

Fault-Tolerant Meshes with Minimal Numbers of Spares

Jehoshua Bruck Robert Cypher Ching-Tien Ho
IBM Almaden Research Center
San Jose, CA 95120

Abstract

This paper presents several techniques for adding fault-tolerance to distributed memory parallel computers. More formally, given a target graph with n nodes, we create a fault-tolerant graph with $n + k$ nodes such that given any set of k or fewer faulty nodes, the remaining graph is guaranteed to contain the target graph as a fault-free subgraph. As a result, any algorithm designed for the target graph will run with no slowdown in the presence of k or fewer node faults, regardless of their distribution. We present fault-tolerant graphs for target graphs which are 2-dimensional meshes, tori, eight-connected meshes and hexagonal meshes. In all cases our fault-tolerant graphs have smaller degree than any previously known graphs with the same properties.

1 Introduction

The advent of microprocessor technology and large scale integration at affordable costs have allowed the design and fabrication of parallel machines hosting a large number of processors. As the number of the components in an architecture becomes larger it is essential to consider the issue of computing in the presence of faults.

A large amount of research has been devoted to creating fault-tolerant parallel architectures. The techniques used in this research can be divided into two main classes. The first class consists of techniques which *do not* add redundancy to the desired architecture. Instead, these techniques attempt to mask the effects of faults by using the healthy part of the architecture to simulate the entire machine. The hope with this approach is to obtain the same functionality with a reasonable slowdown factor. This approach is taken, for example, by [1,6,11,13]. While this approach yields interesting theoretical results, even a constant factor slowdown in performance can be very significant in practice. Furthermore, this approach requires that

some healthy processors simulate several processors. As a result, each simulated processor can have only a fraction of the memory present in a healthy processor.

The second class consists of techniques which *do* add redundancy to the desired architecture. These techniques attempt to isolate the faults, usually by disabling certain links or disallowing certain switch settings, while maintaining the complete desired architecture. Examples of such techniques are given in [2,3,8,9,12,14,15,16,17,18,19,21,23,24]. Many of these techniques require either a fault-free switching mechanism [3,14,15,18,19,21,23] or a non-minimal number of spare processors [2,3,14,15,21,23]. In contrast, all of the results presented in this paper require neither a switching mechanism nor more spare processors than can be faulty. Furthermore, we assume a worst case distribution of faults, while many of the above approaches do not work in a worst case scenario.

Our approach is based on a graph model. In this model a distributed memory parallel computer is viewed as being a graph in which the nodes represent the processors and the edges represent the communication links. A target graph with n nodes is first selected. Then a fault-tolerant graph with $n + k$ nodes is defined with the property that given any set of k or fewer faulty nodes, the remaining graph is guaranteed to contain the target graph as a subgraph. This approach guarantees that any algorithm designed for the target graph will run *with no slowdown* in the presence of k or fewer node faults, regardless of their distribution. Note that in our approach the spare nodes are fully utilized. Hence, minimizing the cost in this model amounts to constructing a fault-tolerant graph with a small degree. Although our results are stated for node faults, it should be noted that they can also be used to tolerate edge faults by viewing a node incident with each faulty edge as being faulty.

This graph model of fault-tolerance has been used by several other researchers. Hayes [12] has used this model with target graphs of cycles, linear arrays and trees. The work by Wong and Wong [24] and Paoli, Wong and Wong [17] relates to cycles. The more re-

cent work by Dutt and Hayes uses trees [8] and hypercubes [9] as target graphs. Rosenberg [20] proves lower bounds on the VLSI area requirements of fault-tolerant graphs which tolerate a constant fraction of the processors being faulty and use a minimum number of spare processors. It should be noted that the constructions presented here tolerate k faults for any value of k , and do not assume that a constant fraction of the processors are faulty.

The main contribution of this paper is the creation of efficient fault-tolerant graphs for several important target graphs. Specifically, we give 4 different constructions for creating fault-tolerant 2-dimensional meshes, as well as constructions for creating fault-tolerant tori, eight-connected meshes and hexagonal meshes. In all cases our fault-tolerant graphs have smaller degree than any previously known graphs with the same properties. In particular, one of the constructions for fault-tolerant 2-dimensional meshes can tolerate k faults and has degree $2k+4$. We also present constructions in which the degree increases by only one per additional fault tolerated, and we present lower bounds for certain types of constructions.

The rest of this paper is organized as follows. Section 2 defines a class of graphs known as “circulant graphs” and reviews some of their properties. Section 3 presents constructions for fault-tolerant 2-dimensional meshes. Fault-tolerant graphs for tori, eight-connected meshes and hexagonal meshes are given in Section 4. Due to space limitations, some of the proofs which appear in the full version of the paper [5] have been omitted here. The following definitions will be used throughout this paper.

Definition: Let k be a nonnegative integer and let $G = (V, E)$ be a graph. We say that the graph $G' = (V', E')$ is (k, G) -tolerant if the subgraph of G' induced by any set of $|V'| - k$ nodes contains G as a subgraph. We note here that throughout this paper the number of spare nodes is minimal, namely $|V'| = |V| + k$.

