COMPUTER SCIENCE PUBLICATION

LOAD BALANCING WITH THE DIFFUSION

METHOD ON K-ARY N-CUBES AND VARIANTS

C.Z. Xu and F.C.M. Lau

Technical Report TR-93-03

March 1993

DEPARTMENT OF COMPUTER SCIENCE
FACULTY OF ENGINEERING
UNIVERSITY OF HONG KONG
POKFULAM ROAD
HONG KONG

Load Balancing with The Diffusion Method on k-ary n-cubes and Variants

Cheng-Zhong Xu and Francis C.M. Lau

Department of Computer Science, The University of Hong Kong

{czxu, fcmlau}@csd.hku.hk

Abstract

The diffusion method is a simple distributed load balancing method for distributed memory multiprocessors. It operates in a relaxation fashion for point-to-point networks. Its convergence to the balanced state relies on the value of a parameter—the *diffusion parameter*. An optimal diffusion parameter would lead to the fastest convergence of the method. Previous results on optimal parameters have existed for the binary n-cubes. In this paper, we derive optimal diffusion parameters for the k-ary n-cube network and its variants—the ring, the torus, the chain and the mesh. We also derive the relationship between the optimal convergence rates of the method in these networks, and conclude that torus has the fastest convergence.

Keywords: diffusion method, distributed scheduling, k-ary n-cube networks, load balancing, message-passing multiprocessors.

1 Introduction

We consider the problem of dynamic load balancing in a distributed memory message-passing multiprocessor with N processors. From time to time during a parallel computation, the system is faced with the following problem of load balancing: to redistribute the system workload $(w_1, w_2, \ldots, w_N)$, where w_i is a nonnegative real number representing the workload in processor i, such that each processor ends up with the same $w = \sum w_i/N$. We study a distributed solution to the problem known as the *diffusion method.* The method is iterative in nature, and is particularly suitable for point-to-point networks with processors that are capable of parallel communications with their neighbors [11]. Willebeek-LeMair and Reeves employed this method in distributed computation of branch-and-bound algorithms on the Intel iPSC/2 [8, 9]. Lüling and Monien implemented it on transputer networks with the deBruijn and ring topology [6]. The efficiency of the diffusion method depends on a *diffusion parameter* which dictates how much of the excess workload between a pair of directly connected processors is to be transferred away at each iteration step. Cybenko was the first one to analyze the method. Under the assumption of synchronous communications, he derived a sufficient and necessary condition for the diffusion parameter to cause convergence [2]. Similar convergence results for hypercube, generalized hypercube, torus, and ring were obtained by Hong *et al.* [5] and Qian and Yang [7]. Cybenko also obtained the optimal diffusion parameter for the binary n-cube, which is the only optimal result for the diffusion method to date.

In this paper, we derive the optimal diffusion parameters for the family of k-ary n-cube networks. A k-ary n-cube is an order-k cube with n dimensions [3]. Besides the binary n-cube, popular networks like the ring, the chain, the torus, and the mesh are either k-ary n-cubes or isomorphic to k-ary n-cubes. In addition, we also establish the relationship between the optimal convergence rates of these networks. In the remainder of this paper, we review the diffusion method in Section 2 followed by derivations of the optimal parameters for k-ary n-cube networks and their variants in the next two sections. The theoretical results are validated through simulation, which we report in Section 5.

2 The diffusion method

The system we consider consists of a finite set of autonomous, homogeneous processors connected by a point-to-point communication network. The links are bi-directional and the processors interact synchronously with one another. Such a system can be depicted as a simple connected graph $G = (V, E)$, where V is a set of vertices (nodes) labeled from 1 to N and $E \subseteq V \times V$ is a set of edges. Each vertex represents a processor and each edge $(i, j) \in E$ represents a communication link between processor i and j. For each vertex i in G, let $A(i)$ denote the set of its adjacent vertices, and $deg(i)$ be the cardinality of the set, *i.e.*, $deg(i) = |A(i)|$.

The execution of the load balancing procedure is divided into a sequence of steps. At each step, a processor would interact and exchange load with all its direct neighbors. Specifically, for processor i, the change of workload is executed as

$$w_i = w_i + \sum_{j \in A(i)} \alpha(w_j - w_i) \tag{1}$$

where w_i and w_j are the current local workloads of processor i and j respectively, and α is the *diffusion parameter* which determines the portion of excess of workload to be diffused away. Let t be the step index, $t = 0, 1, 3, \ldots$, and $w_i^t (1 \leq i \leq N)$ be the local workload of processor i at step t. Then the overall workload distribution at step t is denoted by the vector $W^t = (w_1^t, w_2^t, \ldots, w_N^t)^T$ in the transposed form. W^0 is the initial workload distribution. The change of the workload distribution in the system at step t can be modeled by the equation

$$W^{t+1} = D(\alpha) W^t \tag{2}$$

where $D(\alpha) = (d_{ij})$, called a *diffusion matrix*, is given by

$$d_{ij} = \begin{cases} \alpha & \text{if } i \neq j \text{, and } i \text{ and } j \text{ are directly connected} \\ 1 - deg(i)\alpha & \text{if } i = j \\ 0 & otherwise \end{cases}$$

For example, the diffusion matrix of a chain of order 4 (*i.e.*, with 4 nodes) is equal to

$$\begin{pmatrix} 1-\alpha & \alpha & & \\ \alpha & 1-2\alpha & \alpha & \\ & \alpha & 1-2\alpha & \alpha \\ & & \alpha & 1-\alpha \end{pmatrix}$$

Given the diffusion matrix, two questions are in order: whether the sequence $\{D^t(\alpha)\}$ is convergent; if it does, what is the asymptotic convergence rate, denoted by $R_\infty(D(\alpha))$.

