Fault-Tolerant Meshes with Small Degree*

Jehoshua Bruck[†] Robert Cypher[‡] Ching-Tien Ho[§]

Abstract

This paper presents constructions for fault-tolerant two-dimensional mesh architectures. The constructions are designed to tolerate k faults while maintaining a healthy n by n mesh as a subgraph. They utilize several novel techniques for obtaining trade-offs between the number of spare nodes and the degree of the fault-tolerant network.

We consider both worst-case and random fault distributions. In terms of worst-case faults, we give a construction that has constant degree and $O(k^3)$ spare nodes. This is the first construction known in which the degree is constant and the number of spare nodes is independent of n. In terms of random faults, we present several new degree-6 and degree-8 constructions and show (both analytically and through simulations) that they can tolerate large numbers of randomly placed faults.

*A preliminary version of this paper appeared in *Proceedings of the Fifth Annual ACM Symposium on Parallel Algorithms and Architectures*, 1993.

[†]California Institute of Technology, Mail Code 116-81, Pasadena, CA 91125, bruck@systems.caltech.edu. This research was performed while the author was at the IBM Almaden Research Center.

[‡]Dept. of Computer Science, Johns Hopkins University, 3400 N. Charles Street, Baltimore, MD 21218, cypher@cs.jhu.edu. This research was performed while the author was at the IBM Almaden Research Center.

[§]IBM Almaden Research Center, 650 Harry Road, San Jose, CA 95120, ho@almaden.ibm.com.

1 Introduction

As the number of processors in parallel machines increases, physical limitations and cost considerations will tend to favor interconnection networks with constant degree and short wires, such as mesh networks [6]. In fact, the two-dimensional mesh is already one of the most important interconnection networks for parallel computers. Examples of existing two-dimensional mesh computers include the MPP (from Goodyear Aerospace), VICTOR (from IBM), and DELTA and Paragon (from Intel).

Another significant issue in the design of massively parallel computers is fault-tolerance. In order to create parallel computers with very large numbers of complex processors, it will become necessary to utilize these machines even when several components have failed. In particular, the ability to tolerate even a small number of faults may allow the machine to continue operation between the occurrence of the first fault and the repair of the 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 [2, 11, 17, 19, 23]. These techniques do not pay any costs for adding fault-tolerance, but they can experience a significant degradation in performance. 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 [1, 3, 4, 7, 8, 9, 10, 13, 14, 15, 18, 20, 22, 24, 25, 26, 28]. The goal with these techniques is to maintain the full performance of the desired architecture while minimizing the cost of the redundant components.

One of the most powerful techniques for adding redundancy is based on a graph-theoretic model of fault-tolerance [18]. In this model, the desired architecture is viewed as a graph (called the *target graph*) and a fault-tolerant graph is created such that after the removal of k faulty nodes, the target graph is still present as a subgraph. This technique yields fault-tolerant networks that can tolerate both node faults and edge faults (by viewing a node incident with the faulty edge as being faulty) and can implement algorithms designed for the target network without any slowdown (due to the simulation of multiple nodes by a single node or the routing of messages through switches or intermediate nodes). Unfortunately, the degree of the fault-tolerant network created with this model can be prohibitably large. In particular, all previously published techniques for creating fault-tolerant meshes have a degree that is linear in the number of faults being tolerated.

In this paper we create fault-tolerant meshes with small degree by trading-off the number of spare nodes with the degree of the fault-tolerant network. We consider both worst-case and random fault distributions. In terms of worst-case faults, we give a construction that tolerates k faults and has constant degree and $O(k^3)$ spares. This is the first construction known in which the degree is constant and the number of spares is independent of n. In terms of random faults, we present several new degree-6 and degree-8 constructions and show (both analytically and through simulations) that they can tolerate large numbers of randomly placed faults.

In addition, the construction for worst-case faults is shown to require only wires of length $O(k^3)$ in Thompson's VLSI model [27], while the constructions for random faults are shown to require only constant length wires. Thus the fault-tolerant constructions maintain much of the scalability of the mesh network. We remark that we use Thompson's VLSI model only because it provides a well-established means for quantifying the locality of an interconnection network; the use of this model does not imply that the constructions presented here are designed for the wafer-scale implementation of a parallel machine. In fact, most existing parallel machines have one, or at most a few, processors per chip. This fact motivates our concern with the degree of the fault-tolerant network (because of the limited number of pins available to connect one chip to another [12]).

The remainder of this paper is organized as follows. Definitions and several previously known results are given in Section 2. The results for worst-case fault distributions and random fault distributions are presented in Sections 3 and 4, respectively.

2 Preliminaries

Definitions: Let k be a nonnegative integer and let $T = (V, E)$ be a graph. The graph $F = (V', E')$ is a *k-fault-tolerant graph with respect to T*, denoted a k-FT T, if the subgraph of F induced by any set of $|V'| - k$ nodes contains T as a subgraph. The graph T will be called the *target graph.* The graph F will be said to contain $|V'| - |V|$ *spare nodes* (or *spares*).

Definition: The cycle with n nodes will be denoted C_n.

Definition: The two-dimensional mesh with $r \geq 2$ rows and $c \geq 2$ columns will be denoted $M_{r,c}$. Each node is $M^{r,c}$ has a unique label of the form (i, j) where $0 \leq ir$ and $0 \leq j < c$. Each node (i, j) is connected to all nodes of the form $(i \pm 1, j)$ and $(i, j \pm 1)$, provided they exist. The node (i, j) will be said to be in *row i* and *column j*.

Definitions: Let n be a positive integer and let S be a set of integers in the range 1 through $n - 1$. The graph $C(n, S)$, called the *n-node circulant graph with connection set S* [16, 14, 10], consists of n nodes numbered $0, 1, \ldots, n - 1$. Each node i is connected to all nodes of the form $(i \pm s)$ mod n where $s \in S$. The graph $D(n, S)$, called the *n-node diagonal graph with connection set S* [10], consists of n nodes numbered $0, 1, \ldots, n - 1$. Each node i is connected to all nodes of the form $i \pm s$ where $s \in S$, provided they exist. (The terms "circulant" and "diagonal" refer to the structure of the adjacency matrix.) The values in a connection set S will be referred to as "jumps" or "offsets" and an edge defined through an offset s will be referred to as *an s-offset edge.*

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

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

The following theorems give constructions for creating fault-tolerant circulant and diagonal graphs. The basic idea is to add offsets so that faulty nodes can be "jumped over". The construction for diagonal target graphs has lower degree because a cluster of faults can be avoided by placing the cluster in the position where the missing wraparound edges would jump over them.

Theorem 2.1 [14] *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 a k-FT $C(n,S)$.*

Theorem 2.2 [10] *Let n be a positive integer, let $y = \lceil n/3 \rceil$, let S be a set of integers in the range 1 through y, let k be a positive integer, and let $T = expand(S, \lfloor k/2 \rfloor)$. The circulant graph $C(n+k,T)$ is a k-FT $D(n,S)$.*

The following theorems relate meshes, circulant graphs and diagonal graphs. Combining these theorems with the two previous theorems yields constructions for fault-tolerant meshes. The first theorem follows immediately from the row-major labeling of the nodes in a mesh. The second theorem follows from a *diagonal-major order* of the nodes in a mesh; see Figure 1 for an example.

Theorem 2.3 *The mesh $M_{r,c}$ is a subgraph of $C(rc, \{1,c\})$ and of $D(rc,\{1,c\})$.*

Theorem 2.4 *The mesh $M_{r,c}$ is a subgraph of $C(rc, \{c-1,c\})$.*

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 $C(rc,\{c-1,c\})$. □

3 Worst Case Faults

In this section we present a graph $\tilde{M}$ that is a k-FT $M_{n,n}$ and has constant degree and $O(k^3)$ spares. Our construction is hierarchical. We first construct a graph M' that is a k-FT $M_{r,c}$ (for some suitably chosen parameters r and c) and has degree which is dependent on k. We then replace each node in M' with a supernode (a graph with certain properties) to obtain a graph $\tilde{M}$ with constant degree.

0	33	26	19	12	5	38	31
8	1	34	27	20	13	6	39
16	9	2	35	28	21	14	7
24	17	10	3	36	29	22	15
32	25	18	11	4	37	30	23

Figure 1: An example of a diagonal-major ordering of a mesh.

3.1 The Basic Construction

We first present a construction for a k-FT cycle with degree 4 and k^2 spare nodes. We will then use this construction to create the graph M' which is a k-FT $M_{r,c}$.

Theorem 3.1 *Let k and N be positive integers where $N \geq k^2 + k + 1$, and let the graph $C' = C(N + k^2, \{1, k+1\})$. The graph C' is a k-FT C_N.*

Proof: First consider the case where $(N + k^2) \bmod (k + 1) = 0$. For each i, $0 \leq i \leq k$, let X_i be the set consisting of all nodes $\{j | j \bmod (k + 1) = i\}$. Because there are only k faults and there are $k + 1$ disjoint sets X_i, at least one of them must be fault-free. Let X be such a fault-free X_i. Note that the nodes in X form a fault-free cycle C'' of length $(N + k^2)/(k + 1)$ using the $(k + 1)$-offset edges. Next, we augment C'' to get a healthy cycle of length at least N. For any two adjacent nodes a and b in C'', if all k of the nodes in C' between a and b are healthy, we traverse all k of these nodes by using the 1-offset edges. On the other hand, if there is a fault between a and b, we skip over all k of the nodes between them by traversing the $(k+1)$-offset edge connecting a and b. It is clear that we will traverse $(k + 1)$-offset edges at most k times, so the resulting augmented cycle will have at least N nodes. If it has more than N nodes, we can choose to traverse additional $(k + 1)$-offset edges, rather than 1-offset edges, until the cycle has length exactly N. An example of a 2-FT cycle is shown in Figure 2.

Now consider the case where $(N + k^2) \bmod (k + 1) = x \neq 0$. Let R be a region of $k + 1 + x$ consecutive healthy nodes in C'. Note that such a region must exist because $N + k^2 \geq 2k^2 + k + 1$, so there must be a region of $2k + 1$ or more consecutive healthy nodes between two faults. Without loss of generality, we will assume that R consists of the $k+1+x$ highest numbered nodes in C'. For each i, $0 \leq i \leq k$, create the cycle C_i'' as follows. First, start at node i and traverse the $(k + 1)$-offset edges until a node in R is reached. Then, traverse the 1-offset edges x times. Finally, traverse one additional $(k + 1)$-offset edge to return to i. Note that these $k + 1$ cycles only share nodes within R. Because all of the nodes in R are healthy, there must exist an i such that C_i'' is healthy. We can augment C_i'' as before to obtain a cycle of length N. $\Box$

Figure 2: A degree-4 2-fault-tolerant cycle with 4 spare nodes.