Definition: Given two graphs G_1 and G_2 , a function ϕ which maps the vertices of G_1 to the vertices of G_2 is called an *embedding of G_1 into G_2* if for any pair of distinct nodes i and j in G_1 , $\phi(i) \neq \phi(j)$, and for any edge (i, j) in G_1 , $(\phi(i), \phi(j))$ is an edge in G_2 .

Definition: For any positive integer n , the set $\{0, 1, \dots, n-1\}$ will be denoted $[n]$.

2 Circulant Graphs

This section discusses a class of graphs known as “circulant graphs” [10].

Definition: Let n be a positive integer and let S be a set of integers in the range 1 through $n-1$. The n -node circulant graph with connection set S , denoted $C_{n,S}$, consists of n nodes. Each node in $C_{n,S}$ has a unique label in the range 0 through $n-1$. Each node i is connected to all nodes of the form $(i \pm s) \bmod n$ where $s \in S$.

Definition: Let n be a positive integer and let S be a set of integers in the range 1 through $n-1$. The closure of S by n , denoted $close(S, n)$, is the set

$$T = \{t \mid t \in S \text{ or } (n-t) \in S\}.$$

Note that the degree of $C_{n,S}$ is $|close(S, n)|$. Also, note that $|S| \leq |close(S, n)| \leq 2|S|$.

Definition: Let S be a set of integers and let k be a nonnegative integer. The expansion of S by k , denoted $expand(S, k)$, is the set T where

$$T = \bigcup_{s \in S} \{s, s+1, \dots, s+k\}.$$

Note that $|expand(S, k)| \leq (k+1)|S|$.

The following theorem is an immediate consequence of a result proven by Dutt and Hayes [9].

Theorem 2.1 *Let n be a positive integer, let S be a set of integers in the range 1 through $n-1$, let k be a nonnegative integer, and let $T = expand(S, k)$. The circulant graph $C_{n+k,T}$ is $(k, C_{n,S})$ -tolerant.*

The idea behind Theorem 2.1 is that given any set of k faulty nodes in $C_{n+k,T}$, we can embed the target graph $C_{n,S}$ into the healthy nodes of the fault-tolerant graph $C_{n+k,T}$ by mapping each node i in the target graph to the i -th healthy node in the fault-tolerant graph. It is clear that any pair of nodes that are x apart in the target graph are mapped to nodes in the fault-tolerant graph that are at least x apart and at most $x+k$ apart (because there are between 0 and k faulty nodes between them). Consider any edge which connects nodes that are x apart in the target graph, where $x \in S$. This edge will be mapped to nodes which are x' apart in the fault-tolerant graph, where $x' \in T$, so it will be mapped to an edge in the fault-tolerant graph. All of our fault-tolerant constructions will make use of Theorem 2.1.

3 2-Dimensional Meshes

In this section we will consider the creation of graphs which can sustain node faults and still be guaranteed to contain a nonfaulty 2-dimensional mesh.

Definition: Let r and c be positive integers. The $r \times c$ 2-dimensional mesh, denoted $M_{r,c}$, consists of rc nodes. Each node in $M_{r,c}$ has a unique label of the form (i, j) where $0 \leq i < r$ and $0 \leq j < c$. Each node (i, j) is connected to the 4 other nodes $(i \pm 1, j)$ and $(i, j \pm 1)$, provided they exist.

We will use Theorem 2.1 to obtain four different constructions for fault-tolerant 2-dimensional meshes. Each construction first defines a circulant graph which is a supergraph of the desired 2-dimensional mesh. Then Theorem 2.1 is used to add fault-tolerance to the supergraph. It is interesting to note that circulant graphs which contain 2-dimensional meshes as subgraphs have been studied in a context unrelated to fault-tolerance [4].

Throughout this section, let r and c be positive integers and let k be a nonnegative integer. Additional constraints on these parameters will be added as needed.

3.1 Mesh Construction 1

The first fault-tolerant mesh construction is based on the fact that, when the nodes in $M_{r,c}$ are labeled in row-major order, the labels of adjacent nodes differ by either 1 or c (see Figure 1a).

Lemma 3.1 Let $S = \{1, c\}$. The mesh $M_{r,c}$ is a subgraph of the circulant graph $C_{rc,S}$.

Proof: Let $\phi(i, j) = ic + j$. It is straightforward to verify that ϕ defines an embedding of $M_{r,c}$ into $C_{rc,S}$. \square

Theorem 3.2 Let $S = \{1, c\}$ and let $T = \text{expand}(S, k)$. The circulant graph $C_{rc+k,T}$ is $(k, M_{r,c})$ -tolerant and has degree at most $4k + 4$.

Proof: From Theorem 2.1, the graph $C_{rc+k,T}$ is $(k, C_{rc,S})$ -tolerant. From Lemma 3.1, the graph $M_{r,c}$ is a subgraph of $C_{rc,S}$. As a result, the graph $C_{rc+k,T}$ is $(k, M_{r,c})$ -tolerant. Because $|S| \leq 2$ and $T = \text{expand}(S, k)$, $|T| \leq 2k + 2$ and the degree of $C_{rc+k,T}$ is at most $4k + 4$. \square

3.2 Mesh Construction 2

While Construction 1 is a very natural application of Theorem 2.1, Theorem 2.1 can also be used to obtain more efficient constructions. In this subsection we will present a construction for obtaining a graph which tolerates k faults, but which has degree only $2k + 4$.