In light of the various properties of $D(\alpha)$ (nonnegative, symmetric and doubly stochastic) and under the constraints of $\alpha \geq 0$ and $1 - deg(i)\alpha \geq 0$ for each i, Cybenko showed that the diffusion method converges if and only if the system graph is not bipartite or $1 - deg(i)\alpha > 0$ for some i [2]. This is true for any given initial load distribution.

Regarding the convergence rate, we need to consider the eigenvalue spectrum of $D(\alpha)$. Let $\mu_j(D(\alpha))$ $(1 \leq j \leq N)$ be the eigenvalues of $D(\alpha)$, and $\rho(D(\alpha))$ and $\gamma(D(\alpha))$ be the dominant (largest) and subdominant (second largest) eigenvalues respectively of $D(\alpha)$ in modulus. Because of the above properties of $D(\alpha)$, $\rho(D(\alpha))$ is unique and equal to 1; therefore $R_\infty(D(\alpha))$ is determined by $\gamma(D(\alpha))$ and is equal to $-\ln\gamma(D(\alpha))$ [1]. Our task is then to choose an α that would minimize $\gamma(D(\alpha))$ while preserving the nonnegativity of $D(\alpha)$. Cybenko proved $\alpha = 1/(n+1)$ to be the optimal choice for binary n-cubes in [2]. In the sequel, we derive the optimal $\alpha_{opt}(D(\alpha))$ for k-ary n-cube networks, $k > 2$, and their variants, and examine the relationship between the optimal convergence rates of these networks.

3 Diffusion method on k-ary n-cube networks

In this section, we analyze the diffusion method as applied to k-ary n-cube networks. The analysis by induction on the dimension $n, n \geq 1$. We begin with the k-ary 1-cube network, *i.e.*, a ring of order k, and then generalize it to n-dimensional $k_1 \times k_2 \times \ldots \times k_n$ torus ($k_i > 2,\ i = 1, 2, \ldots, n$). A k-ary n-cube is a special case of the torus network with the same order in each dimension. Notice that an even torus (even order in every dimension) is bipartite and therefore, according to Cybenko's theorem, the diagonal elements of the corresponding diffusion matrix must be positive in order that the diffusion process would converge—that is, $\alpha < 1/(2n)$.

Our derivation relies on the theory of *block circulant matrices* [4]. Circulant matrices are most appropriate here because the diffusion matrices of the ring and the torus are in circulant and block circulant form respectively. Let $A_1, A_2, \ldots, A_m$ be square matrices of order n. Then a *block circulant matrix* is a matrix of the form

$$\Phi_{m,n}(A_1, A_2, \ldots, A_m) = \begin{pmatrix} A_1 & A_2 & \ldots & \ldots & A_m \\ A_m & A_1 & \ldots & \ldots & A_{m-1} \\ & \ldots & & \ldots & \\ A_2 & A_3 & \ldots & \cdots & A_1 \end{pmatrix}$$

If $n = 1$, a block circulant matrix degenerates to a circulant matrix. The following useful lemma can be easily derived based on circulant matrix theory [4].

Lemma 3.1 *Let matrix* $A = \Phi_{m,n}(A_1, A_2, \ldots, A_m)$. *Then, the eigenvalues of the matrix* A *are those of matrices* $A_1 + \omega^j A_2 + \ldots + \omega^{j(m-1)} A_m$, $j = 0, 1, \ldots, m-1$, *where* $\omega^j = \cos\frac{2\pi j}{m} + \mathrm{i}\sin\frac{2\pi j}{m}$, $\mathrm{i} = \sqrt{-1}$. *In particular, if* $n = 1$, $\mu_j(A) = A_1 + \omega^j A_2 + \ldots + \omega^{j(m-1)} A_m$.

3.1 k-ary 1-cube networks

By the definition of the diffusion matrix, the following is easily seen.

Lemma 3.2 *Let* DR_k *be the diffusion matrix of the k-ary 1-cube network. Then,* $DR_k = \Phi_{k,1}(1 - 2\alpha, \alpha, 0, \ldots, 0, \alpha)$.

Theorem 3.1 *The optimal diffusion parameter for the k-ary 1-cube network,* $\alpha_{opt}(DR_k)$, *is equal to* $1/(3 - \cos(2\pi/k))$ *if* k *is even, and* $1/(2 + \cos(\pi/k) - \cos(2\pi/k))$ *otherwise. Moreover,* $R_\infty(DR_k) > R_\infty(DR_{k+2})$.