Theorem 3.2 *Let k, r and c be positive integers where $r, c \geq 2$ and $rc \geq k^2 + k + 1$, let $N = rc$, and let $M' = C(N + k^2, \{1, k+1\} \cup \{c + ik | 0 \leq i \leq k\})$. The graph M' is a k-FT $M_{r,c}$.*

Proof: Let $T = C(N, \{1, c\})$. We will prove that M' is a k-FT T. Applying Theorem 2.3 will complete the proof. First, it follows from Theorem 3.1 that in the presence of k faults, M' contains a cycle of N healthy nodes. Let C'' denote a cycle of healthy nodes constructed according to the proof of Theorem 3.1 and number the nodes in C'' from 0 through $N - 1$. We will now prove that any two nodes numbered a and b in C'', where $(a + c) \bmod N = b$, are connected in M'. Let a' and b' be the labels of a and b in M', and assume without loss of generality that $a' < b'$. We know that it is possible to traverse the cycle C'' from a to b by traversing 1-offset edges and at most k $(k+1)$-offset edges. Therefore, $b' - a' = c + jk$ for some integer j where $0 \leq j \leq k$, which implies that a and b are connected in M'. $\Box$

3.2 Hierarchical Constructions

In the previous subsection we described a construction of a k-FT cycle with k^2 spare nodes and degree 4 and a construction of a k-FT 2-dimensional mesh with k^2 spare nodes and degree $2k + 6$. In this subsection we will present techniques for reducing the degree of these FT graphs. The general idea is to replace each node in the original FT graph by a small graph (which we call a supernode). Then, for each edge (a, b) in the original graph, one or more nodes in the supernode corresponding to a is connected to one or more nodes in the supernode corresponding to b. This approach results in a FT graph with lower degree than the original graph, although it does increase the number of spare nodes that are required.

3.2.1 Hierarchical fault-tolerant cycles

We illustrate the concept of a supernode by creating a hierarchical FT cycle.

Theorem 3.3 *Let k and N be positive integers where $N \geq k^2 + k + 1$, and let $\hat{C}$ be the graph with $2N + 2k^2$ nodes, numbered 0 through $2N + 2k^2 - 1$, and with edges specified as follows: each odd numbered node i is connected to nodes $(i+1)$, $(i-1)$ and $i + 2k + 1$, and each even numbered node i is connected to nodes $(i+1)$, $(i-1)$ and $i - 2k - 1$, where all of the arithmetic is performed modulo $(2N + 2k^2)$. Then $\hat{C}$ is a k-FT C_{2N}.*

Proof: The graph $\hat{C}$ can be obtained from the graph C' of Theorem 3.1 by replacing each node with a supernode consisting of a pair of nodes connected to one another. The edges that correspond to the positive direction connections in C' are connected to odd nodes in $\hat{C}$ while the edges that correspond to negative direction connections in C' are connected to even nodes in $\hat{C}$. Consider the graph C' in which a node a is faulty iff at least one of the nodes in the supernode corresponding to a in $\hat{C}$ is faulty. It follows from Theorem 3.1 that C' contains a cycle of N healthy nodes. Therefore, $\hat{C}$ must contain a cycle of $2N$ healthy nodes corresponding to the cycle of N healthy nodes in C'. □

Figure 3 shows an example of a 2-FT cycle of degree 3 with $2k^2 = 8$ spares.

Figure 3: A degree-3 2-fault-tolerant cycle with 8 spare nodes.

3.2.2 Hierarchical fault-tolerant meshes

We will now show how hierarchical constructions can be used to reduce the degree of the graph M' of Theorem 3.2. We will start with an approach that reduces the degree to $\Theta(\sqrt{k})$. We will then consider a more powerful technique that reduces the degree to a constant. The first approach uses the following graph as a supernode.

Definition: Let H_n be a graph with n nodes and degree 3 (if n is even) or degree 4 (if n is odd) such that for every pair of distinct nodes in H_n, there is a Hamiltonian path that has those nodes as endpoints. H_n graphs have been created for all $n \geq 2$ [5]. See Figure 4 for an example.

Figure 4: Examples of Hamiltonian graphs by Moon with minimal degree. The number of nodes is even in (a) and is odd in (b).

Construction 3.4 *Let k, r, c, n and s be positive integers where $r, c, s \geq 2$, $rc \geq k^2+k+1$, and $2rs = c = n$, let $V = \{c + ik | 0 \leq i \leq k\}$), and let the graph $M' = C(rc + k^2, \{1, k+1\} \cup V)$. Let $\hat{M}$ be the hierarchical graph obtained from M' by replacing each node in M' by a supernode H_{2s}. Divide the nodes in each supernode arbitrarily into two halves of s nodes each. Add connections between supernodes as follows:*

1. *Connect each node in each supernode i to every node in supernodes $i-1$, $i+1$, $i-k-1$ and $i+k+1$ (all modulo $rc+k^2$). These edges, called horizontal edges, contribute $8s$ to the degree of each node.*
2. *For each offset $v \in V$ and for every supernode i, connect one of the nodes in the second half of supernode i to one of the nodes in the first half of supernode $(i+v) \bmod (rc+k^2)$. These edges, called vertical edges, should be evenly distributed among the nodes in each half of each supernode, so they contribute at most $\lceil (k+1)/s \rceil$ to the degree of each node.*

Note that the degree of $\hat{M}$ is at most $8s + \lceil (k+1)/s \rceil + 3$. Choosing $s = \Theta(\sqrt{k})$ yields a graph $\hat{M}$ with degree $O(\sqrt{k})$ and with $O(k^{5/2})$ spare nodes.

Theorem 3.5 *The graph $\hat{M}$ defined in Construction 3.4 is a k-FT $M_{n,n}$.*

Proof: Consider the graph M' of Theorem 3.2 in which a node a' is faulty iff at least one of the nodes in the supernode corresponding to a' in $\hat{M}$ is faulty. It follows from Theorem 3.2 that M' contains a healthy $M_{r,c}$ subgraph. We will show that this implies that $\hat{M}$ contains a healthy $M_{n,n}$ subgraph.

Let a' be any node in the healthy $M_{r,c}$ subgraph of M' and let $\hat{a}$ be the supernode in $\hat{M}$ corresponding to a'. We will view $\hat{a}$ as a column of $2s$ nodes in $M_{n,n}$. Note that a' has vertical neighbors $a' - v_1 \bmod (rc + k^2)$ and $a' + v_2 \bmod (rc + k^2)$, where v_1 and v_2 are in V. Let t be the node in the first half of $\hat{a}$ that is connected to a node in supernode $\hat{a} - v_1 \bmod (rc + k^2)$ and let b be the node in the second half of $\hat{a}$ that is connected to a node in supernode $\hat{a} + v_2 \bmod (rc + k^2)$. We will view t as being the top node and b as being the bottom node in the column of $2s$ nodes formed by $\hat{a}$. Recall that for every pair of nodes in H_{2s}, there is a Hamiltonian path that has those nodes as endpoints. Therefore, we can use the Hamiltonian path with endpoints t and b as the vertical connections within $\hat{a}$. Furthermore, the connections between the node b in one supernode and the node t in the next supernode provide the vertical connections between supernodes. Finally, note that a' has horizontal neighbors $a' - x_1 \bmod (rc + k^2)$ and $a' + x_2 \bmod (rc + k^2)$ where x_1 and x_2 are in $\{1, k+1\}$. Because each node in $\hat{a}$ is connected to every node in supernodes $\hat{a} - x_1 \bmod (rc+k^2)$ and $\hat{a} + x_2 \bmod (rc+k^2)$, the horizontal connections between supernodes are also present. □

We will now show how the use of a different supernode graph can yield a k-FT mesh with $O(k^3)$ spare nodes and constant degree. The following graph will be used as the supernode graph.

Definition: The graph P_k consists of $2k + 4$ nodes. This graph consists of two parts, denoted S_1 and S_2, each of which is the graph $C(k + 2, \{1, 2\})$, plus an edge connecting node $k + 1$ in S_1 with node $k + 1$ in S_2. See Figure 5 for an example of P_6.

Now we describe the construction of a k-FT mesh based on the graph P_k as a supernode.

Construction 3.6 *Let k, r, c, and n be positive integers where $r, c \geq 2$, $rc \geq k^2 + k + 1$, and $(2k+4)r = c = n$, let $V = \{c + ik | 0 \leq i \leq k\})$, and let $M' = C(rc + k^2, \{1, k+1\} \cup V)$.*

Let $\tilde{M}$ be the hierarchical graph obtained from M' by replacing each node in M' by the supernode P_k. Add connections between supernodes as follows:

1. *Connect each node $j \in S_1$ of supernode i to nodes $\{j-2, j-1, j, j+1, j+2\} \bmod (k+2)$ in S_1 of supernodes $\{i-1, i+1, i-k-1, i+k+1\} \bmod (rc+k^2)$. These edges, called horizontal edges, contribute 20 to the degree of each node in S_1.*
2. *Connect each node $j \in S_2$ of supernode i to nodes $\{j-2, j-1, j, j+1, j+2\} \bmod (k+2)$ in S_2 of supernodes $\{i-1, i+1, i-k-1, i+k+1\} \bmod (rc+k^2)$. These edges, also called horizontal edges, contribute 20 to the degree of each node in S_2.*
3. *Connect each node $j \in S_2$ of supernode i, where $0 \leq j \leq k$, to node $j \in S_1$ of supernode $(i + c + jk) \bmod (rc + k^2)$. These edges, called vertical edges, correspond to*

Figure 5: An example of the graph P_6.

the $k+1$ offsets in V and contribute 1 to the degree of each node numbered less than $k+1$ in each half of each supernode.

Note that the degree of $\tilde{M}$ is 25. The fact that $\tilde{M}$ is a k-FT mesh relies on the following lemmas.