This construction is based on an ordering of the nodes in the mesh which we call the *antidiagonal-major order* (see Figure 1b). The advantage of antidiagonal-major order is that it leads to a circulant graph that has a connection set consisting of 2 consecutive integers. As a result, fault-tolerance can be added to the circulant graph in an efficient manner.

Lemma 3.3 Let $S = \{c, c + 1\}$. The mesh $M_{r,c}$ is a subgraph of the circulant graph $C_{rc,S}$.

Proof: Let $\phi(i, j) = ((i+j) \bmod r)c + j$. It is straightforward to verify that ϕ defines an embedding of $M_{r,c}$ into $C_{rc,S}$. \square

Theorem 3.4 Let $S = \{c, c + 1\}$ and let $T = \text{expand}(S, k)$. The circulant graph $C_{rc+k,T}$ is $(k, M_{r,c})$ -tolerant and has degree at most $2k + 4$.

Proof: From Theorem 2.1, the graph $C_{rc+k,T}$ is $(k, C_{rc,S})$ -tolerant. From Lemma 3.3, the graph $M_{r,c}$ is a subgraph of $C_{rc,S}$. As a result, the graph $C_{rc+k,T}$ is $(k, M_{r,c})$ -tolerant. Because $T = \{c, c + 1, \dots, c + k + 1\}$, $|T| \leq k + 2$ and the degree of $C_{rc+k,T}$ is at most $2k + 4$. \square

3.3 Mesh Construction 3

The fault-tolerant meshes produced by Construction 2 require 2 additional edges per node for each additional fault which is tolerated. In this subsection we will give a construction that requires only 1 additional edge per node for each additional fault which is tolerated. However, this reduced rate of growth in the degree requires a larger initial degree.

The construction is based on an ordering of the nodes in the mesh which we call the *interleaved antidiagonal-major order* (see Figure 1c). The interleaved antidiagonal-major order assigns the numbers 0 through $rc - 1$ to the nodes in $M_{r,c}$. Node $(0, 0)$ (the upper left corner) is assigned the value 0, and successive values are assigned to the nodes in every other antidiagonal. Then node $(1, 0)$ (the node immediately below the upper left corner) is assigned the value $\lceil rc/2 \rceil$, and successive values are assigned to the nodes in the remaining antidiagonals. The advantage of interleaved antidiagonal-major order is that it leads to a circulant graph with rc nodes that has a connection set which is clustered about the value $rc/2$ (see Figure 1d). The proofs of the following lemma and theorem have been omitted.

Lemma 3.5 Let r and c be integers greater than 2, let $a = \lceil rc/2 \rceil - \lfloor r/2 \rfloor$, let $b = \lceil rc/2 \rceil + \lfloor r/2 \rfloor$, and let S be the set of integers in the range a through b . The mesh $M_{r,c}$ is a subgraph of the circulant graph $C_{rc,S}$.

Theorem 3.6 Let r and c be integers greater than 2, let $a = \lceil rc/2 \rceil - \lfloor r/2 \rfloor$, let $b = \lceil rc/2 \rceil + \lfloor r/2 \rfloor$, let S be the set of integers in the range a through b , let k be a nonnegative integer, and let $T = \text{expand}(S, k)$. The circulant graph $C_{rc+k,T}$ is $(k, M_{r,c})$ -tolerant and has degree at most $k + r + 1$ when r is odd and c is even, and at most $k + r$ otherwise.

Theorem 3.6 is based on Lemma 3.5 which showed that the mesh $M_{r,c}$ is a subgraph of a circulant graph with rc nodes and a connection set which has values that are all near $rc/2$. Specifically, when r is odd and c is even all of the values in the connection set are within $(r+1)/2$ of $rc/2$, and in all other cases all of the values in the connection set are within $r/2$ of $rc/2$. If Lemma 3.5 could be improved by finding a circulant graph with a connection set that is more tightly clustered around $rc/2$, the degree of the construction in Theorem 3.6 could be reduced. However, as we will see in Theorem 3.7, no such improvement in Lemma 3.5 is possible. The proof of Theorem 3.7 is omitted.

Theorem 3.7 Let r and c be integers where $4 \leq r \leq c$ and let $C_{rc,S}$ be a circulant graph which contains the mesh $M_{r,c}$ as a subgraph. There exists an $s \in S$ such that $|s - (rc/2)| \geq (r+1)/2$ if r is odd and c is even, and such that $|s - (rc/2)| \geq r/2$ otherwise.

3.4 Mesh Construction 4

In this subsection we will present constructions of $(k, M_{r,c})$ -tolerant graphs which combine the advantages of Constructions 2 and 3. More precisely, the degree of the constructions given here increases at the rate of 2 per fault up to some number of faults, at which point it increases at the rate of 1 per fault. The cut-off point at which the rate of growth in the degree slows depends on a value called the *gap*, which will be defined later.

Lemma 3.8 The following properties hold. (i) If r is odd then r and $(r-1)/2$ are relatively prime. (ii) If $r \bmod 4 = 0$ then r and $(r/2) - 1$ are relatively prime. (iii) If $r \bmod 4 = 2$ then r and $(r/2) - 2$ are relatively prime.

