappings with the smallest average Hamming distance.¹ In contrast, a greedy algorithm of [5] generally produced the poorest mappings (excluding the so-called default or random mapping schemes, which map processes to nodes independent of the given process communication pattern).

In the present paper, simulation studies are used to provide an indication of how the quality of mappings produced by hypersphere mapper compare to the quality of mappings produced by the other known embedding heuristics. The simulation study carried out in this subsection involves mapping 128 processes onto the 128 nodes of a seven-dimensional hypercube. The assumed communication pattern is generated by selecting process source-destination pairs at random. The expected number of randomly selected source-destination pairs is set at 448, i.e., $E[|\mathcal{W}|] = 448$. This exact same simulation study was one of the experiments used in evaluating the heuristics in [7].

The results of simulation studies for the hypersphere mapper, simulated annealing, greedy, and default algorithms (averaged over 100 runs) are summarized in Table II. The column labeled $\hat{f}_H(z)$ denotes the average Hamming distance between mapped communicating processes. The column labeled $\hat{\sigma}_{\text{pnn}}^2$ denotes the sample variance of the number of processes mapped to each node,

¹ Refer to [9] for general background material on the simulated annealing approach.

TABLE II
Comparison between Hypersphere Mapper and Other Heuristics: 128 Processes onto 128 Nodes with $E[|\mathcal{W}|] = 448$

Algorithm	$\hat{f}_H(z)$	$\hat{\sigma}_{\text{pnn}}^2$
Hypersphere mapper	1.889	1.68
Simulated annealing	2.042	0.00
Greedy	2.867	0.00
Default	3.503	0.00

which is defined by

$$\hat{\sigma}_{\text{pnn}}^2 = \frac{1}{N} \sum_{i=0}^N \left(\text{pnn}(i) - \frac{P}{N} \right)^2,$$

where $\text{pnn}(i)$ denotes the number of processes mapped to node i . The term P/N represents the average number of nodes mapped to each node, which equals unity for the simulation of this subsection because there are $P = 128$ processes being mapped onto $N = 128$ nodes. (In general, the average number of processes mapped onto each node is P/N , regardless of the mapping.)

From Table II, note that the value of $\hat{\sigma}_{\text{pnn}}^2$ is nonzero only for the hypersphere mapper because the mappings produced by hypersphere mapper are not (generically) one-to-one. On the other hand, the competing algorithms always produce one-to-one mappings and thus have zero values for $\hat{\sigma}_{\text{pnn}}^2$ (i.e., exactly one process is mapped to each node, thus the variance of the number of processes mapped to each node is zero). The results tabulated for the simulated annealing and greedy algorithms were taken from Ref. [7]. The last row of Table II gives the results for a default mapping algorithm, which simply maps process i to node i , for all $i = 0, 1, \dots, 127$.

With regard to average Hamming distance, the mappings produced by hypersphere mapper are indeed superior to those produced by the simulated annealing algorithm (i.e., 1.889 versus 2.042). However, the mappings produced by hypersphere mapper are not generically one-to-one; thus, the corresponding value of $\hat{\sigma}_{\text{pnn}}^2$ is not equal to zero.

Table III shows the effect of applying the spreading procedure to the hypersphere mapper solution. The average Hamming distance of the mapping produced after spreading phase 1 is slightly better than that of the simulated annealing algorithm (i.e., 2.020 versus 2.042) and the value of the variance of the number of processes mapped to each node, relative to the hypersphere mapper solution with no spreading, is reduced by about 60% (i.e., from 1.68 down to 1.07). It is also noted that the *one-to-one* mapping produced after spreading phase 7 is signifi-

TABLE III
Application of the Spreading Procedure to Hypersphere Mapper: 128 Processes onto 128 Nodes with $E\{|W|\} = 448$

Spreading	$\hat{f}_H(z)$	$\hat{\sigma}_{pmn}^2$
None	1.889	1.68
Phase 1	2.020	1.07
Phase 2	2.283	0.40
Phase 3	2.440	0.15
Phase 4	2.524	0.06
Phase 5	2.558	0.03
Phase 6	2.580	0.01
Phase 7	2.587	0.00

cantly better than the mapping produced by the greedy algorithm (i.e., 2.587 versus 2.867).