Proof. From Lemma 3.1, it follows that

$$\begin{aligned} \mu_j(DR_k) &= 1 - 2\alpha + \alpha e^j + \alpha e^{k-j} \\ &= 1 - 2\alpha + 2\alpha\cos(\frac{2\pi j}{k}),\ j = 0, 1, \ldots, k-1. \end{aligned}$$

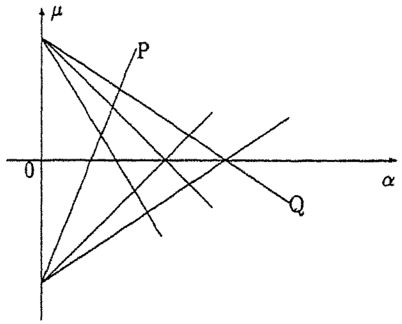

Figure 1: The eigenvalues in modulus μ of $DR_k(\alpha)$ versus the diffusion parameter α

Note that $e^{k-j} = \cos(\frac{2\pi j}{k}) - \mathrm{i}\sin(\frac{2\pi j}{k})$. We want to determine the value of α such that the subdominant eigenvalue in modulus, $\gamma(DR_k)$, is minimized. Since each eigenvalue in modulus μ is linearly dependent on α, it is easy to see that $\gamma(DR_k)$ is minimized at the intersection of the lines P and Q

$$P: \quad \begin{cases} \mu = 4\alpha - 1 & \text{if } k \text{ is even} \\ \mu = 2\alpha + 2\alpha\cos(\frac{\pi}{k}) - 1 & \text{if } k \text{ is odd} \end{cases}$$

$$Q: \quad \mu = 1 - 2\alpha + 2\alpha\cos(\tfrac{2\pi}{k})$$

which are the lines with the steepest and the flattest slopes respectively in the plot of μ versus α, as illustrated in Figure 1. It is then clear that the subdominant (second largest) eigenvalue in modulus $\gamma(DR_k)$ of DR_k is minimized at the intersection point of lines P and Q whose abscissa corresponds to

$$\alpha = \begin{cases} 1/(3 - \cos(\frac{2\pi}{k})) & \text{if } k \text{ is even} \\ 1/(2 + \cos(\frac{\pi}{k}) - \cos(\frac{2\pi}{k})) & \text{if } k \text{ is odd} \end{cases}$$

These values of α preserve the nonnegativity of DR_k because both are less than 1/2. Substituting these values for α in the equation for the eigenvalues yields the optimal diffusion parameter

$$\gamma(DR_k) = \begin{cases} 4/(3 - \cos(\frac{2\pi}{k})) - 1 & \text{if } k \text{ is even} \\ 2/(3 - 2\cos(\frac{\pi}{k})) - 1 & \text{if } k \text{ is odd} \end{cases}$$

Evidently, $\gamma(DR_{k+2}) > \gamma(DR_k)$, and hence $R_\infty(DR_{k+2}) < R_\infty(DR_k)$. □

For example, the optimal diffusion parameter of the 4-ary 1-cube (equivalent to a binary 2-cube) is 1/3 according to this theorem, which matches Cybenko's binary n-cube result [2]. This theorem also says that the more nodes the ring has, the slower the convergence of the load balancing procedure, which is not unexpected.[1]

3.2 k-ary n-cube networks

On the basis of the above result for k-ary 1-cubes (rings), we now consider two-dimensional $k_1 \times k_2$ tori ($k_1 > 2$, $k_2 > 2$) as depicted in Figure 2. For simplicity, we assume both k_1 and k_2 are either even or odd. The omitted cases of k_1 even and k_2 odd and vice versa can be analyzed in much the same way. As the spectrum of eigenvalues of the diffusion matrix of a network is invariant under any permutation of the node labels, we therefore label the nodes in the "row major" fashion as shown in Figure 2. In the following, I_k denotes the identity matrix of order k. A two-dimensional torus can be viewed as a stack of rings of order k_1; so we can express its diffusion matrix in terms of the diffusion matrix of ring, as follows.

[1] We proved $R_\infty(DR_{k+2}) < R_\infty(DR_k)$ here, proving $R_\infty(DR_{k+1}) < R_\infty(DR_k)$ requires solving $1 + \cos^2(\pi/k) - 2\cos(\pi/(k+1)) \leq 0$ for $k \geq 2$

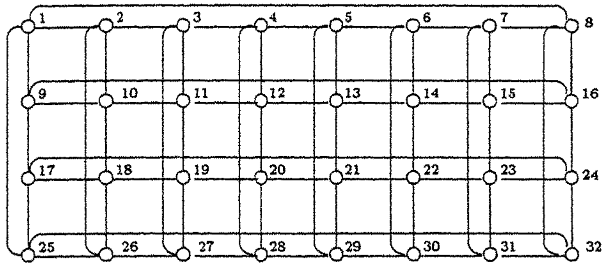

Figure 2: An example of a two-dimensional 8×4 torus

Lemma 3.3 *Let DT_{k_1,k_2} be the diffusion matrix of a two-dimensional $k_1 \times k_2$ torus. Then, $DT_{k_1,k_2} = \Phi_{k_2,k_1}(DR_{k_1} - 2\alpha I_{k_1}, \alpha I_{k_1}, 0, \ldots, 0, \alpha I_{k_1})$.*