Proof: We first prove (i). Let $r = 2x + 1$ where x is an integer. Thus $\gcd(r, (r-1)/2) = \gcd(2x+1, x) = \gcd(1, x) = 1$, so r and $(r-1)/2$ are relatively

prime. To prove (ii), let $r = 4x$ where x is an integer. Thus $\gcd(r, (r/2) - 1) = \gcd(4x, 2x - 1) = \gcd(2, 2x - 1) = 1$. To prove (iii), let $r = 4x + 2$ where x is an integer. Thus, $\gcd(r, (r/2) - 2) = \gcd(4x + 2, 2x - 1) = \gcd(4, 2x - 1) = 1$. \square

Figures 2a to 2c show examples of the mapping ϕ for Lemmas 3.9 through 3.11, respectively.

Lemma 3.9 Let r be odd and let $S = \{(r-1)c/2, (r-1)c/2 + 1\}$. The mesh $M_{r,c}$ is a subgraph of the circulant graph $C_{rc,S}$.

Proof: Let $f(i) = (i(r-1)/2) \bmod r$ and let $\phi(i, j) = f(i+j)c + j$. We will first show that ϕ maps distinct nodes to distinct values. From Lemma 3.8, r and $(r-1)/2$ are relatively prime, so for any integers x and x' , if $f(x) = f(x')$ it follows that $x \bmod r = x' \bmod r$. Let (i, j) and (i', j') be any nodes in $M_{r,c}$. Note that $f(i+j)c$ and $f(i'+j')c$ are multiples of c , and that $j \in [c]$ and $j' \in [c]$. Therefore, if $\phi(i, j) = \phi(i', j')$, it follows that $j = j'$, which implies that $f(i+j) = f(i'+j')$, so $i \bmod r = i' \bmod r$ and $i = i'$.

We will now show that ϕ maps edges in $M_{r,c}$ to edges in $C_{rc,S}$. We will show this by proving that for any integers i and j , (i) $|\phi(i+1, j) - \phi(i, j)| \in \text{close}(S, rc)$, and (ii) $|\phi(i, j+1) - \phi(i, j)| \in \text{close}(S, rc)$. Let $s = (r-1)c/2$ and let $s' = (r+1)c/2$, and note that $\text{close}(S, rc) = \{s, s+1, s'-1, s'\}$. Also, note that for any integer x , $f(x+1) - f(x)$ equals either $(r-1)/2$ or $-(r+1)/2$. Therefore, $|\phi(i+1, j) - \phi(i, j)| = |f(i+1+j) - f(i+j)|c$ which equals either s or s' , and property (i) holds. Let $\phi(i, j) = y$. Clearly, $\phi(i, j+1) = \phi(i+1, j) + 1$, which equals either $y + s + 1$ or $y - s' + 1$, so property (ii) holds as well. As a result, ϕ is an embedding of $M_{r,c}$ into $C_{rc,S}$. \square

Lemma 3.10 Let $r \bmod 4 = 0$ and let $S = \{(r/2 - 1)c, (r/2 - 1)c + 1\}$. The mesh $M_{r,c}$ is a subgraph of the circulant graph $C_{rc,S}$.

Proof: Let $f(i) = ((r/2 - 1)i) \bmod r$ and let $\phi(i, j) = f(i+j)c + j$. The proof is analogous to that of Lemma 3.9 and will not be repeated here. \square

Lemma 3.11 Let $r \bmod 4 = 2$ and let $S = \{(r/2 - 2)c, (r/2 - 2)c + 1\}$. The mesh $M_{r,c}$ is a subgraph of the circulant graph $C_{rc,S}$.

Proof: Let $f(i) = ((r/2 - 2)i) \bmod r$ and let $\phi(i, j) = f(i+j)c + j$. The proof is analogous to that of Lemma 3.9 and will not be repeated here. \square

Definition: Let n, s and x be integers where $0 \leq s < \lfloor n/2 \rfloor$ and $0 \leq x \leq s$, and let $S = \{s-x, s-x+1, \dots, s\}$. Then $\text{gap}(S, n) = n-1-2s$.

Intuitively, $\text{gap}(S, n)$ is the length of the “gap” between the two consecutive groups of offsets in $\text{close}(S, n)$. For instance, if $S = \{5, 6\}$, $\text{close}(S, 16) = \{5, 6, 10, 11\}$ and $\text{gap}(S, 16) = 3$.

Lemma 3.12 *Let n, s and x be integers where $0 \leq s < \lfloor n/2 \rfloor$ and $0 \leq x \leq s$, let $S = \{s-x, s-x+1, \dots, s\}$ and let $T = \text{expand}(S, k)$. The circulant graph $C_{n+k, T}$ has degree at most*

$$d = \begin{cases} 2k + 2|S|, & \text{if } k \leq \text{gap}(S, n), \\ k + 2|S| + \text{gap}(S, n), & \text{if } k > \text{gap}(S, n). \end{cases}$$