B. The Case of Having More Processes than Nodes

A primary attribute of hypersphere mapper is that problems involving more processes than nodes (i.e., $P > N$) can be solved directly. In contrast, the competing hypercube embedding heuristics (including the twelve evaluated in [7]) require $P \leq N$.

Simulation results are reported here for the case of mapping 256 processes onto 64 nodes. The assumed process communication patterns are randomly generated (with a specified expected number of source-destination pairs). Simulation studies are conducted for cases where the expected number of source-destination pairs is 128, 256, 512, and 1024. Table IV shows the results of the simulation studies in terms of average Hamming distance, $\hat{f}_H(z)$, and the sample variance of the number of processes mapped onto each node, $\hat{\sigma}_{pmn}^2$. As would be expected, the average Hamming distance increase as the expected number of source-destination pairs increases. Also, for a fixed value of the expected number of source-destination pairs, the average Hamming distance increases with successive applications of the spreading procedure. Simulations for a simple default mapper were also conducted, where the default mapper simply

mapped process i onto node j , where j is defined by masking (to zero) all but lower six bits of the binary representations of i . The simulations results for the default mapper were $\hat{f}_H(z) \approx 3.00$ and $\hat{\sigma}_{pmn}^2 = 0.00$, for all cases.

VII. CONCLUSIONS

A. Summary

A new hypercube embedding algorithm called hypersphere mapper was introduced. The hypersphere mapper approach is novel in that it utilizes techniques from nonlinear programming as a means of solving the mapping problem. In contrast, competing mapping algorithms are based on combinatoric heuristics. An important advantage of hypersphere mapper over existing mapping heuristics is that the case of mapping more than N processes onto N nodes is solved directly. In contrast, the competing heuristics are limited to problems where the number of processes is no larger than the number of nodes.

The variance of the number of processes mapped to each node is introduced as a means of quantifying the measure of "computational load balance" associated with a particular mapping. The average distance between mapped communicating processes is used to measure the communication load requirement on the interconnection network.

B. Ongoing and Future Work

By altering, slightly, the proposed objective function for hypersphere mapper, the resultant mappings can be tuned to have desired characteristics. For instance, by simply multiplying the average Euclidean distance component of the objective function (refer to Eq. (V.4)) by a scalar weighting factor, say γ , the nature of the resultant mappings can be changed. In particular, the mappings produced with $\gamma > 1$ tend to have smaller values of $\hat{f}_H(z)$ and larger values of $\hat{\sigma}_{pmn}^2$ (with no spreading), than the mappings produced with $\gamma \leq 1$. Also, the objective function can easily be adopted to accommodate a weight associated with each element in \mathcal{W} , which corresponds to the frequency with which the associated pair of processes

TABLE IV
Hypersphere Mapper Results: 256 Processes onto 64 Nodes

Spreading	$E\{ W \} = 128$		$E\{ W \} = 256$		$E\{ W \} = 512$		$E\{ W \} = 1024$	
	$\hat{f}_H(z)$	$\hat{\sigma}_{pmn}^2$	$\hat{f}_H(z)$	$\hat{\sigma}_{pmn}^2$	$\hat{f}_H(z)$	$\hat{\sigma}_{pmn}^2$	$\hat{f}_H(z)$	$\hat{\sigma}_{pmn}^2$
None	0.619	3.06	0.852	4.10	1.340	4.16	1.763	7.73
Phase 1	0.850	0.45	1.000	0.87	1.433	1.28	1.850	4.05
Phase 6	0.973	0.00	1.168	0.00	1.598	0.00	2.110	0.00

request communication. These weights correspond to an increased penalty for separating those process pairs that communicate frequently. Work is currently underway in evaluating a variety of different objective functions.

For cases where P^2 is large relative to $|W|$, the dominant computational time associated with each of the gradient projection iterations comes in computing the component of the gradient associated with the average inverted Euclidean distance between all pairs of processes (refer to Eqs. (V.4) and (V.11)). Instead of averaging the inverted Euclidean distance over all pairs of processes, an average could be taken over a smaller set of pairs of processes. The set W appears to be a natural candidate. Another possibility is to average over a randomly generated collection of process pairs. Some preliminary studies along these lines are currently underway. Work is also currently underway in integrating hypersphere mapper into an automated software engineering tool for high-performance computing systems.