Proof. Proved by induction on the order of the second dimension k_2. □

Theorem 3.2 *Let $k = \max\{k_1, k_2\}$. Then, the optimal diffusion parameter for the two-dimensional $k_1 \times k_2$ torus is*

$$\alpha_{opt}(DT_{k_1,k_2}) = \begin{cases} 1/(5 - \cos(2\pi/k)) & \text{if both } k_1 \text{ and } k_2 \text{ are even} \\ 1/(3 + \sum_{i=1}^{2} \cos(\pi/k_i) - \cos(2\pi/k)) & \text{if both } k_1 \text{ and } k_2 \text{ are odd} \end{cases}$$

Moreover, the convergence rate $R_\infty(DT_{k_1,k_2})$ is equal to $R_\infty(DT_{k,k})$ if both k_1 and k_2 are even.

Proof. From Lemma 3.1, the eigenvalues of the block circulant matrix DT_{k_1,k_2} are those of the matrices

$$\begin{aligned} & DR_{k_1} - 2\alpha I_{k_1} + \alpha e^{j_2} I_{k_1} + \alpha e^{k_2 - j_2} I_{k_1} \\ = \; & DR_{k_1} - 2\alpha I_{k_1} + 2\alpha \cos(\frac{2\pi j_2}{k_2}) I_{k_1}, \; j_2 = 0, 1, \ldots, k_2 - 1 \end{aligned}$$

Therefore,

$$\begin{aligned}\mu(DT_{k_1,k_2}) &= \mu(DR_{k_1}) - 2\alpha + 2\alpha\cos(\frac{2\pi j_2}{k_2}) \\ &= 1 - 4\alpha + 2\alpha\cos(\frac{2\pi j_1}{k_1}) + 2\alpha\cos(\frac{2\pi j_2}{k_2}),\end{aligned}$$

where $j_1 = 0, 1, \ldots, k_1 - 1$ and $j_2 = 0, 1, \ldots, k_2 - 1$. Then, as illustrated in Figure 1, the subdominant eigenvalue in modulus $\gamma(DT_{k_1,k_2})$ is minimized at the intersection point of the lines P and Q

$$P: \begin{cases} \mu = 8\alpha - 1 & \text{if both } k_1 \text{ and } k_2 \text{ are even} \\ \mu = 4\alpha + 2\alpha\cos(\pi/k_1) + 2\alpha\cos(\pi/k_2) - 1 & \text{if both } k_1 \text{ and } k_2 \text{ are odd} \end{cases}$$

$$Q: \quad \mu = 1 - 2\alpha + 2\alpha\cos(2\pi/k)$$

where $k = \max\{k_1, k_2\}$. The α values corresponding to this intersection point in each case are as that stated in the theorem. Substituting these optimal values for the diffusion parameter in the equation for the eigenvalues yields

$$\gamma(DT_{k_1,k_2}) = \begin{cases} \frac{8}{5-\cos(2\pi/k)} - 1 & \text{if both } k_1 \text{ and } k_2 \text{ are even} \\ \frac{4+2\sum_{i=1}^{2}\cos(\pi/k_i)}{3+\sum_{i=1}^{2}\cos(\pi/k_i)-\cos(2\pi/k)} - 1 & \text{if both } k_1 \text{ are } k_2 \text{ are odd} \end{cases}$$

Clearly, $\gamma(DT_{k_1,k_2}) = \gamma(DT_{k,k})$ in the even case. Hence, the theorem is proved. □

From the theorem, we see that the convergence of a two-dimensional torus depends only on the larger dimension when both k_1 and k_2 are even. For example, the tori $DT_{8,j}$, $j = 4, 6, 8$, all have the same optimal convergence rate with the optimal diffusion parameter $\alpha_{opt} = 2/(10 - \sqrt{2}) \approx 0.11647$.

The results of two-dimensional tori can be generalized to multi-dimensional tori. Consider an n-dimensional $k_1 \times k_2 \times \ldots \times k_n$ torus, $k_i > 2$ $(i = 1, 2, \ldots, n)$. Given any labeling of the nodes, by permutation, we can bring the diffusion matrix into the following iterative form.

$$DT_{k_1,k_2,\ldots,k_n} = \Phi_{k_n,\tilde{N}}(DT_{k_1,k_2,\ldots,k_{n-1}} - 2\alpha I_{\tilde{N}}, \alpha I_{\tilde{N}}, 0, \ldots, 0, \alpha I_{\tilde{N}})$$

where $\tilde{N} = k_1 \times k_2 \times \ldots \times k_{n-1}$. By induction on the number of dimensions n, it follows that

$$\mu(DT_{k_1,k_2,\ldots,k_n}) = 1 - 2n\alpha + 2\alpha \sum_{i=1}^{i=n} \cos(\frac{2\pi j_i}{k_i}),\ j_i = 0, 1, \ldots, k_i - 1$$

Using the technique in the proofs of the above two theorems, we obtain the following result.