Proof: The fact that $d \leq 2k + 2|S|$ follows immediately from the definitions of expansion and closure. Now consider the case where $k > \text{gap}(S, n)$. Note that $T = \{s-x, s-x+1, \dots, s+k\}$ so we will compare the values of $s+k$ and $n+k-(s+k) = n-s$ to see if there is a gap between the two groups of offsets in $\text{close}(T, n+k)$. Note that $(n-s)-(s+k) = n-2s+k$, which is less than $n-2s-(n-1-2s) = 1$ because $k > \text{gap}(S, n) = n-1-2s$. Therefore, there is no gap between the two groups of offsets in $\text{close}(T, n+k)$, so $\text{close}(T, n+k) = \{s-x, s-x+1, \dots, n+k-(s-x)\}$ and $|\text{close}(T, n+k)| = n+k-(s-x)-(s-x-1) = n+k-2s+2x+1$. Since $k + 2|S| + \text{gap}(S, n) = k + 2(x+1) + n-1-2s$, we have $|\text{close}(T, n+k)| = k + 2|S| + \text{gap}(S, n)$. \square

Theorem 3.13 *Let*

$$S = \begin{cases} \left\{ \left\lfloor \frac{(r-1)c}{2} \right\rfloor, \left\lfloor \frac{(r-1)c}{2} \right\rfloor + 1 \right\}, & \text{if } r \text{ is odd,} \\ \left\{ \left(\frac{r}{2} - 1 \right) c, \left(\frac{r}{2} - 1 \right) c + 1 \right\}, & \text{if } r \bmod 4 = 0, \\ \left\{ \left(\frac{r}{2} - 2 \right) c, \left(\frac{r}{2} - 2 \right) c + 1 \right\}, & \text{if } r \bmod 4 = 2, \end{cases}$$

and let $T = \text{expand}(S, k)$. Then the circulant graph $C_{rc+k, T}$ is $(k, M_{r,c})$ -tolerant and has degree at most

$$d(k, r, c) = \begin{cases} 2k + 4, & r \text{ is odd and } k \leq c - 3, \\ c + k + 1, & r \text{ is odd and } k > c - 3, \\ 2k + 4, & r \bmod 4 = 0 \text{ and } k \leq 2c - 3, \\ 2c + k + 1, & r \bmod 4 = 0 \text{ and } k > 2c - 3, \\ 2k + 4, & r \bmod 4 = 2 \text{ and } k \leq 4c - 3, \\ 4c + k + 1, & r \bmod 4 = 2 \text{ and } k > 4c - 3. \end{cases}$$

Proof: From Theorem 2.1, the graph $C_{rc+k, T}$ is $(k, C_{r,c,S})$ -tolerant. From Lemmas 3.9, 3.10 and 3.11, the graph $M_{r,c}$ is a subgraph of $C_{r,c,S}$. Thus, the

graph $C_{rc+k, T}$ is $(k, M_{r,c})$ -tolerant. The degree follows directly from Lemma 3.12 and the fact that $\text{gap}(S, rc) = c-3$ if r is odd, $2c-3$ if $r \bmod 4 = 0$, and $4c-3$ if $r \bmod 4 = 2$. \square

Note that if the numbers $0, 1, \dots, rc-1$ are assigned to meshes in the diagonal manner instead of the antidiagonal manner, then the new $\text{gap}(S, rc)$ is greater than the original $\text{gap}(S, rc)$ by two. It should be noted that Theorem 3.13 also applies to meshes of the same shape with additional vertical wraparound connections. For instance, Lemma 3.9 actually shows that the mesh $M_{r,c}$ with vertical wraparound connections is a subgraph of the circulant graph $C_{rc,S}$ for the case that r is odd.

4 Other Graphs

In this section we will present fault-tolerant graphs for target graphs which are tori, 8-connected meshes and hexagonal meshes.

4.1 Torus Construction

An r by c torus, denoted $M_{r,c}^T$, is an r by c mesh $M_{r,c}$ to which “wraparound” connections have been added which connect the first and last nodes in each row and the top and bottom nodes in each column. In this subsection, we will show that given any $r \times c$ torus $M_{r,c}^T$, we can construct a $(k, M_{r,c}^T)$ -tolerant graph with $rc+k$ nodes and degree at most

$$\begin{cases} 2k + 4, & \text{if } r \text{ and } c \text{ are relatively prime,} \\ 2k + 6, & \text{if at least one of } r \text{ and } c \text{ is odd,} \\ 4k + 6, & \text{if both } r \text{ and } c \text{ are even.} \end{cases}$$

4.1.1 Case 1: r and c are relatively prime

The construction of a fault-tolerant torus $M_{r,c}^T$ for which r and c are relatively prime is based on the wraparound diagonal-major order (see Figure 3a).

Lemma 4.1 *Let r and c be relatively prime, let x be the integer satisfying $xc \bmod r = 1$ and $1 \leq x < r$, and let $S = \{xc-1, xc\}$. The torus $M_{r,c}^T$ is a subgraph of the circulant graph $C_{rc,S}$.*

Proof: Let $f(i) = xi \bmod r$ and let $\phi(i, j) = f(i-j)c + j$. The proof is analogous to that of Lemma 3.9 and will not be repeated here. \square

Theorem 4.2 Let r and c be relatively prime, let x be the integer satisfying $xc \bmod r = 1$ and $1 \leq x < r$, let $S = \{xc - 1, xc\}$, and let $T = \text{expand}(S, k)$. The circulant graph $C_{rc+k, T}$ is $(k, M_{r,c}^T)$ -tolerant and has degree at most $2k + 4$.