Lemma 3.7 *Consider the subgraph $S = S_1$ (or equivalently, $S = S_2$) of P_k. There exists a set of paths, $\{Q_0, Q_1, \ldots, Q_k\}$, such that for each i, $0 \leq i \leq k$, Q_i is a Hamiltonian path through S with endpoints i and $k+1$, and for each i, $0 \leq i < k$, and for each j, $0 \leq j \leq k+1$, if a is the j-th node in Q_i and b is the j-th node in Q_{i+1}, then $(a-b) \equiv x \pmod{k+2}$ where $x \in \{-2,-1,0,1,2\}$.*

Proof: For each i, $0 \leq i \leq k$, define Q_i as follows. Start at i and traverse the 1-offset edges in the positive direction until node k is reached. Then traverse the 2-offset edges in the positive direction until either node $i-1$ or $i-2$ is reached. If node $i-1$ is reached, traverse the 1-offset edge to node $i-2$ and then traverse the 2-offset edges in the negative direction until node $k+1$ is reached. On the other hand, if node $i-2$ is reached before node $i-1$, traverse the 1-offset edge to node $i-1$ and then traverse the 2-offset edges in the negative direction until node $k+1$ is reached. See Figure 6 for an example of the paths Q_4 and Q_5 in S_1 of P_6.

If $i \leq a \leq k-1$, then $b = a+1$. If $a = k$, then $b = 0$. If $a = k+1$, then $b = a$. If $0 \leq a \leq i-1$, we have the following cases: (i) if a is even and $0 \leq a \leq i-2$, then $b = a+2$, (ii) if a is even and $a = i-1$, then $b = a+1$, and (iii) if a is odd and $1 \leq a \leq i-1$, then $b = a$. Therefore, in every case $(a-b) \equiv x \pmod{k+2}$ where $x \in \{-2,-1,0,1,2\}$. □

Figure 6: An example of two paths in S_1 of P_6 starting from nodes 4 and 5, respectively.

Lemma 3.8 *Let k, r, c, n, V, and M' be as defined in Construction 3.6. Consider any set of k faulty nodes in M' and let M be the healthy mesh $M_{r,c}$ in M' that is obtained by applying Theorem 3.2. Let a, b, a', and b' be any nodes in M' such that a and b are horizontal neighbors in M, a' and b' are horizontal neighbors in M, a and a' are vertical neighbors in M, and b and b' are vertical neighbors in M. If $a' = (a + c + ik) \bmod (rc + k^2)$ and $b' = (b + c + jk) \bmod (rc + k^2)$ where $0 \le i, j \le k$, then $|i - j| \le 1$.*

Proof: Assume without loss of generality that a is to the left of b in M and a' is to the left of b' in M. Note that $(b - a) \equiv x \pmod{rc + k^2}$ where $x \in \{1, k + 1\}$ and $(b' - a') \equiv x' \pmod{rc + k^2}$ where $x' \in \{1, k+1\}$. Therefore, $(i - j)k \equiv (c + ik) - (c + jk) \equiv (a' - a) - (b' - b) \equiv x - x' \pmod{rc + k^2}$, which implies that $|i - j| \le 1$. □

Theorem 3.9 *The graph $\tilde{M}$ defined in Construction 3.6 is a k-FT $M_{n,n}$ and has constant degree and $2k^3 + 4k^2$ spare nodes.*

Proof: The proof is analogous to that of Theorem 3.5. In particular, as in the proof of Theorem 3.5, we project the faults in $\tilde{M}$ onto M' and use Theorem 3.2 to find a healthy $M_{r,c}$ subgraph of M'.

Let a' be any node in the healthy $M_{r,c}$ subgraph of M' and let $\tilde{a}$ be the corresponding supernode in $\tilde{M}$. We view $\tilde{a}$ as a column of $2k + 4$ nodes in $M_{n,n}$. We find top and bottom nodes t and b in $\tilde{a}$ as in the proof of Theorem 3.5, and we use Lemma 3.7 (twice) to create a Hamiltonian path through $\tilde{a}$ with endpoints t and b. Then, let b' be a node that is horizontally adjacent to a' in the healthy $M_{r,c}$ subgraph of M', and let $\tilde{b}$ be the corresponding supernode in $\tilde{M}$. It follows from Lemma 3.8 that the top nodes in $\tilde{a}$ and $\tilde{b}$ have positions within their supernodes that differ by at most one. A similar argument applies to the bottom nodes in $\tilde{a}$ and $\tilde{b}$. Therefore, it follows from Lemma 3.7 that for each i, $0 \le i < 2k + 4$, the i-th node in the Hamiltonian path in $\tilde{a}$ has a horizontal connection to the i-th node in the Hamiltonian path in $\tilde{b}$, which completes the proof. □

Hence, we have obtained a construction of a k-FT two-dimensional mesh with constant degree and $O(k^3)$ spare nodes. Although the construction given above is for a k-FT $M_{n,n}$ where n is a multiple of $2k+4$, it is straightforward to generalize the construction to arbitrary values of n as follows.

Construction 3.10 *Let k, r, c, and n be positive integers where $r, c \geq 2$, $rc \geq k^2+k+1$, $r = \lfloor n/(2k+4) \rfloor$, and $c = n$, let $V = \{c + ik | 0 \leq i \leq k\})$, and let $M' = C(rc+k^2, \{1, k+1\} \cup V)$.*

Let $n \bmod (2k+4) = \alpha$. If $\alpha = 0$, let $\hat{M}$ be the graph $\tilde{M}$ defined in Construction 3.6. If $\alpha \neq 0$, first define the graph P'_k from P_k as follows. Add a node, denoted x, to P_k, connect node x to node $k+1$ of S_1 in P_k, and connect node x to node $k+1$ of S_2 in P_k. Let $\hat{M}$ be the graph obtained by replacing each of the first $\alpha n + k^2$ nodes in M' by the supernode P'_k and replacing each of the remaining nodes in M' by the supernode P_k. Add connections between supernodes as follows:

1. *Ignore the x nodes in the P'_k supernodes and add connections between supernodes as required by Construction 3.6.*
2. *For each supernode i, where $0 \leq i < \alpha n + k^2$, connect node x in supernode i to node x in supernode j, where $j \in \{i-1, i+1, i-k-1, i+k+1\}$ and $0 \leq j < \alpha n + k^2$.*

The following theorem is immediate from the preceding construction.

Theorem 3.11 *Let k and n be positive integers, let $r = \lfloor n/(2k+4) \rfloor$, and let $c = n$. If $r, c \geq 2$ and $rc \geq k^2+k+1$, then there exists a k-FT $M_{n,n}$ with constant degree and $2k^3+5k^2$ spare nodes.*

Although the degree of $\hat{M}$ is increased to 26 (as both node $k+1$ of S_1 and node $k+1$ of S_2 have an edge to node x in the same supernode), one can easily reduce the number of horizontal edges of node $k+1$ to 4 (as opposed 20 of the current definition) so that the degree of $\hat{M}$ remains 25. In fact, we remark that it is possible to reduce the degree still further by using a different graph for each supernode. Specifically, if each supernode is defined to be the product graph of P_k and a 4-node linear array, and if each supernode plays the role of a $(2k+4) \times 4$ submesh, it is possible to obtain a k-FT mesh with degree 12 and $8k^3+16k^2$ spare nodes. The details are omitted.

Finally, we will consider laying out the fault-tolerant graph $\hat{M}$ using Thompson's VLSI model [27]. One of the greatest advantages of two-dimensional mesh networks is that they can be laid out using only short (constant length) wires. The following theorem shows that the fault-tolerant graph $\hat{M}$ may require somewhat longer wires, but the wire lengths are still independent of n.

Theorem 3.12 *It is possible to lay out the graph $\hat{M}$ defined in Construction 3.10 using only wires with length $O(k^3)$.*

Proof: We will begin by presenting a mapping from the nodes in M' to the nodes in a torus network which maintains locality. We will then use standard techniques for laying out torus networks to obtain the final layout of $\hat{M}$. First, consider the case where k^2 is a multiple of c. In this case, lay out the nodes in M' in row-major order on an $(rc+k^2)/c$ by c torus. It is straightforward to verify that any pair of nodes that are connected in M' map to nodes that are in columns of the torus that differ by at most $O(k^2)$ and in rows of the torus that differ by at most $O(1)$. This torus can then be mapped to on an $(rc+k^2)/c$ by c grid by using the standard technique of placing the first half of the torus columns (rows) in increasing order in the even numbered columns (rows) of the grid and the remaining torus columns (rows) in decreasing order in the odd numbered columns (rows) of the grid (see, for example, [21, p. 246]). Finally, each node in M' can be layed out using an $O(k)$ by $O(k)$ square. The vertical tracks between grid columns are $O(k)$ wide and the horizontal tracks between grid rows are $O(k^3)$ wide (to accommodate wires that traverse $O(k^2)$ nodes, each of which is $O(k)$ wide). Thus each wire is of length $O(k^3)$.

Now consider the case where c does not evenly divide k^2. In this case, let $\alpha = k^2 \bmod c$ and use a $(rc+k^2)/c$ by $c+1$ torus. The nodes of M' are placed in the torus in row-major order, with the first α rows receiving $c+1$ nodes and all remaining rows receiving only c nodes. Again, it is straightforward to verify that any pair of nodes that are connected in M' map to nodes that are in columns of the torus that differ by at most $O(k^2)$ and in rows of the torus that differ by at most $O(1)$. This torus can then be laid out as described for the other case. $\Box$

4 Random Faults

In this section we consider random fault distributions. More specifically, we will assume that the fault-tolerant graph contains k faults, and that every configuration of k faulty nodes is equally likely. We will focus on the problem of creating fault-tolerant graphs for the mesh $M_{n,n}$. We will present six constructions for fault-tolerant meshes, analyze their asymptotic fault-tolerance, and study their fault-tolerance for realistic values of n.

The first three constructions are simple generalizations of previously known constructions [10] designed to tolerate worst-case fault distributions, while the remaining three constructions are new. In particular, the fourth construction introduces the concept of adding "dummy faults" in order to provide a fairly regular fault pattern. The fifth construction introduces the use of a 2 by 2 "submeshes", and the sixth construction combines the use of dummy faults with the use of submeshes.