Theorem 3.3 *Let $DT_{k_1,k_2,\ldots,k_n}$ be the diffusion matrix of an n-dimensional torus of $k_1 \times k_2 \times \ldots \times k_n$, and let $k = \max_{1\le i\le n}\{k_i\}$. Then, the optimal diffusion parameter is as follows*

$$\alpha_{opt}(DT_{k_1,k_2,\ldots,k_n}) = \begin{cases} \frac{1}{2n+1-\cos(2\pi/k)} & \text{if } k_i, i = 1, 2, \ldots, n, \text{ are even} \\ \min\{\frac{1}{2n}, \frac{1}{n+1+\sum_{i=1}^{n}\cos(\pi/k_i) - \cos(2\pi/k)}\} & \text{if } k_i, i = 1, 2, \ldots, n, \text{ are odd} \end{cases}$$

Moreover, the convergence rate $R_\infty(DT_{k_1,k_2,\ldots,k_n})$ is equal to $R_\infty(DT_{k,k,\ldots,k})$ in the even case.

We omit the tedious proof which is quite similar to those above. Notice that the alternative choice of $1/2n$ for $\alpha_{opt}(DT_{k_1,k_2,\ldots,k_n})$ in the odd case is for preserving the nonnegativity of the diffusion matrix. Since a k-ary n-cube network is a special case of an n-dimensional torus, we have the following.

Corollary 3.1 *Let $DT_{k;n}$ be the diffusion matrix of a k-ary n-cube network, $k > 2$. Then, its optimal diffusion parameter is*

$$\alpha_{opt}(DT_{k;n}) = \begin{cases} \frac{1}{2n+1-\cos(2\pi/k)} & \text{if } k \text{ is even} \\ \frac{1}{n+1+n\cos(\pi/k)-\cos(2\pi/k)} & \text{if } k \text{ is odd and } n \le 3 \\ \frac{1}{2n} & \text{if } k \text{ is odd and } n \ge 4 \end{cases}$$

4 Diffusion method on variants of k-ary n-cube networks

In this section, we consider chains and meshes which are variants of k-ary n-cubes without the end-round connections. Our analysis makes use of the concept of *direct product* of matrices [4]. Let A_m and B_n be square matrices of order m and n respectively. Then the *direct product* of A_m and B_n is a square matrix defined by

$$A_m \otimes B_n = \begin{pmatrix} a_{0,0}B_n & a_{0,1}B_n & \cdots & a_{0,m-1}B_n \\ a_{1,0}B_n & a_{1,1}B_n & \cdots & a_{1,m-1}B_n \\ & \cdots & & \cdots \\ a_{m-1,0}B_n & a_{m-1,1}B_n & \cdots & a_{m-1,m-1}B_n \end{pmatrix}$$

The following useful lemma follows from this definition.

Lemma 4.1 *Let A_m, B_n be square matrices of order m and n with eigenvalues $\mu_i(A_m)$ and $\mu_j(B_n)$, $i = 1, 2, \ldots, m$, $j = 1, 2, \ldots, n$, respectively. Then, the eigenvalues of $I_n \otimes A_m + B_n \otimes I_m$ are $\mu_i(A_m) + \mu_j(B_n)$, where I_n (I_m) is the identity matrix of order n (m).*

4.1 The chain

We first consider the chain network of order k, $k > 2$, and present its diffusion matrix in the following form.

Lemma 4.2 *Let DC_k be the diffusion matrix of a chain network of order k. $DC_k = (1-2\alpha)I_k + 2\alpha T_k$, where*

$$T_k = \begin{pmatrix} 1/2 & 1/2 & & & & \\ 1/2 & 0 & 1/2 & & & \\ & & \ddots & \ddots & & \\ & & & 1/2 & 0 & 1/2 \\ & & & & 1/2 & 1/2 \end{pmatrix}_{k\times k}$$

Observe that the matrix T_k is in the form of the transition matrix of the k-state elastic random walk with equal transition probabilities (1/2) in a Markov chain [1]. It follows that

$$\mu_j(T_k) = \cos(\pi j/k),\ j = 0, 1, \ldots, k-1$$

Theorem 4.1 *The optimal diffusion parameter for the chain network, $\alpha_{opt}(DC_k)$, is equal to 1/2. Moreover, the convergence rate $R_\infty(DC_k)$ is inversely proportional to k.*

Proof. By Lemma 4.2,

$$\begin{aligned} \mu_j(DC_k) &= 1 - 2\alpha + 2\alpha\mu_j(T_k) \\ &= 1 - 2\alpha + 2\alpha\cos(\pi j/k),\ j = 0, 1, \ldots, k-1 \end{aligned}$$

Then, as illustrated in Figure 1, the subdominant eigenvalue in modulus $\gamma(DC_k)$ is minimized at the intersection point of the lines of P and Q

$$P: \quad \mu = 2\alpha + 2\alpha\cos(\pi/k) - 1$$

$$Q: \quad \mu = 1 - 2\alpha + 2\alpha\cos(\pi/k)$$

It follows that $\alpha_{opt} = 1/2$, and $\gamma(DC_k) = \cos(\pi/k)$. And hence $R_\infty(DC_k)$ decreases with the increase of k. □

By comparing $\gamma(DC_k)$ with the minimum subdominant eigenvalue for the ring, $\gamma(DR_k)$, shown in Theorem 3.1, we obtain the following result immediately.