Proof: Follows from Theorem 2.1 and Lemma 4.3. \square

Note that $\text{gap}(S, rc) = (r - 2x)c - 1$ if $x < r/2$, and $(2x - r)c - 3$ otherwise. It is possible to reduce the *gap* by two for the former case by doing an antidiagonal traversal instead of a diagonal traversal. As before, for $k \geq \text{gap}(S, rc)$ the additional edge per node is one for each additional fault (i.e., for each increment in k).

4.1.2 Case 2: r is odd and c is even

The construction of fault-tolerant torus $M_{r,c}^T$ for which r is odd and c is even is based on the *interleaved zigzag-major order* (see Figure 3b).

Lemma 4.3 Let r be odd, let c be even, and let $S = \{(r - 1)c/2 - 1, (r - 1)c/2, (r - 1)c/2 + 1\}$. The torus $M_{r,c}^T$ is a subgraph of the circulant graph $C_{rc, S}$.

Proof: Let $\phi(i, j) = \{([i - (j \bmod 2)] \frac{r+1}{2}) \bmod r\}c + j$. The proof is analogous to that of Lemma 3.9 and will not be repeated here. \square

Theorem 4.4 Let r be odd, let c be even, let $S = \{(r - 1)c/2 - 1, (r - 1)c/2, (r - 1)c/2 + 1\}$, and let $T = \text{expand}(S, k)$. The circulant graph $C_{rc+k, T}$ is $(k, M_{r,c}^T)$ -tolerant and has degree at most

$$\begin{cases} 2k + 6, & \text{if } k \leq c - 3, \\ k + c + 3, & \text{otherwise.} \end{cases}$$

Proof: The fact that the circulant graph $C_{rc+k, T}$ is $(k, M_{r,c}^T)$ -tolerant follows from Theorem 2.1 and Lemma 4.3. The degree follows from the fact that $\text{gap}(S, rc) = c - 3$, $|S| = 3$, and Lemma 3.12. \square

Note that if the same zigzag style ordering is applied without interleaving between successive zigzag rows, the degree is still $2k + 6$ but the gap is much larger.

4.1.3 Case 3: both r and c are odd

Assume without loss of generality that $r \leq c$. The construction is based on a hybrid method combining the wraparound diagonal-major order (of case 1) with the zigzag-major order (similar to case 2). See Figure 3c for an example.

Lemma 4.5 Let r and c be odd integers with $r \leq c$ and let $S = \{2c - 1, 2c, 2c + 1\}$. The torus $M_{r,c}^T$ is a subgraph of the circulant graph $C_{rc, S}$.

Proof: The zigzag ordering is applied to the $(r + 1)/2$ leftmost columns if $r \bmod 4 = 3$, and to the $(r - 1)/2$ leftmost columns if $r \bmod 4 = 1$. The wraparound diagonal ordering is applied to the rest of columns on the right.

Formally, let $f_1(i) = i(r - 2) \bmod r$, let $f_2(i) = 2i \bmod r$ and let $\phi(i, j)$ equal $f_1(i - (j \bmod 2))c + j$ if $r \bmod 4 = 3$ and $j < \frac{r+1}{2}$, $f_1(i - j + \frac{r+1}{2})c + j$ if $r \bmod 4 = 3$ and $j \geq \frac{r+1}{2}$, $f_2(i - (j \bmod 2))c + j$ if $r \bmod 4 = 1$ and $j < \frac{r-1}{2}$, or $f_2(i - j + \frac{r-1}{2})c + j$ if $r \bmod 4 = 1$ and $j \geq \frac{r-1}{2}$. It is straightforward to verify that ϕ defines an embedding of $M_{r,c}^T$ into $C_{rc, S}$. \square

Theorem 4.6 Let r and c be odd and $r \leq c$, let $S = \{2c - 1, 2c, 2c + 1\}$, and let $T = \text{expand}(S, k)$. The circulant graph $C_{rc+k, T}$ is $(k, M_{r,c}^T)$ -tolerant and has degree at most $2k + 6$.

Proof: The fault-tolerance of $C_{rc+k, T}$ follows from Theorem 2.1 and Lemma 4.5. Since $T = \{2c - 1, 2c, \dots, 2c + k + 1\}$, the degree of $C_{rc+k, T}$ is $|\text{close}(T, rc + k)| = 2k + 6$. \square

4.1.4 Case 4: Both r and c are even

In this case, we simply use *row-major order* as used in Section 3.1 for the first fault-tolerant mesh construction. The proof of the following lemma is analogous to that of Lemma 3.1 and will not be included.

Lemma 4.7 Let $S = \{1, c - 1, c\}$. The torus $M_{r,c}^T$ is a subgraph of the circulant graph $C_{rc, S}$.

Theorem 4.8 Let $S = \{1, c - 1, c\}$ and let $T = \text{expand}(S, k)$. The circulant graph $C_{rc+k, T}$ is $(k, M_{r,c}^T)$ -tolerant and has degree at most $4k + 6$.

Proof: Follows from Theorem 2.1 and Lemma 4.7. \square

Note that Theorem 4.8 does not require that r and c have any special properties.