REFERENCES

1. Antonio, J. K., and Metzger, R. C. Hypersphere mapper: A nonlinear programming approach to the hypercube embedding problem. *Proceedings, 7th International Parallel Processing Symposium, Newport Beach, CA, April 1993*, pp. 538–547.
2. Bazaraa, M. S., and Shetty, C. M. *Nonlinear Programming: Theory and Applications*. Wiley, New York, 1979.
3. Bhat, K., On the complexity of testing a graph for N -cube. *Inform. Process. Lett.* **11** (1980), 16–19.
4. Bettayeb, S., Miller, Z., and Sudborough, I. Embedding grids into hypercubes. *VLSI Algorithms and Architectures: 3rd Aegean Workshop on Computing*. Lecture Notes in Computer Science, Vol. 319. Springer-Verlag, New York/Berlin, 1988, pp. 201–211.
5. Chen, W.-K., and Gehringer, E. A graph oriented mapping strategy for a hypercube. *Proceedings, Third Conference on Hypercube Concurrent Computers and Applications, 1988*, pp. 200–209.
6. Cybenko, G., Krumme, D., and Venkataraman, K. Fixed hypercube embedding. *Inform. Process. Lett.* **25** (1987), 35–39.
7. Chen, W.-K., Stallman, M. F. M., and Gehringer, E. F. Hypercube embedding heuristics: An evaluation. *Int. J. Parallel Programming* **18**, 6 (1989), 505–549.
8. Fox, G. A graphical approach to load balancing and sparse matrix vector multiplication on the hypercube. Martin Schults (Ed.). *The IMA Volumes in Mathematics and its Applications*, Vol. 13, Springer-Verlag, New York/Berlin, 1988.
9. Kirkpatrick, S., Gelatt, C., Jr., and Vecchi, M. Optimization by simulated annealing. *Science* (1983), 671–680.
10. Wagner, A. Embedding arbitrary binary trees in a hypercube. *J. Parallel Distrib. Comput.* **7** (1989), 503–520.
11. Wu, A. Embedding of tree networks into hypercubes. *Int. J. Parallel Distrib. Comput.* **2** (1985), 238–249.

JOHN K. ANTONIO received the B.S. degree (with honors), the M.S. degree, and the Ph.D. all in electrical engineering from Texas A&M University, College Station in 1984, 1986, and 1989, respectively. Dr. Antonio worked as a control systems engineer with the Digital Flight Control Group, General Dynamics/Fort Worth Division, during the summers of 1983 and 1984. During the summer of 1986, he was employed by General Motors Research Laboratories, Warren, Michigan, as a research associate in the Control Systems Group. From the fall of 1986 until 1989, he held the position of Lecturer in the Electrical Engineering Department at Texas A&M University, while pursuing the Ph.D. After graduation, he accepted employment at Purdue University, where he currently holds the position of Assistant Professor of Electrical Engineering. During the summers of 1991–93, he participated in an AFOSR sponsored research programs at Rome Laboratory, Rome, NY, where he conducted research in the area of high performance computing. His current research interests include analysis and design of fast distributed algorithms for high-speed computer networks, iterative methods for solving large scale computational problems, computational aspects of control and optimization. He is a member of the IEEE computer society, control systems society, and communications society. Also, he is a member of Tau Beta Pi, Eta Kappa Nu, and Phi Kappa Phi. In 1989 he received the Ruth and Joel Spira "Outstanding Teacher Award," from the School of Electrical Engineering at Purdue University.

RICHARD C. METZGER is a computer engineer at Rome Laboratory, where he is a member of the Software Engineering for High Performance Computing group. His research interests include real-time computing on parallel systems, task allocation, and software engineering. He obtained a B.S. in medical technology from Upstate Medical Center in 1986, and a dual B.S. in computer science and information science from Syracuse University in 1990. He is currently pursuing an M.S. degree in computer engineering from Syracuse University. Metzger is a member of Phi Kappa Phi and Alpha Sigma Lambda.