Corollary 4.1 *The optimal diffusion method for the ring network converges faster than that for the chain network of the same order.*

The above theorem reveals that the optimal diffusion parameter of a chain is fixed at 1/2 regardless of how long the chain may be. With the additional end-round connection, the ring converges faster than the chain, which is expected.

The chain is a special case of the mesh and will serve as a building block for our analysis of the mesh. In the next section, we first study the two-dimensional mesh, and then the n-dimensional mesh by induction on the number of dimensions n.

4.2 The mesh

We consider the two-dimensional $k_1 \times k_2$ mesh, $k_1 > 2$, $k_2 > 2$. Without loss of generality, we assume the nodes are indexed in the "row major" fashion. Then, by induction on the second dimension k_2, we obtain the diffusion matrix

$$DM_{k_1,k_2} = \begin{pmatrix} A_1 & A_3 & & & \\ A_3 & A_2 & A_3 & & \\ & & \ddots & \ddots & \\ & & A_3 & A_2 & A_3 \\ & & & A_3 & A_1 \end{pmatrix}$$

where $A_1 = DC_{k_1}(\alpha) - \alpha I_{k_1}$, $A_2 = DC_{k_1}(\alpha) - 2\alpha I_{k_1}$, $A_3 = \alpha I_{k_1}$. This matrix has $k_2 \times k_2$ block elements each of which is a matrix of $k_1 \times k_1$ nonnegative reals. We rewrite it in a more concise form in terms of direct products of matrices in the following lemma.

Lemma 4.3 *Let DC_k be the diffusion matrix of a chain of order k. Then,*

$$DM_{k_1,k_2} = I_{k_2} \otimes (DC_{k_1} - 2\alpha I_{k_1}) + T_{k_2} \otimes 2\alpha I_{k_1}$$

where I_{k_1} and I_{k_2} are identity matrices, and T_{k_2} is defined as in Lemma 4.2.

Theorem 4.2 *The optimal diffusion parameter for the $k_1 \times k_2$ mesh, $\alpha_{opt}(DM_{k_1,k_2})$, is equal to $1/4$. Moreover, the convergence rate $R_\infty(DM_{k_1,k_2})$ is equal to $R_\infty(DM_{k,k})$, where $k = \max\{k_1, k_2\}$.*

Proof. From Lemma 4.1 and Theorem 4.1, it follows that

$$\mu(DM_{k_1,k_2}) = 1 - 4\alpha + 2\alpha\cos(\frac{\pi j_1}{k_1}) + 2\alpha\cos(\frac{\pi j_2}{k_2})$$

where $j_1 = 0, 1, \ldots, k_1 - 1$ and $j_2 = 0, 1, \ldots, k_2 - 1$. Without loss of generality, assume $k_1 \geq k_2$. Then, the subdominant eigenvalue in modulus, $\gamma(DM_{k_1,k_2})$, is minimized at the intersection point of the lines P and Q:

$$P: \quad \mu = 4\alpha + 2\alpha\cos(\pi/k_1) + 2\alpha\cos(\pi/k_2) - 1$$

$$Q: \qquad\qquad \mu = 1 - 2\alpha + 2\alpha\cos(\pi/k_1)$$

That is, $\alpha = 1/(3+\cos(\pi/k_2))$. But this choice of α would lead to a negative element $1-4\alpha$ (*i.e.*, a node with four links) in DM_{k_1,k_2}. To preserve the nonnegativity of the diffusion matrix, we pick a value of α which is closest to the above α and which would make $1-4\alpha$ nonnegative. Hence, $\alpha_{opt} = 1/4$. Substituting this into the equation for $\mu(DM_{k_1,k_2})$ gives $\gamma(DM_{k_1,k_2}) = \frac{1}{2}+\frac{1}{2}\cos(\frac{\pi}{k_1})$. Therefore, $R_\infty(DM_{k_1,k_2}) = R_\infty(DM_{k,k})$, where $k = \max\{k_1, k_2\}$. □

Corollary 4.2 *The optimal diffusion method in a torus converges faster than that in a mesh of the same dimensions.*

That is, $\gamma(DM_{k_1,k_2}) > \gamma(DT_{k_1,k_2})$ or $R_\infty(DM_{k_1,k_2}) < R_\infty(DT_{k_1,k_2})$, which follows by comparing the result here with the result for torus in Theorem 3.2. Again, we see that the end-round connections help.

The above theorem says that the convergence in a mesh depends only on its larger dimension. For example, the meshes $DM_{8,j}$, $j = 4, 6, 8$, all have the same convergence rate for the fixed optimal diffusion parameter $\alpha = 1/4$.