4.2 Eight-Connected Meshes

An r by c eight-connected mesh, denoted $M_{r,c}^8$, is an r by c mesh $M_{r,c}$ to which connections between nodes which are diagonal or antidiagonal neighbors have been added. We will use *row-major order* to construct its fault-tolerant graph. The proofs are analogous to those of the previous section and are omitted.

Lemma 4.9 Let $S = \{1, c - 1, c, c + 1\}$. The 8-connected mesh $M_{r,c}^8$ is a subgraph of the circulant graph $C_{r,c,S}$.

Theorem 4.10 Let $S = \{1, c - 1, c, c + 1\}$ and let $T = \text{expand}(S, k)$. The circulant graph $C_{rc+k,T}$ is $(k, M_{r,c}^8)$ -tolerant and has degree at most $4k + 8$. Furthermore, the same circulant graph $C_{rc+k,T}$ is also (k, G) -tolerant where G is any one of the following graphs: (i) an $r \times c$ torus, (ii) an $r \times c$ horizontally twisted torus, (iii) an $r \times c$ vertically twisted torus, and (iv) an $r \times c$ double twisted torus.

4.3 Hexagonal Meshes

A hexagonal mesh (H-mesh) of order c is a 6-connected mesh with hexagonal boundary. Each node is connected to 2 horizontal neighbors, 2 diagonal neighbors and 2 antidiagonal neighbors, if they exist. The order is the length of one coordinate. Chen et al. [7] defined the wraparound connection of H-meshes, termed C-type wrapping, such that they become node symmetric graphs. In the C-type wrapping, the rightmost node at row i , where $0 \leq i < 2c - 1$, is connected to the leftmost node at row $(i + c) \bmod (2c - 1)$. The same wrapping scheme is applied to two other coordinates after rotating the H-mesh. Chen et al. [7] also showed the isomorphism between the C-type wrapping H-meshes and a family of circulant graphs (as described by the lemma follows), which is useful in constructing the fault-tolerant graphs for H-meshes. In the following, we denote M_c^H the C-type H-mesh of order c and $N(c) = 3c^2 - 3c + 1$ the number of nodes in M_c^H .

Lemma 4.11 [7] Let $S = \{1, 3c - 2, 3c - 1\}$ and let $N(c) = 3c^2 - 3c + 1$. The wraparound hexagonal mesh M_c^H is topologically equivalent to the circulant graph $C_{N(c),S}$.

Theorem 4.12 Let $S = \{1, 3c - 2, 3c - 1\}$, let $N(c) = 3c^2 - 3c + 1$, and let $T = \text{expand}(S, k)$. The circulant graph $C_{N(c)+k,T}$ is (k, M_c^H) -tolerant and has degree at most $4k + 6$.

Note that an H-mesh of order c without wraparound is also a subgraph of a $(2c - 1) \times (2c - 1)$ 8-connected mesh. However, the latter has $c^2 - c$ more nodes than the former.

References

[1] F. Annexstein, *Fault Tolerance in Hypercube-Derivative Networks*, Proc. ACM Symp. on Par-

allel Algorithms and Architectures, pp. 179-188, 1989.

[2] V. Balasubramanian and P. Banerjee, *A Fault Tolerant Massively Parallel Processing Architecture*, J. of Parallel and Distributed Computing, vol. 4, pp. 363-383, 1987.

[3] K. E. Batcher, *Design of a Massively Parallel Processor*, IEEE Trans. on Computers, vol. C-29, no. 9, pp. 836-840, 1980.

[4] R. Beivide, E. Herrada, J. L. Balcázar and J. Labarta, *Optimized Mesh-Connected Networks for SIMD and MIMD Architectures*, Proc. Intl. Symp. on Computer Architecture, pp. 163-170, 1987.

[5] J. Bruck, R. Cypher and C.T. Ho, *Fault-Tolerant Meshes with Minimal Numbers of Spares*, Tech. Rep. RJ8137, IBM Almaden Res. Ctr., 1991.

[6] J. Bruck, R. Cypher and D. Soroker, *Running Algorithms Efficiently on Faulty Hypercubes*, Proc. ACM Symp. on Parallel Algorithms and Architectures, pp. 37-44, 1990.

[7] M.-S. Chen, K.G. Shin and D.D. Kandlur, *Addressing, Routing, and Broadcasting in Hexagonal Mesh Multiprocessors*, IEEE Trans. on Computers, vol. C-39, no. 1, pp. 10-18, 1990.

[8] S. Dutt and J. P. Hayes, *On Designing and Reconfiguring k-Fault-Tolerant Tree Architectures*, IEEE Trans. on Computers, vol. C-39, no. 4, pp. 490-503, 1990.

[9] S. Dutt and J. P. Hayes, *An Automorphic Approach to the Design of Fault-Tolerant Multiprocessors*, Proc. Intl. Symp. on Fault-Tolerant Computing, pp. 496-503, 1989.

[10] B. Elspas and J. Turner, *Graphs with Circulant Adjacency Matrices*, J. of Combinatorial Theory, no. 9, pp. 297-307, 1970.

[11] J. Hastad, F. T. Leighton and M. Newman, *Fast Computations using Faulty Hypercubes*, Proc. ACM Symp. on Theory of Computing, pp. 251-284, 1989.