Throughout this section let $T_1(n)$ denote $C(n^2, \{n-1, n\})$ and let $T_2(n)$ denote $D(n^2, \{1, n\})$. Recall from Theorems 2.4 and 2.3 that both $T_1(n)$ and $T_2(n)$ contain $M_{n,n}$ as a subgraph.

4.1 The Constructions

In this subsection we will define the six constructions, $M_1(n,k)$ through $M_6(n,k)$. The first two constructions are based on the target graph $T_1(n)$.

Definition: The graph $M_1(n,k) = C(n^2+k, \{n-1, n, n+1\})$.

Definition: The graph $M_2(n,k) = C(n^2+k, \{n-1, n, n+1, n+2\})$.

Note that $M_1(n,k)$ has degree 6 and $M_2(n,k)$ has degree 8. The idea behind both of these constructions is that they can tolerate faults by using the larger offsets ($n+1$ and $n+2$) to jump over faults. The remaining constructions are based on the target graph $T_2(n)$.

Definition: The graph $M_3(n,k) = C(n^2+k, \{1, 2, n, n+1\})$.

Definition: The graph $M_4(n,k) = C(n^2+n+k, \{1, 2, n+1, n+2\})$.

Note that both $M_3(n,k)$ and $M_4(n,k)$ have degree 8. In both constructions, the 1-offset and 2-offset edges of the fault-tolerant graphs implement the 1-offset edges of the target graph and the n-offset, $(n+1)$-offset, and $(n+2)$-offset edges of the fault-tolerant graph implement the n-offset edges of the target graph. In particular, the $(n+1)$-offset and $(n+2)$-offset edges of $M_4(n,k)$ can implement the n-offset edges of the target graph, provided that each window of $n+1$ consecutive nodes contains at least 1 fault and each window of $n+2$ consecutive nodes contains at most 2 faults. Although it is very unlikely (or impossible) that each window of $n+1$ consecutive nodes contains at least 1 fault, we can view up to n healthy nodes as being "dummy faults" (because there are $n+k$ spares) in order to satisfy this requirement.

Finally, constructions $M_5(n,k)$ and $M_6(n,k)$ are hierarchical constructions based on $M_3(n/2,k)$ and $M_4(n/2,k)$, respectively. They are defined only for even values of n.

Definition: The graph $M_5(n,k)$ is created from $M_3(n/2,k)$ as follows:

1. Create $n' = (n/2)(n/2)+k = n^2/4+k$ squares (that is, cycles of length 4) numbered 0 through $n'-1$.
2. For each square i, connect the upper right corner of i to the upper left corners of $(i+1) \bmod n'$ and $(i+2) \bmod n'$, and connect the lower right corner of i to the lower left corners of $(i+1) \bmod n'$ and $(i+2) \bmod n'$.
3. For each square i, connect the lower left corner of i to the upper left corners of $(i+n) \bmod n'$ and $(i+n+1) \bmod n'$, and connect the lower right corner of i to the upper right corners of $(i+n) \bmod n'$ and $(i+n+1) \bmod n'$.

Definition: The graph $M_6(n,k)$ is created from $M_4(n/2,k)$ as follows:

1. Create $n' = (n/2)(n/2)+(n/2)+k = n^2/4+n/2+k$ squares (that is, cycles of length 4) numbered 0 through $n'-1$.

2. For each square i, connect the upper right corner of i to the upper left corners of $(i+1) \bmod n'$ and $(i+2) \bmod n'$, and connect the lower right corner of i to the lower left corners of $(i+1) \bmod n'$ and $(i+2) \bmod n'$.

3. For each square i, connect the lower left corner of i to the upper left corners of $(i+n+1) \bmod n'$ and $(i+n+2) \bmod n'$, and connect the lower right corner of i to the upper right corners of $(i+n+1) \bmod n'$ and $(i+n+2) \bmod n'$.

Note that both $M_5(n,k)$ and $M_6(n,k)$ have degree 6. The idea behind these constructions is that the squares act as 2 by 2 submeshes and the graphs can be reconfigured if the corresponding fault-tolerant graph (namely $M_3(n/2,k)$ or $M_4(n/2,k)$) can tolerate faults located in the positions corresponding to the faulty squares (see the proof of Theorem 3.5 for a description of hierarchical fault-tolerant graphs).

4.2 Asymptotic Techniques

In this section we will present several definitions and lemmas that will be useful in establishing the asymptotic fault-tolerance of the six constructions given above. It will be assumed throughout that $k \leq n/2$ and $k = o(n)$.

Definition: A graph *tolerates* $\Theta(f(n))$ *random faults* iff $o(f(n))$ random faults can be tolerated with a probability that is $1 - o(1)$ and $\omega(f(n))$ random faults can be tolerated with a probability that is $o(1)$.

Definition: Given a circulant graph with x nodes, and given integers y and z where $0 \leq y, z < x$, the y-node window starting at z, denoted $W(y,z)$, consists of the y nodes in the graph numbered $z, z+1 \bmod x, \ldots, z+y-1 \bmod x$.

Definition: Given a circulant graph with x nodes, and given integers y and z where $0 \leq y, z < x$, the distance between y and z, denoted $dist(y,z)$, is the minimum of $z-y \bmod x$ and $y-z \bmod x$, and nodes x and y are *consecutive* iff $dist(y,z) = 1$.

Definition: Given a circulant graph with x nodes, and given integers y and z where $1 \leq y < x$ and $0 \leq z < x$, *the y-th healthy node following (respectively, preceding)* z is the healthy node a such that there are exactly y healthy nodes in the set $\{z+1 \bmod x, z+2 \bmod x, \ldots, a\}$ (respectively, $\{z-1 \bmod x, z-2 \bmod x, \ldots, a\}$).

We will consider the fault-tolerance of $M_1(n,k)$ and $M_2(n,k)$ with respect to the target graph $T_1(n)$, and the fault-tolerance of $M_3(n,k)$, $M_4(n,k)$, $M_5(n,k)$ and $M_6(n,k)$ with respect to the target graph $T_2(n)$. For $M_1(n,k)$, $M_2(n,k)$, $M_3(n,k)$ and $M_4(n,k)$, we will consider only embeddings of the target graph in which node 0 of the target graph maps to some healthy node h in the fault-tolerant graph, and for each i, node i in the target graph maps to the i-th healthy node following node h. For $M_5(n,k)$ and $M_6(n,k)$, we will consider only embeddings obtained by viewing squares that contain faults as representing faulty nodes in the corresponding fault-tolerant graph (namely $M_3(n/2,k)$ or $M_4(n/2,k)$).

Lemma 4.1 *Let M' be a circulant graph with $\Theta(n^2)$ nodes and let $y = \Theta(n)$. Assume that M' contains k randomly located faulty nodes. If k is $o(n^{1/2})$, the probability that there exists a node i such that $W(y, i)$ contains two or more faults is $o(1)$.*

Proof: Given any two faults a and b, the probability that there exists a node i such that both a and b lie in $W(y, i)$ is $\Theta(n^{-1})$. There are $o(n)$ distinct pairs of faults, so the probability that there exists a node i such that $W(y, i)$ contains two or more faults is $o(n^{-1}n) = o(1)$. □

Lemma 4.2 *Let M' be a circulant graph with $\Theta(n^2)$ nodes and let $W = W(y, z_1), W(y, z_2), \ldots, W(y, z_q)$ be a collection of q mutually disjoint y-node windows in M', where $y = \Theta(n)$ and $q = \Theta(n)$. Assume that M' contains k randomly located faulty nodes. If k is $\omega(n^{1/2})$, the probability that there exists a window in W that contains two or more faults is $1 - o(1)$.*

Proof: Divide the faults into halves. After the first half of the faults have been placed, if no window in W contains two or more faults then there must be $\omega(n^{3/2})$ healthy locations, each of which lies within a window in W that contains a fault. Therefore, the probability that any given fault in the second half will lie in a window in W that contains another fault is $\omega(n^{-1/2})$. As a result, the probability that no window in W contains two or more faults after all of the faults have been placed is at most $(1 - n^{-1/2})^{\omega(n^{1/2})} = (1 - n^{-1/2})^{n^{1/2}\omega(1)} = (1/e)^{\omega(1)} = o(1)$. □

Lemma 4.3 *Let M' be a circulant graph with $\Theta(n^2)$ nodes and let $y = \Theta(n)$. Assume that M' contains k randomly located faulty nodes. If k is $o(n^{2/3})$, the probability that there exists a node i such that $W(y, i)$ contains three or more faults is $o(1)$.*

Proof: Given any three faults, the probability that there exists a node i such that all three faults lie in $W(y, i)$ is $\Theta(n^{-2})$. There are $o(n^2)$ distinct sets of three faults, so the probability that there exists a node i such that $W(y, i)$ contains three or more faults is $o(n^{-2}n^2) = o(1)$. □

Lemma 4.4 *Let $f(n)$ be any function such that $1 \leq f(n) \leq n$. Given $\omega(n)$ independent Bernoulli trials, each of which has a probability of success of at least $1/f(n)$, the probability of at least $n/f(n)$ successes is $1 - o(1)$.*

Proof: Divide the trials into $\omega(n/f(n))$ groups, each of which contains at least $\lceil f(n) \rceil$ trials. Given any one group of trials, the probability of at least one success in that group is at least $1/2$. Therefore, given any $2\lceil n/f(n) \rceil$ groups, the probability of at least $n/f(n)$ successes is at least $1/2$. This implies that the probability that the entire set of $\omega(n)$ trials contains at least $n/f(n)$ successes is at least $1 - (1/2)^{\omega(1)} = 1 - o(1)$. □

Lemma 4.5 *Let M' be a circulant graph with $\Theta(n^2)$ nodes and let $W = W(y, z_1), W(y, z_2), \ldots, W(y, z_q)$ be a collection of q mutually disjoint y-node windows in M', where $y = \Theta(n)$ and $q = \Theta(n)$. Assume that M' contains k randomly located faulty nodes. If k is $\omega(n^{2/3})$, the probability that there exists a window in W that contains three or more faults is $1 - o(1)$.*

Proof: Divide the faults into three groups, each of which contains $\omega(n^{2/3})$ faults. Consider the three following statements:

Statement 1: At least $n^{2/3}$ windows in W contain at least one fault each.

Statement 2: At least $n^{1/3}$ windows in W contain at least two faults each.

Statement 3: There exists a window in W that contains three or more faults.

For all sufficiently large n, after the first group of faults has been placed, at least one of the three statements above must be true.