These results for two-dimensional meshes can be generalized to n-dimensional $k_1 \times k_2 \times \ldots k_n$ ($k_i > 2, i = 1, 2, \ldots, n$) meshes whose diffusion matrix can be written in the following recursive form.

$$DM_{k_1,k_2,\ldots,k_n} = I_{k_n} \otimes (DM_{k_1,k_2,\ldots,k_{n-1}} - 2\alpha I_{\tilde{N}}) + DC_{k_n} \otimes I_{\tilde{N}}$$

where $\tilde{N} = k_1 \times k_2 \times \ldots \times k_{n-1}$. By induction on the number of dimensions n, it follows that

$$\mu(DM_{k_1,k_2,\ldots,k_n}) = 1 - 2n\alpha + 2\alpha \sum_{i=1}^{i=n} \cos(\frac{\pi j_i}{k_i}),\ j_i = 0, 1, \ldots, k_n - 1$$

Hence, we obtain the following results.

Theorem 4.3 *Let $DM_{k_1,k_2,\ldots,k_n}$ be the diffusion matrix of an n-dimensional $k_1 \times k_2 \times \ldots \times k_n$ mesh. Then, the optimal diffusion parameter $\alpha_{opt}(DM_{k_1,k_2,\ldots,k_n})$ is equal to $1/(2n)$. Moreover, the convergence rate $R_\infty(DM_{k_1,k_2,\ldots,k_n})$ is equal to $R_\infty(DM_{k,k,\ldots,k})$, where $k = \max\{k_1, k_2, \ldots, k_n\}$.*

Corollary 4.3 *The optimal diffusion method for an n-dimensional torus converges faster than that for an n-dimensional mesh of the same dimensions.*

5 Simulation

We have derived the optimal diffusion parameters which would lead to the fastest asymptotic convergence rate. For actual computations, it would be of considerable value to estimate the number of iterations required for the system to arrive at its load balanced state. Define ϵ^t to be the error vector from the load balanced state at the t-th iteration, *i.e.*, $\epsilon^t = W^t - \overline{W}$, where $\overline{W}$ is the balanced workload distribution. Then, from Equation 2, it follows that $||\epsilon^t||_2 \leq \gamma^t(D(\alpha))||\epsilon^0||_2$. Hence, from an initial workload distribution with error vector $||\epsilon^0||_2$, the number of iterations t required to reduce the error to some prescribed bound δ satisfies

$$t \geq \frac{\log \delta - \log ||\epsilon^0||_2}{\log \gamma(D(\alpha))} \tag{3}$$

From this inequality, we see that the iteration number is dependent upon such factors as the initial workload distribution, the topology and size of the underlying system structure, the prescribed relative error bound δ and the diffusion parameter α, as is illustrated in the following simulation results.

We simulated a few cases to obtain an idea of the iteration numbers required by the load balancing procedure under various choices of the diffusion parameters. We denote this number by NI. In addition to revealing the actual efficiency of the method, the results of the simulation also validate the theoretical results derived in the preceding sections.

The initial workload distribution is a random vector, each element of which is drawn independently from an identical uniform distribution. The amount of workload a processor gets is thus determined by the distribution mean. The relative error bound δ can be tuned to achieve the desired performance in practice. In our simulation experiments, this value is set to one. That is, the load balancing procedure continues until the Euclidean norm of the error vector is less than one. Figures 3–6 plot the expected iteration numbers in various networks for reaching the balanced state from initial workload distribution with a mean 128 as α varies in steps of 0.05 from 0.10 to the maximum value which preserves the nonnegativity of the corresponding diffusion matrix. The maximum value is 0.50 in the cases of rings and chains, and 0.25 in the cases of two dimensional tori and meshes. Notice that even rings and tori are bipartite graphes, and according to the Cybenko's necessary and sufficient conditions for convergence, the value of α should be less than the graph degree. Hence, we set the upper bound of α to 0.49 in even rings and 0.24 in even tori in our experiments.

To reduce the effect of the variance of the initial load distribution on the iteration numbers, we take the average of 100 runs for each data point, each run using a different random initial load distribution.

As can be seen from the figures, the expected number of iterations in each case for different α's vary with the value of $\gamma(\alpha)$. Specifically, the theoretically-proven optimal diffusion parameter of each case yields the best result in terms of the expected number of iterations. Also, the expected number of iterations agrees with $\gamma(\alpha)$ in their dependent relationships on the topologies and size of networks. In particular, from Figure 5, it is evident that the expected number of iterations of even tori is insensitive to the smaller dimensions. Finally, we point out that the relatively large values of NI are partially due to the high accuracy (error bound = 1) we prescribed in our experiments.