[12] J. P. Hayes, *A Graph Model for Fault-Tolerant Computing Systems*, IEEE Trans. on Computers, vol. C-25, no. 9, pp. 875-884, 1976.

- [13] C. Kaklamanis, A. R. Karlin, F. T. Leighton, V. Milenkovic, P. Raghavan, S. Rao, C. Thomborson and A. Tsantilas, *Asymptotically Tight Bounds for Computing with Faulty Arrays of Processors*, Proc. IEEE Symp. on Foundations of Computer Science, pp. 285–296, 1990.
- [14] S. Y. Kung, S. N. Jean and C. W. Chang, *Fault-Tolerant Array Processors Using Single-Track Switches*, IEEE Trans. on Computers, vol. C-38, no. 4, pp. 501–514, 1989.
- [15] S.-Y. Kuo and W. K. Fuchs, *Efficient Spare Allocation for Reconfigurable Arrays*, IEEE Design and Test, pp. 24–31, Feb. 1987.
- [16] F. T. Leighton and C. E. Leiserson, *Wafer Scale Integration of Systolic Arrays*, IEEE Trans. on Computers, vol. C-34, no. 5, pp. 448–461, 1985.
- [17] M. Paoli, W. W. Wong and C. K. Wong, *Minimum k -Hamiltonian Graphs, II*, J. of Graph Theory, vol. 10, pp. 79–95, 1986.
- [18] A. L. Rosenberg, *The Diogenes Approach to Testable Fault-Tolerant VLSI Processor Arrays*, IEEE Trans. on Computers, vol. C-32, no. 10, pp. 902–910, 1983.
- [19] A. L. Rosenberg, *Routing with Permuters: Toward Reconfigurable and Fault-Tolerant Networks*, Tech. Rep. CS-1981-13, Duke U. CS Dept., 1981.
- [20] A. L. Rosenberg, *On Designing Fault-Tolerant Arrays of Processors*, Tech. Rep. CS-1982-14, Duke U. CS Dept., 1982.
- [21] V. P. Roychowdhury, J. Bruck and T. Kailath, *Efficient Algorithms for Reconfiguration in VLSI/WSI Arrays*, IEEE Trans. on Computers, vol. C-39, no. 4, pp. 480–489, 1990.
- [22] M. R. Samatham and D. K. Pradhan, *The de Bruijn Multiprocessor Network: A Versatile Network for Parallel Computation*, IEEE Trans. on Computers, vol. 38, no. 4, pp. 567–581, 1989.
- [23] M. Sami and R. Stefanelli, *Reconfigurable Architectures for VLSI Processing Arrays*, Proc. of IEEE, vol. 74, no. 5, 1986.
- [24] W. W. Wong and C. K. Wong, *Minimum k -Hamiltonian Graphs*, J. of Graph Theory, vol. 8, pp. 155–165, 1984.

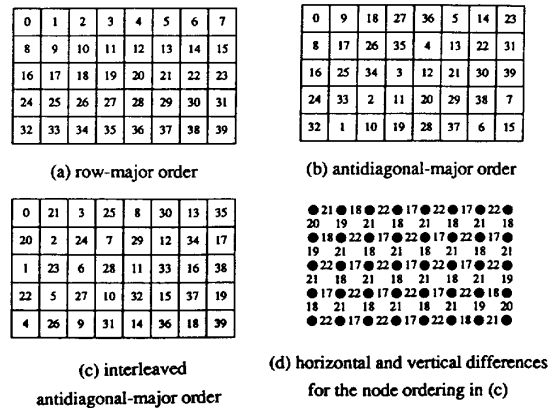


Figure 1: Mesh node orderings.

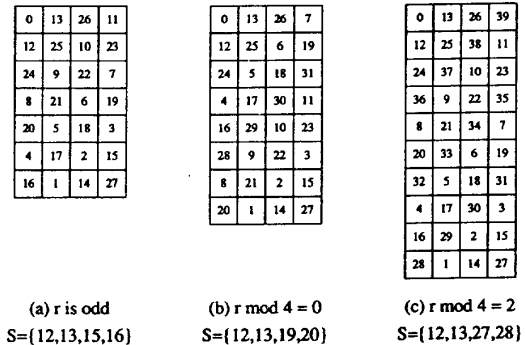


Figure 2: Mesh node orderings.

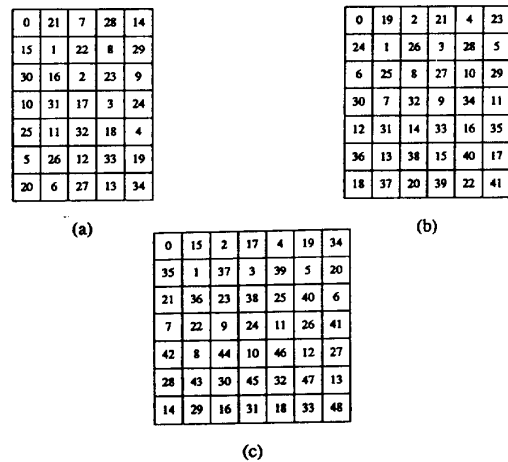


Figure 3: Torus node orderings.