First, consider the situation in which Statement 1 is true after the first group of faults has been placed. For each fault in the second group, consider that fault to be a success iff it lies in a window in W that contains a fault from the first group. Given any fault in the second group, the probability that it is a success is $\Omega(n^{-1/3})$. It follows from Lemma 4.4 that with probability $1 - o(1)$ at least $n^{1/3}$ faults in the second group are successes.

Therefore, regardless of which statement is true after the first group of faults is placed, there is a probability of at least $1 - o(1)$ that after the second group of faults is placed, either Statement 2 or Statement 3 (or both) is true. Now consider the situation in which Statement 2 is true and Statement 3 is false after the second group of faults is placed. For each fault in the third group, consider that fault to be a success iff it lies in a window in W that contains at least two faults from the union of the first and second groups. Given any fault in the third group, the probability that it is a success is $\Omega(n^{-2/3})$. It follows from Lemma 4.4 that with probability $1 - o(1)$ at least one fault in the third group is a success.

As a result, in any case there is a probability of at least $1 - o(1)$ that after all of the faults have been placed, Statement 3 holds. $\Box$

Lemma 4.6 *Let M' be a circulant graph with $\Theta(n^2)$ nodes. Assume that M' contains k randomly located faulty nodes. If k is $o(n)$, the probability that there exists a pair of faults that are consecutive is $o(1)$.*

Proof: Given any two faults, the probability that they are consecutive is $\Theta(n^{-2})$. There are $o(n^2)$ distinct pairs of faults, so the probability that there exists a pair of faults that are consecutive is $o(n^{-2}n^2) = o(1)$. $\Box$

4.3 Asymptotic Results

Given the previous definitions and lemmas, we are now ready to establish the asymptotic fault-tolerance of the six constructions.

Lemma 4.7 *Assume that $M_1(n, k)$ contains k faulty nodes. $M_1(n, k)$ tolerates the faults iff for each i, $0 \le i < n^2 + k$, $W(n + 1, i)$ contains at most one fault.*

Proof: First, assume that for each i, $0 \le i < n^2+k$, $W(n+1,i)$ contains at most one fault. In this case, given any healthy node i, $W(n+2,i)$ contains at most one fault. Therefore, there is an edge between each healthy node and both the $(n-1)$-st healthy node following it and the n-th healthy node following it. As a result, $M_1(n,k)$ contains a healthy copy of $T_1(n)$.

Next, assume that there exists an i, $0 \le i < n^2+k$, such that $W(n+1,i)$ contains two or more faults. Let a be the first healthy node preceding i. Let a' be the node in $T_1(n)$ that maps to a, let b' be node $a'+n \bmod n^2$ in $T_1(n)$, and let b be the node to which b' maps. Note that b is the n-th healthy node following a, so $b \notin W(n+1,i)$ and a and b are not connected to one another. $\Box$

Theorem 4.8 *The graph $M_1(n,k)$ tolerates $\Theta(n^{1/2})$ random faults.*

Proof: The proof is immediate from Lemmas 4.1, 4.2 and 4.7. $\Box$

The proof of the following lemma is analogous to that of Lemma 4.7 and is omitted.

Lemma 4.9 *Assume that $M_2(n,k)$ contains k faulty nodes. $M_2(n,k)$ tolerates the faults iff for each i, $0 \le i < n^2+k$, $W(n+2,i)$ contains at most two faults.*

Theorem 4.10 *The graph $M_2(n,k)$ tolerates $\Theta(n^{2/3})$ random faults.*

Proof: The proof is immediate from Lemmas 4.3, 4.5 and 4.9. $\Box$

Lemma 4.11 *Assume that $M_3(n,k)$ contains k faulty nodes. $M_3(n,k)$ tolerates the faults if for each i, $0 \le i < n^2+k$, $W(n+1,i)$ contains at most one fault.*

Proof: Given any healthy node i, $W(n+2,i)$ contains at most one fault. Therefore, there is an edge between each healthy node and both the first healthy node following it and the n-th healthy node following it. As a result, $M_3(n,k)$ contains a healthy copy of $T_2(n)$. $\Box$

Lemma 4.12 *Assume that $M_3(n,k)$ contains k faulty nodes. $M_3(n,k)$ does not tolerate the faults if there exist x and y, where $0 \le x,y < n^2+k$, $dist(x,y) \ge 2n$, $W(n+1,x)$ contains at least two faults, and $W(n+1,y)$ contains at least two faults.*

Proof: Assume for the sake of contradiction that the faults can be tolerated. Let a be the first healthy node preceding x and let a' be the node in $T_2(n)$ that maps to a. If there exists a node b' in $T_2(n)$ where $b' = a'+n$, let b be the node to which b' maps. Note that b is the n-th healthy node following a, so $b \notin W(n+1,x)$ and a and b are not connected to one another. Therefore, no such node b exists, which implies that $a' \ge n^2-n$.

Let c be the first healthy node preceding y and let c' be the node in $T_2(n)$ that maps to c. A similar argument shows that $c' \ge n^2-n$. As a result, $|a'-c'| \le n-1$, so $dist(a,c) \le n-1+k$ and $dist(x,y) \le n-1+2k < 2n$, which is a contradiction. $\Box$

Theorem 4.13 *The graph $M_3(n,k)$ tolerates $\Theta(n^{1/2})$ random faults.*

Proof: If the number of faults is $o(n^{1/2})$, it follows from Lemmas 4.1 and 4.11 that the probability of tolerating the faults is $1-o(1)$. If the number of faults is $\omega(n^{1/2})$, divide the faults into halves. Let $W_1 = W(y,0), W(y,y), W(y,2y), \ldots, W(y,qy)$ and let $W_2 = W(y,(q+2)y), W(y,(q+3)y), W(y,(q+4)y)\ldots, W(y,2qy)$ where $y = n+1$ and $q = \lfloor n/4 \rfloor$. Apply Lemma 4.2 to the first half of the faults with $W = W_1$, apply Lemma 4.2 to the second half of the faults with $W = W_2$, and apply Lemma 4.12 to complete the proof. $\Box$

Definition: Given a circulant graph with x nodes, a *block of healthy nodes* is a window $W(y,i)$, where $1 \leq y < x$ and $0 \leq i < x$, consisting solely of healthy nodes such that both node $i-1 \bmod x$ and node $i+y \bmod x$ are faulty.

Consider the following algorithm for adding dummy faults to $M_4(n,k)$:

Algorithm A: Consider each block of healthy nodes separately. Assume a block consists of y healthy nodes. There are three cases based on the value of y.

Case 1: $y \leq n$. In this case, do not add any dummy faults to the block.

Case 2: $n+1 \leq y \leq 2n$. In this case, add one dummy fault to the block. Place the dummy fault in the middle of the block so that it divides the block into two subblocks of healthy nodes, the first of which has $\lceil (y-1)/2 \rceil$ nodes and the second of which has $\lfloor (y-1)/2 \rfloor$ nodes.

Case 3: $2n+1 \leq y$. In this case, add two dummy faults that divide the block into three subblocks of healthy nodes, the first of which has $n-1$ nodes, the second of which has $z = y-2n$ nodes, and the third of which has $n-1$ nodes. Let a and b denote these two dummy nodes. Then add an additional $x = \lfloor z/(n+1) \rfloor$ dummy faults between a and b. This leaves $w = z - x$ healthy nodes in the block, which are divided into $x+1$ subblocks of healthy nodes by the x dummy faults. Distribute the dummy faults so that each subblock has length $\lfloor w/(x+1) \rfloor$ or $\lceil w/(x+1) \rceil$.

The following lemmas establish properties of Algorithm A.

Lemma 4.14 *Given w, x and z in Case 3 above, $xn \leq w \leq (x+1)n$.*

Proof: Because $x = \lfloor z/(n+1) \rfloor$, $z \geq x(n+1)$ and $w = z - x \geq xn$. Because $x = \lfloor z/(n+1) \rfloor$, $z \leq xn + n + x$ and $w = z - x \leq xn + n = (x+1)n$. $\Box$

Lemma 4.15 *After applying Algorithm A, no block of $n+1$ or more healthy nodes exists.*

Proof: If there is a block of $n+1 \leq y \leq 2n$ healthy nodes prior to applying Algorithm A, the algorithm adds a dummy node that divides the block into subblocks of at most $\lceil (y-1)/2 \rceil \leq n$ healthy nodes each. If there is a block of $2n+1 \leq y$ healthy nodes prior to applying Algorithm A, the algorithm adds dummy nodes a and b that divide the block into subblocks of $n-1$, $z = y - 2n$, and $n-1$ healthy nodes, each. Then $x = \lfloor z/(n+1) \rfloor$ dummy faults are added to the subblock of z healthy nodes, leaving $w = z - x$ healthy nodes. These w healthy nodes occur in subblocks of length at most $\lceil w/(x+1) \rceil \leq \lceil (x+1)n/(x+1) \rceil \leq n$. □

Lemma 4.16 *After applying Algorithm A, no dummy fault is consecutive with another (actual or dummy) fault, provided that $n \geq 2$.*

Proof: If there is a block of $n+1 \leq y \leq 2n$ healthy nodes prior to applying Algorithm A, the algorithm adds a dummy node that divides the block into subblocks of at least $\lfloor (y-1)/2 \rfloor \geq \lfloor n/2 \rfloor \geq 1$ healthy nodes each. If there is a block of $2n+1 \leq y$ healthy nodes prior to applying Algorithm A, the algorithm adds dummy nodes a and b that divide the block into subblocks of $n-1$, $z = y - 2n$, and $n-1$ healthy nodes, each. Then $x = \lfloor z/(n+1) \rfloor$ dummy faults are added to the subblock of z healthy nodes, leaving $w = z - x$ healthy nodes. If $x = 0$, there are $w = z - 0 = y - 2n \geq 1$ healthy nodes between dummy faults a and b. If $x \geq 1$, the w healthy nodes occur in subblocks of at least $\lfloor w/(x+1) \rfloor \geq \lfloor xn/(x+1) \rfloor \geq \lfloor n/2 \rfloor \geq 1$ nodes each. □

Lemma 4.17 *Consider any configuration of actual faults such that no two faults are consecutive and there does not exist a node i such that $W(2n+3, i)$ contains three or more faults, where $n \geq 2$. After applying Algorithm A to this configuration of faults, no two (actual or dummy) faults will be consecutive and there will not exist a node j such that $W(n+2, j)$ contains three or more (actual or dummy) faults.*

Proof: The fact that no two faults will be consecutive follows immediately from the preceding lemma. Now assume for the sake of contradiction that after applying Algorithm A, there exists a node j such that $W(n+2, j)$ contains three or more faults. Clearly, $W(n+2, j)$ must contain at least one dummy fault. Select one such dummy fault and denote it as d, and let C denote the block of y originally healthy nodes containing d. Clearly, $y \geq n+1$.