6 Conclusion

We have analyzed the diffusion method for dynamic load balancing as applied to k-ary n-cube networks and their variants—the ring, the chain, the mesh, and the torus. We have derived the optimal diffusion parameters in closed form, which maximize the convergence rates of the iterative load balancing procedure in these networks. We also showed that the torus network performs better than the mesh with the same number of processors. These theoretical results have been validated through simulation experiments. The diffusion method is one of the few iterative methods for dynamic load balancing [10]. The other well-known method in the category is the *dimension exchange* method. Cybenko proved that the dimension exchange method

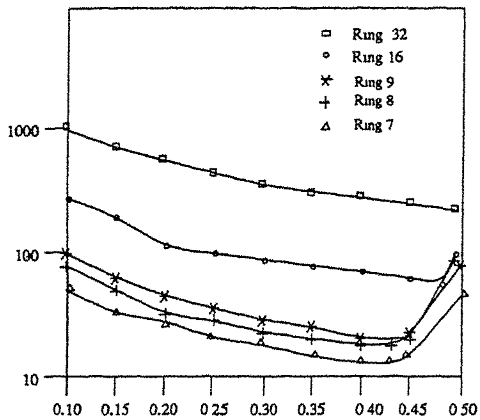

Figure 3: Iteration numbers versus α in rings

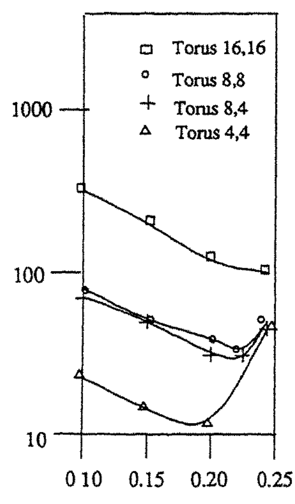

Figure 4: Iteration numbers versus α in torus

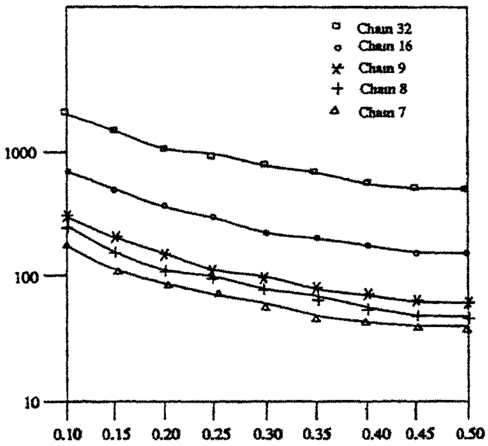

Figure 5: Iteration numbers versus α in chains

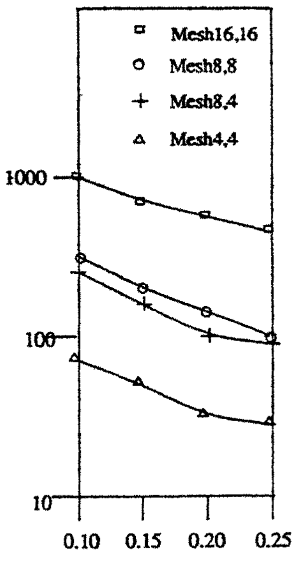

Figure 6: Iteration numbers versus α in meshes

is superior to the diffusion method in the case of the hypercube [2]. We tuned the method to achieve optimal performance in various structures including the mesh and the torus [10]. A comparison between the diffusion method and the dimension exchange method in practical situations is now underway.

References

[1] A. Berman and R.J. Plemmons. *Nonnegative matrices in the mathematical sciences.* Academic Press, 1979.

[2] G. Cybenko. Load balancing for distributed memory multiprocessors. *Journal of Parallel and Distributed Computing*, 7:279–301, 1989.

[3] W. J. Dally. Performance analysis of k-ary n-cube interconnection networks. *IEEE Transactions on Computers*, 39(6):775–785, June 1990.

[4] P. J. Davis. *Circulant matrices.* John Wiley and Sons Inc., 1979.

[5] J.-W. Hong, X.-N. Tan, and M. Chen. From local to global: an analysis of nearest neighbor balancing on hypercube. In *Proceedings of ACM-SIGMETRICS*, pages 73–82, May 1988.

[6] R. Lüling and B. Monien. Load balancing for distributed branch and bound algorithm. In *Proc. of 6th. Int. Parallel Processing Symposium*, pages 543–5448, California, USA, March 1992.

[7] X.-S. Qian and Q. Yang. Load balancing on generalized hypercube and mesh multiprocessors with lal. In *Proceedings of 11th International Conference on Distributed Computing Systems*, pages 402–409, 1991.

[8] M. Willebeek-LeMair and A. P. Reeves. Distributed dynamic load balancing. In *Proceedings of Conference on Hypercube Concurrent Computers and Applications*, pages 609–612, 1989.

[9] M. Willebeek-LeMair and A. P. Reeves. Local vs. global strategies for dynamic load balancing. In *Proceedings of International Conference on Parallel Processing*, pages 569–570(I), 1990.

[10] C.-Z. Xu and F.C.M. Lau. Analysis of the generalized dimension exchange method for dynamic load balancing. *Journal of Parallel and Distributed Computing*, 16(4):385–393, December 1992.

[11] C.-Z. Xu and F.C.M. Lau. Iterative dynamic load balancing in multicomputers. Technical Report TR-92-09, Dept. of Computer Science, The University of Hong Kong, 1992.