If $n+1 \leq y \leq 2n$, then d is the only dummy fault in C, so either $W(n+2, j)$ contains two actual faults or $W(n+2, j)$ contains some other dummy fault located in some other block of originally healthy nodes. First, consider the case where $W(n+2, j)$ contains two actual faults. Let e and f denote these actual faults. Either y lies between e and f or it does not. If y lies between e and f, $W(n+2, j)$ must contain at least $y+2 \geq n+3$ nodes, which is a contradiction. Thus y does not lie between e and f. Now let y' denote the number of nodes between e and f. Because $W(n+2, j)$ contains only $n+2$ nodes and because there are at least $\lfloor (y-1)/2 \rfloor \geq y/2 - 1$ nodes between d and every actual fault, it follows that $y' + y/2 + 2 \leq n + 2$, which implies that $y' \leq n - y/2$ and there were three actual faults within a window of $y + y' + 3 \leq n + y/2 + 3 \leq 2n + 3$ nodes, which is a contradiction.

Now, consider the case where $W(n+2, j)$ contains a dummy fault located in another block of originally healthy nodes. Let d' denote such a dummy fault and let C' denote the block of

y' originally healthy nodes containing d'. Clearly, $y' \leq 2n$, because otherwise there would be at least $n-1$ healthy nodes between d' and the nearest actual fault. However, note that if C' follows C then there are at least $\lfloor (y-1)/2 \rfloor \geq \lfloor n/2 \rfloor$ consecutive healthy nodes following d and at least $\lceil (y'-1)/2 \rceil \geq \lceil n/2 \rceil$ consecutive healthy nodes preceding d', which implies that $W(n+2,j)$ contains at least $n+3$ nodes, which is a contradiction. The case in which C' precedes C is analogous.

If $2n+1 \leq y$, then either $W(n+2,j)$ contains at least one actual fault and at least one dummy fault, or else $W(n+2,j)$ contains three dummy faults and no actual faults. If $W(n+2,j)$ contains at least one actual fault and at least one dummy fault, then it must contain the $n-1$ healthy nodes which separate the dummy faults in C from the actual faults. Furthermore, because no two (actual or dummy) faults are consecutive, $W(n+2,j)$ must contain at least $n+3$ nodes, which is a contradiction. On the other hand, if $W(n+2,j)$ contains three dummy faults and no actual faults, let a, b, w, x, and z be as defined in Case 3 of Algorithm A. It follows that $x \geq 1$ and that $W(n+2,j)$ contains at least two blocks of $\lfloor w/(x+1) \rfloor$ or more healthy nodes in addition to the three dummy faults. However, the fact that $x \geq 1$ implies that $z \geq n+1$. Therefore, the dummy faults designated a and b cannot both be in $W(n+2,j)$, so it follows that $x \geq 2$. Therefore, $\lfloor w/(x+1) \rfloor \geq \lfloor xn/(x+1) \rfloor \geq \lfloor 2n/3 \rfloor \geq n/2$, so $W(n+2,j)$ contains at least n healthy nodes and three dummy faults, which is a contradiction. □

Lemma 4.18 *After applying Algorithm A, at least n^2 healthy nodes remain.*

Proof: First, we will show that Algorithm A adds at most one dummy fault per $n+1/3$ originally healthy nodes. In Case 2 of Algorithm A, one dummy fault is added to a block of at least $n+1$ originally healthy nodes. In Case 3 of Algorithm A, if two dummy faults are added there are at least $2n+1$ originally healthy nodes in the block, so at most one dummy fault is added per $n+1/2$ originally healthy nodes. In Case 3 of Algorithm A, if $i \geq 3$ dummy faults are added there are at least $in+i-2$ originally healthy nodes in the block, so at most one dummy fault is added per $n+(i-2)/i$ originally healthy nodes. This quantity is minimized when $i=3$, at which point one dummy fault is added per $n+1/3$ originally healthy nodes.

Now consider the case in which exactly k actual faults exist. In this case there must be n^2+n originally healthy nodes, so at most $\lfloor (n^2+n)/(n+1/3) \rfloor \leq n$ dummy faults are added, and at least n^2 healthy nodes remain. Now consider the case in which $k-x$ actual faults exist, where $x \geq 1$. At most $\lfloor (n^2+n+x)/(n+1/3) \rfloor \leq \lfloor (n^2+n)/(n+1/3) \rfloor + \lceil x/(n+1/3) \rceil \leq n+x$ dummy faults are added, and at least n^2 healthy nodes remain. □

The proofs of the following two lemmas are analogous to those of Lemmas 4.11 and 4.12, and are omitted.

Lemma 4.19 *Assume that $M_4(n,k)$ contains $f \leq n+k$ (actual or dummy) faults. $M_4(n,k)$ tolerates the faults if no two faults are consecutive and for each i, $0 \leq i < n^2+n+k$, $W(n+1,i)$ contains at least one fault and $W(n+2,i)$ contains at most two faults.*

Lemma 4.20 *Assume that $M_4(n,k)$ contains f faulty nodes. $M_4(n,k)$ does not tolerate the faults if there exist x and y, where $0 \leq x, y < n^2+n+k$, $dist(x,y) \geq 4n$, $W(n+2,x)$ contains at least three faults, and $W(n+2,y)$ contains at least three faults.*

Theorem 4.21 *The graph $M_4(n,k)$ tolerates $\Theta(n^{2/3})$ random faults.*

Proof: First, consider the case where the number of faults is $o(n^{2/3})$. It follows from Lemmas 4.3 and 4.6 that with probability $1-o(1)$, no two actual faults are consecutive and there does not exist a node i such that $W(2n+3,i)$ contains three or more faults. Therefore, it follows from Lemmas 4.18, 4.17 and 4.19 that after applying Algorithm A, the faults can be tolerated with probability $1-o(1)$.

Next, consider the case where the number of faults is $\omega(n^{2/3})$. In this case, divide the faults into halves. Let $W_1 = W(y,0), W(y,y), W(y,2y), \ldots, W(y,qy)$ and let $W_2 = W(y,(q+4)y), W(y,(q+5)y), W(y,(q+6)y) \ldots, W(y,2qy)$ where $y = n+2$ and $q = \lfloor n/4 \rfloor$. Apply Lemma 4.5 to the first half of the faults with $W = W_1$, apply Lemma 4.5 to the second half of the faults with $W = W_2$, and apply Lemma 4.20 to complete the proof. □

The following theorems follow immediately from Theorems 4.13 and 4.21.

Theorem 4.22 *The graph $M_5(n,k)$ tolerates $\Theta(n^{1/2})$ random faults.*

Theorem 4.23 *The graph $M_6(n,k)$ tolerates $\Theta(n^{2/3})$ random faults.*

Table 4.3 summarizes various characteristics, including the asymptotic fault-tolerance, of the six fault-tolerant constructions.

Construction	Symbol	Deg.	No. spares	Offsets	Asymp. FT
M_1	circ6	6	k	$\{n-1, n, n+1\}$	$\Theta(n^{1/2})$
M_2	circ8	8	k	$\{n-1, n, n+1, n+2\}$	$\Theta(n^{2/3})$
M_3	diag8	8	k	$\{1, 2, n, n+1\}$	$\Theta(n^{1/2})$
M_4	diag8r	8	$k+n$	$\{1, 2, n+1, n+2\}$	$\Theta(n^{2/3})$
M_5	diag6	6	$4k$	M_3 + submesh	$\Theta(n^{1/2})$
M_6	diag6r	6	$4k+2n$	M_4 + submesh	$\Theta(n^{2/3})$

Table 1: Comparison of characteristics of the 6 FT meshes.

Notice that $M_2(n,k)$ and $M_4(n,k)$ both have degree 8 and tolerate $\Theta(n^{2/3})$ faults, but $M_4(n,k)$ requires more spares than does $M_2(n,k)$. Thus, the technique of adding dummy faults does not in itself provide a more practical fault-tolerant network. Similarly, notice that $M_1(n,k)$ and $M_5(n,k)$ both have degree 5 and tolerate $\Theta(n^{1/2})$ faults, but $M_5(n,k)$ requires more spares than does $M_1(n,k)$. Thus, the technique of using 2 by 2 submeshes does not in itself provide a more practical fault-tolerant network. However, by combining these two techniques, $M_6(n,k)$ is the only degree 6 network that is capable of tolerating $\Theta(n^{2/3})$ faults.

Finally, we will consider laying out the fault-tolerant graphs presented in this section using Thompson's VLSI model [27]. The following theorem shows that, just like the mesh itself, all of the fault-tolerant constructions can be laid out with constant length wires.

Theorem 4.24 *It is possible to lay out each of the graphs $M_i(n,k)$ where $1 \leq i \leq 6$ using only wires with length $O(1)$.*

Proof: The layouts for graphs $M_1(n,k)$, $M_2(n,k)$, $M_3(n,k)$ and $M_4(n,k)$ follow immediately from the techniques presented in the proof of Theorem 3.12. The layouts for graphs $M_5(n,k)$ and $M_6(n,k)$ follow from the layouts for $M_3(n/2,k)$ and $M_4(n/2,k)$, respectively, by replacing each node by a square of four nodes. $\square$

4.4 Simulation Results

Figures 7 to 9 show the simulation results for the fault tolerance of an $n \times n$ target mesh for $n = 16$, 64 and 256, respectively. When $n = 16$ or 64, the probability given for each construction and each value of k is the result of 10,000 simulation trials. When $n = 256$, the probability given for each construction and each value of k is the result of 1,000 simulation trials.

For each figure, the probability of reconfiguration for each construction of the FT meshes, $M_i(n,k)$ where $1 \leq i \leq 6$, is plotted as a functions of k. Each curve has a name of the form "xyz", where "x" is either "circ" for circulant graph or "diag" for diagonal graph (as the basic target graph), "y" denotes the degree (6 or 8), and "z" is either "r" (designating an extra row of spare nodes or supernodes) or an empty string. The solid lines denote the degree-6 FT meshes while the dotted lines denote the degree-8 FT meshes.

Note that the FT meshes for the three curves from the left tolerate $\Theta(n^{1/2})$ random faults, while the remaining three curves on the right can tolerate $\Theta(n^{2/3})$ random faults. Thus the asymptotic bounds proven above do appear to describe the behavior of these networks for realistic values of n. Also, note that the graph $M_6(n,k)$ (designated "diag6r" in the figures) performs the best out of the degree-6 networks studied, and that it has over a 90% chance of tolerating 12 faults when $n = 64$.

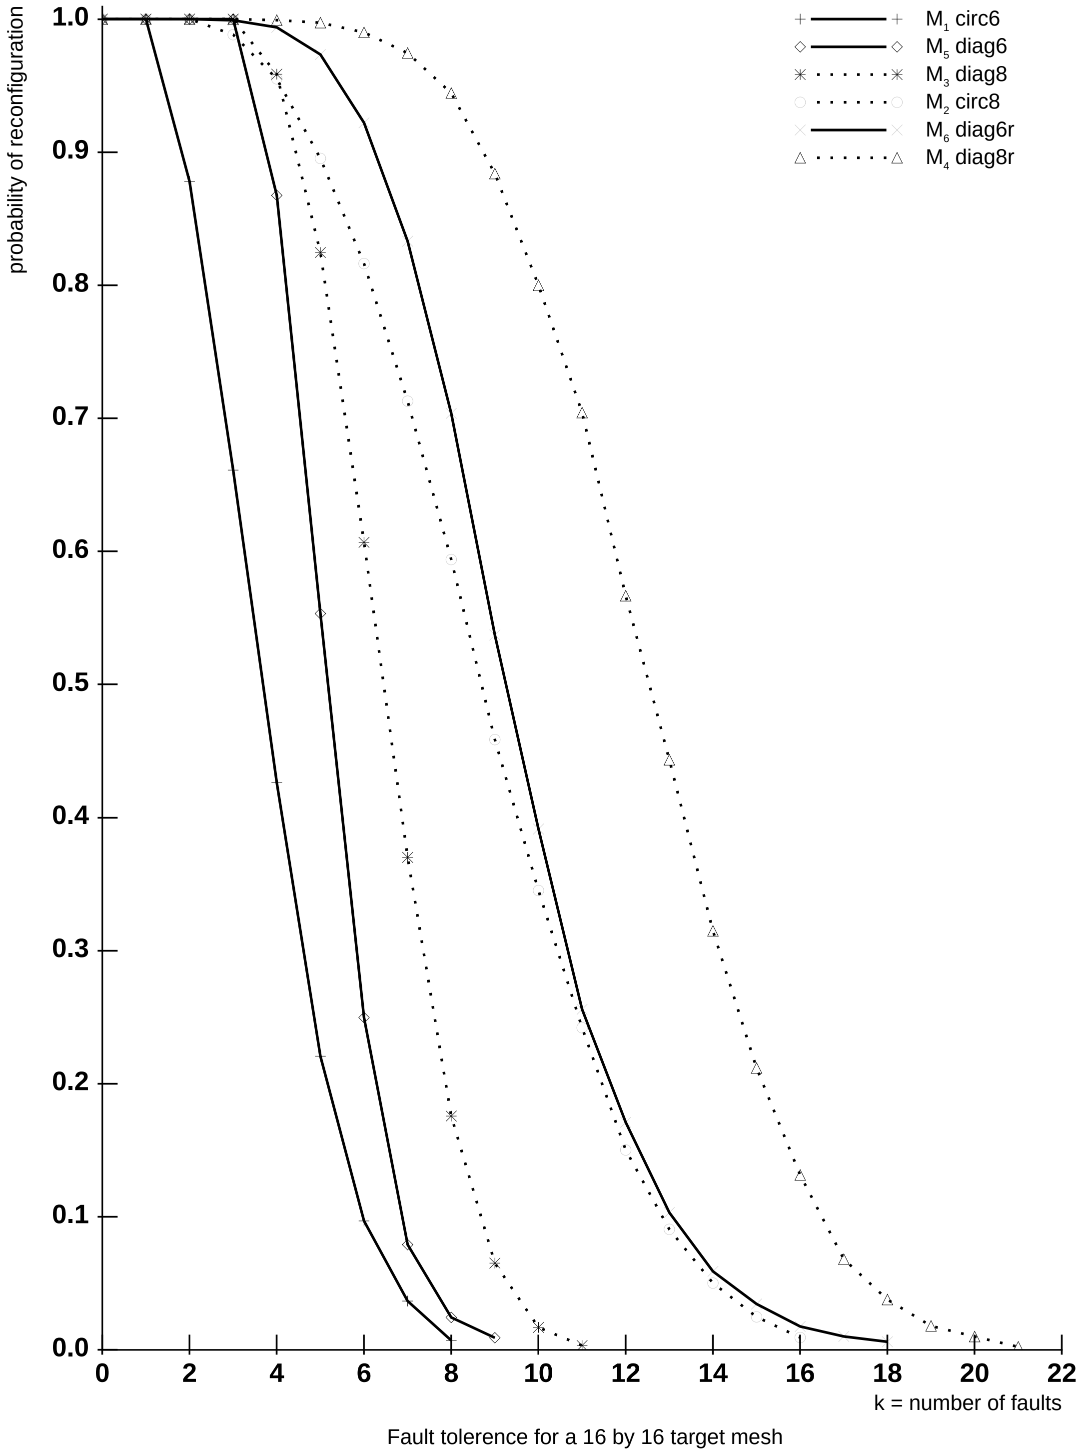

Figure 7: Simulation results of fault tolerance for a 16×16 target mesh.

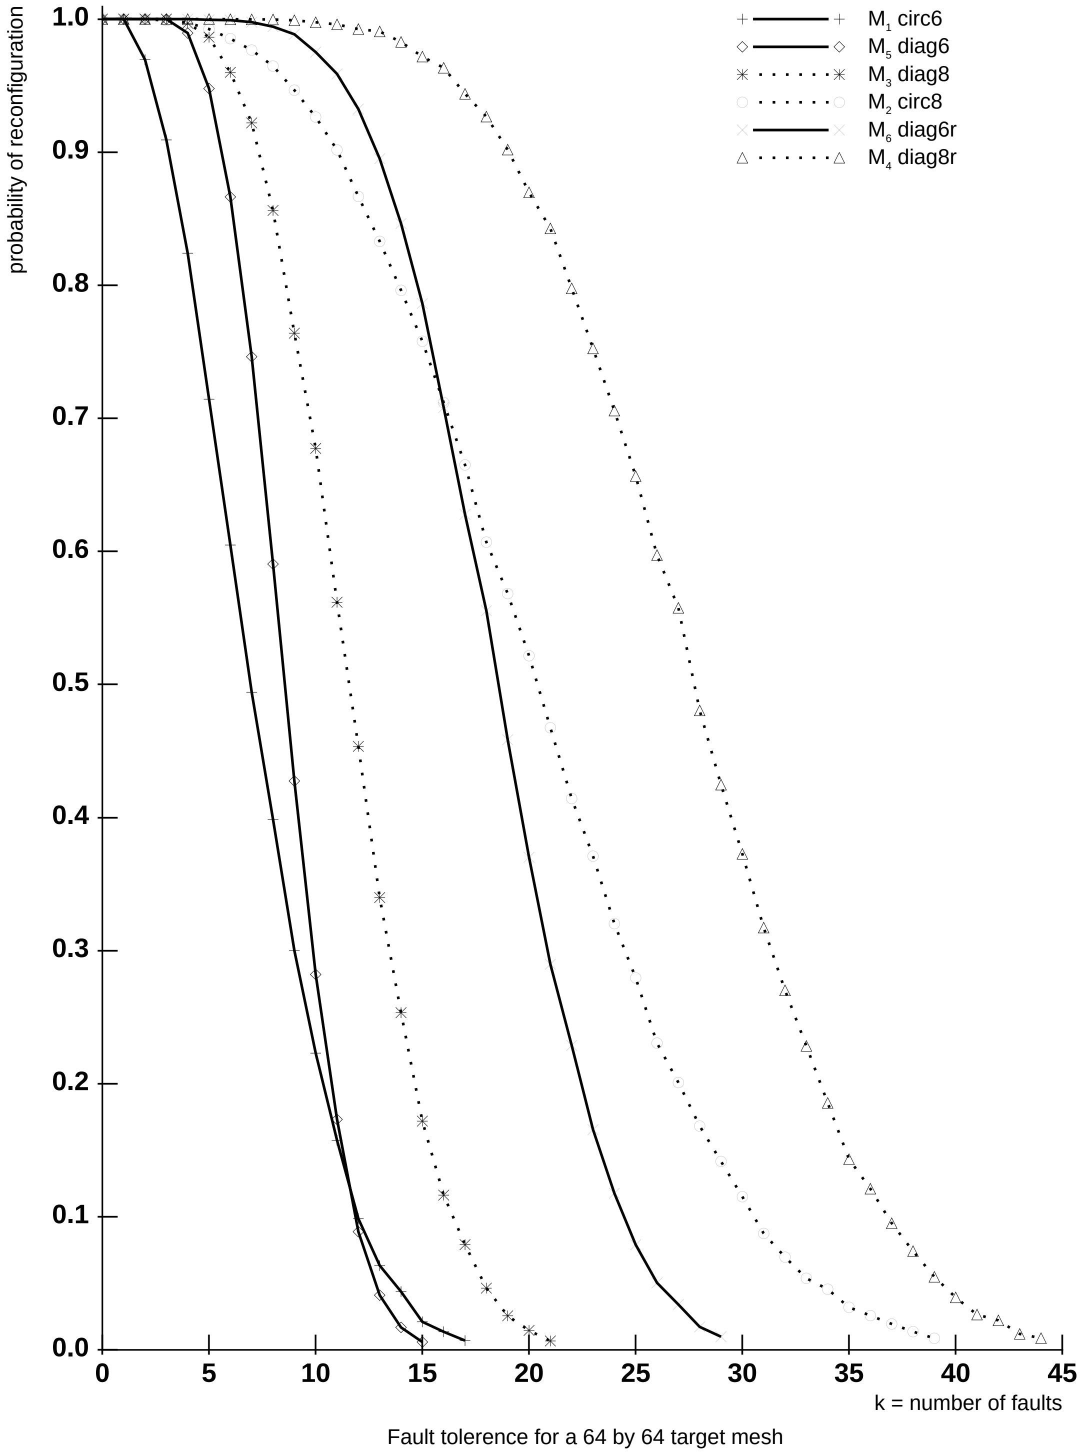

Figure 8: Simulation results of fault tolerance for a 64×64 target mesh.

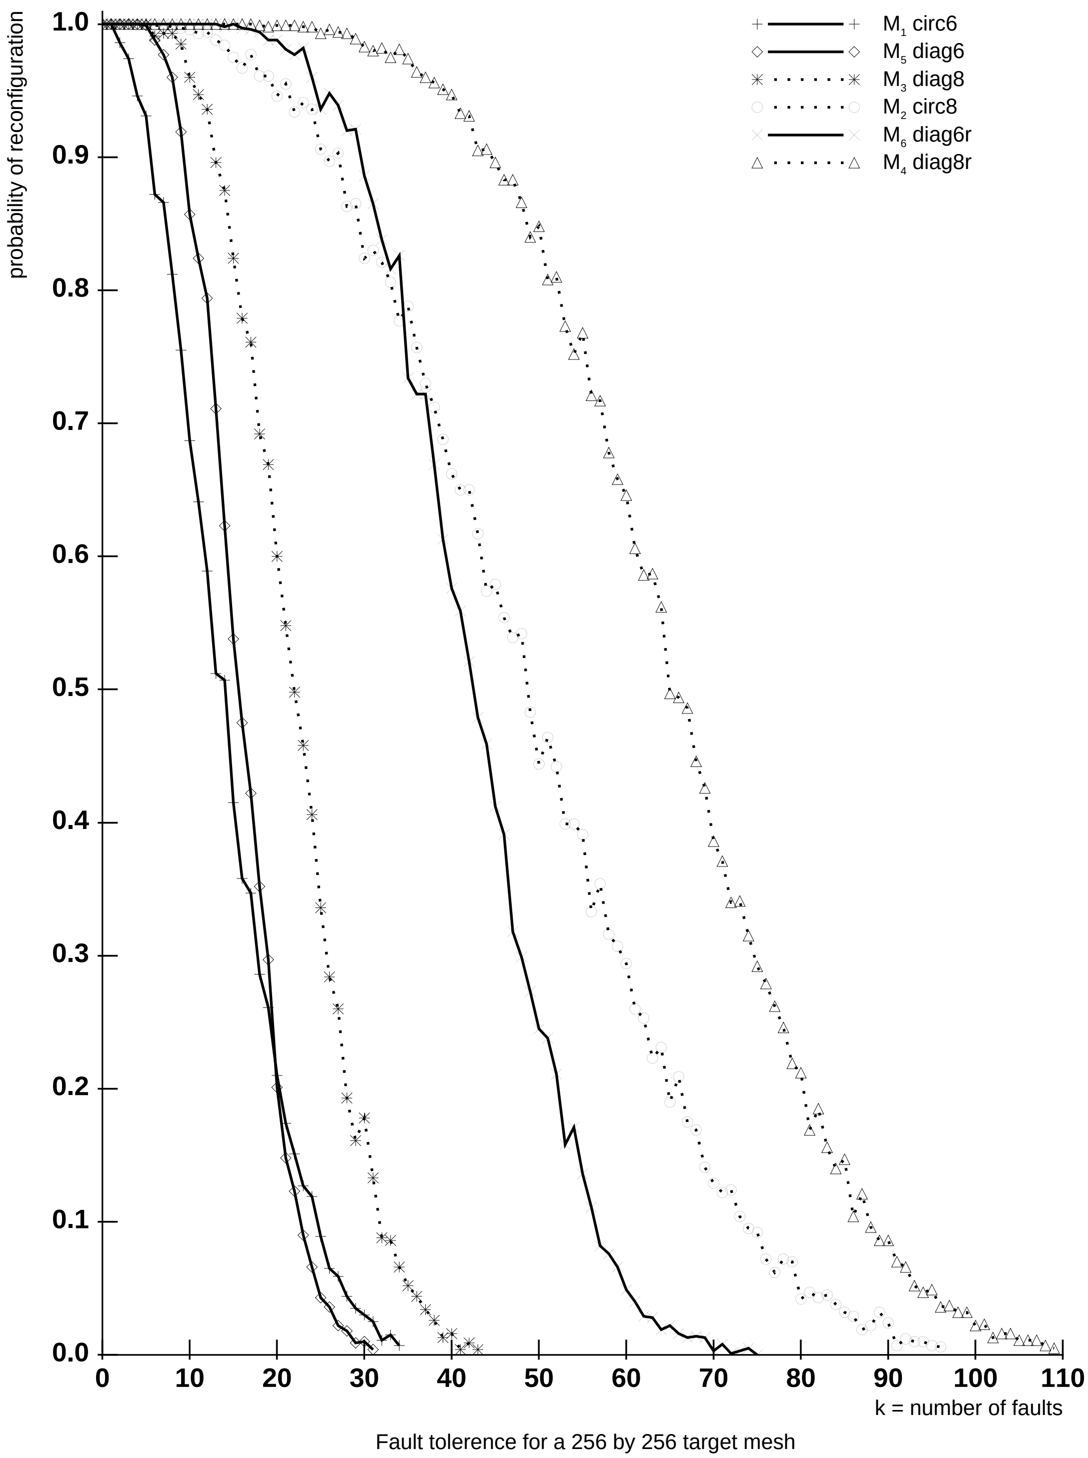

Figure 9: Simulation results of fault tolerance for a 256×256 target mesh.

References

[1] M. Ajtai, N. Alon, J. Bruck, R. Cypher, C.T. Ho, M. Noar and E. Szemerédi, *Fault Tolerant Graphs, Perfect Hash Functions and Disjoint Paths*, Proc. of 33rd Annual IEEE Symp. on Foundations of Computer Science, pp. 693–702, 1992.

[2] F. Annexstein, *Fault Tolerance in Hypercube-Derivative Networks*, Proceedings of the 1st Annual ACM Symposium on Parallel Algorithms and Architectures, pp. 179–188, 1989.

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

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

[5] C. Berge, *Graphs*, page 218, a Theorem attributed to Moon, North-Holland, 1985.

[6] G. Bilardi and F.P. Preparata, *Horizons of Parallel Computing*, Future Tendencies in Computer Science and Applied Mathematics, A. Bensoussan and J.P. Verjus, Eds., pp. 155–174, 1992.

[7] J. Bruck, R. Cypher and C.-T. Ho, *Fault-Tolerant de Bruijn and Shuffle-Exchange Networks*, IEEE Transactions on Parallel and Distributed Systems, vol. 5, no. 5, pp. 548–553, 1994.

[8] J. Bruck, R. Cypher and C.-T. Ho, *Tolerating Faults in a Mesh with a Row of Spare Nodes*, Theoretical Computer Science, vol. 128, pp. 241–252, 1994.

[9] J. Bruck, R. Cypher and C.-T. Ho, *Wildcard Dimensions, Coding Theory and Fault-Tolerant Meshes and Hypercubes*, IEEE Transactions on Computers (to appear). Also appeared in Proceedings of the 23rd International Symposium on Fault-Tolerant Computing, pp. 260–267, June 1993.

[10] J. Bruck, R. Cypher and C.-T. Ho, *Fault-Tolerant Meshes and Hypercubes with Minimal Numbers of Spares*, IEEE Transactions on Computers, vol. 42, no. 9, pp. 1089–1104, September 1993.

[11] J. Bruck, R. Cypher and D. Soroker, *Tolerating Faults in Hypercubes Using Subcube Partitioning*, IEEE Transactions on Computers, vol. 41, no. 5, pp. 599–605, 1992.

[12] R. Cypher, *Theoretical Aspects of VLSI Pin Limitations*, SIAM Journal on Computing, vol. 22, no. 2, pp. 356–378, 1993.

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

[14] S. Dutt and J. P. Hayes, *Designing Fault-Tolerant Systems Using Automorphisms*, Journal of Parallel and Distributed Computing, vol. 12, pp. 249–268, 1991.

[15] S. Dutt and J. P. Hayes, *Some Practical Issues in the Design of Fault-Tolerant Multiprocessors*, Proceedings of the 21st International Symposium on Fault-Tolerant Computing, pp. 292–299, June 1991.

[16] B. Elspas and J. Turner, "Graphs with circulant adjacency matrices", *Journal of Combinatorial Theory*, No. 9, 1970, pp. 297–307.

[17] J. Hastad, F. T. Leighton and M. Newman, *Fast Computations using Faulty Hypercubes*, Proceedings of 21st Annual ACM Symposium on Theory of Computing, pp. 251–284, 1989.

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

[19] 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. of 31st Annual IEEE Symp. on Foundations of Computer Science, pp. 285–296, October 1990.

[20] S.-Y. Kuo and W. K. Fuchs, *Efficient Spare Allocation for Reconfigurable Arrays*, IEEE Design and Test, pp. 24–31, February 1987.

[21] F. T. Leighton, *Introduction to Parallel Algorithms and Architectures: Arrays, Trees, Hypercubes.* Morgan Kaufmann, San Mateo, CA, 1992.

[22] 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, May 1985.

[23] T. Leighton, B. Maggs and R. Sitaraman, *On the Fault Tolerance of Some Popular Bounded-Degree Networks*, Proc. of 33rd Annual IEEE Symp. on Foundations of Computer Science, pp. 542–552, 1992.

[24] M. Paoli, W. W. Wong and C. K. Wong, *Minimum k-Hamiltonian Graphs, II,* J. of Graph Theory, Vol. 10, pp. 79–95, 1986.

[25] 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, October 1983.

[26] 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, April 1990.

[27] C. Thompson, *A Complexity Theory for VLSI*, Ph.D. Thesis, Dept. of Computer Science, Carnegie–Mellon University, Pittsburgh, PA, 1980.

[28] W. W. Wong and C. K. Wong, *Minimum k-Hamiltonian Graphs*, J. of Graph Theory, Vol. 8, pp. 155-165, 1984.