Short Notes

Operationally Enhanced Folded Hypercubes

Jong Kim and Kang G. Shin

Abstract— **Recently, several variations of the hypercube have been proposed to enhance its performance and reliability. The folded hypercube is one of these variations, in which an extra link is added to each node providing a direct connection to the node located farthest from it. In this short note, we propose a new operation mode of the folded hypercube to enhance its performance and fault-tolerance. There are $\binom{n+1}{k}$ regular k-cubes within a folded hypercube of dimension n, denoted by FQ_n. We introduce another type of hypercube, called the *twisted hypercube*, to improve the performance and fault tolerance of the folded hypercube. The problems of finding a subcube of given size in an FQ_n and routing messages within the subcube are addressed for the proposed operation mode. The advantages of the proposed operation mode over the regular-hypercube operation mode are analyzed in terms of dependability and robustness. The proposed operation mode is shown to make significant improvements over the regular-hypercube operation mode in both dependability and robustness. Because the new operation mode can be applied to only an $(n-1)$-subcube level for a given FQ_n, we present a general form of folded hypercube, thus enhancing the availability of subcubes of any dimension $m < n$.**

Index Terms— **Distributed architectures; regular, folded and twisted hypercubes; communication diameter; average message distance; subcube fault tolerance/availability**

I. Introduction

Selection of an appropriate interconnection network is the key to the design of any distributed/multiprocessor system, because the speed of internode communication, rather than that of computation, is known to be the bottleneck in accomplishing speedup with multiple processors. Over the past two decades, an overwhelming number of interconnection networks have been reported in the literature. Examples include crossbars, multiple buses, multistage interconnection networks, and hypercubes, to name a few. Among these, the hypercube has received considerable attention due mainly to its rich topological properties. The hypercube is a regular structure, has a small diameter, and offers good connectivity with a relatively small node degree. Moreover, a number of other well-known topologies, such as rings, trees, and meshes, can be mapped onto the hypercube [1].

The recent surge of interest in the hypercube has also led to the development of several variations of the hypercube, mainly to enhance the performance and fault tolerance of the original hypercube [2]–[5]. Preparata [2] introduced cube-connected cycles to provide a fixed degree of connection for each node. Esfahanian [3] introduced the twisted hypercube to reduce the communication diameter. Meyer and Pradhan [4] discussed the concept of the folded hypercube as the topology that has a reduced communication diameter, and Latifi [5] analyzed its properties and performance. Among these variations, the folded hypercube is shown to have advantages over the others in subcube fault tolerance/availability; i.e., one can find, with a higher probability, subcubes of a given dimension in the folded hypercube than in the other hypercube variations in case nodes fail. The main objective of this short note is to analyze and enhance the performance and fault tolerance of the folded hypercube.

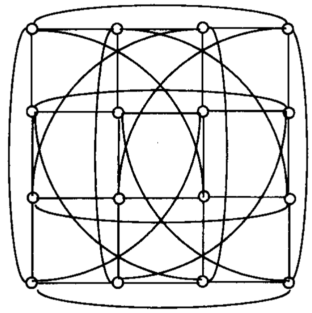

Fig. 1. A folded hypercube of dimension 4, FQ_4.

A folded hypercube [5]–[8] is the hypercube with an extra link connecting each node to its antidote node or to the node located farthest from it. Fig. 1 shows a 4-D folded hypercube, or FQ_4. Latifi and El-Amawy [5]–[7] showed that a folded hypercube has all the salient properties of the regular hypercube and some additional advantages. For example, the node degree of an FQ_n is $n+1$, and its communication diameter is $\lceil \frac{n}{2} \rceil$. An FQ_n has $\binom{n+1}{m}2^{n-m}$ different regular m-cubes, or Q_m's, whereas there are $\binom{n}{m}2^{n-m}$ Q_m's in a Q_n, where $n > m$. Additional attractive properties of the folded hypercube can be found in [5]–[8].

The operation of an FQ_n in the presence of faulty components was analyzed by Latifi [7] under the assumption that an injured FQ_n "functions" or "works," as long as one can find a Q_m, $m < n$, that includes all faulty components. This implies that in order to function properly, an injured FQ_n must have at least one healthy Q_{n-1} in it. Because all subcubes under this operation mode are the regular hypercubes, identification of these subcubes and routing in each of them are straightforward. However, using only regular hypercubes does not exploit the full advantages of the folded hypercube, because, as mentioned earlier, it contains other hypercube variations.

We therefore present a new operation mode under which the folded hypercube can be fully exploited. First, we uncover new topological properties of the folded hypercube that have not been reported in the literature, including the existence of twisted hypercubes in a folded hypercube. Next we analyze the advantages of the folded hypercube over the regular hypercube in terms of the number of distinct subcubes, average message distance, and subcube availability. Finally, we propose a new operation mode to fully exploit the advantages of the folded hypercube. The subcube dependability and robustness under the proposed operation mode are analyzed. We also propose a folding method that enhances the availability of healthy subcubes.

Manuscript received June 11, 1992; revised May 7, 1994. This work was supported in part by the Office of Naval Research under Grants N00014-92-J-1080 and N00014-94-1-80.

J. Kim is with the Department of Computer Science and Engineering, Pohang University of Science and Technology, Pohang, Kyungpook 790-784, Korea.

K.G. Shin is with the Real-Time Computing Laboratory, Division of Computer Science and Engineering, Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109-2122 USA.

IEEE Log Number 9405018.

1045-9219/94$04.00 © 1994 IEEE

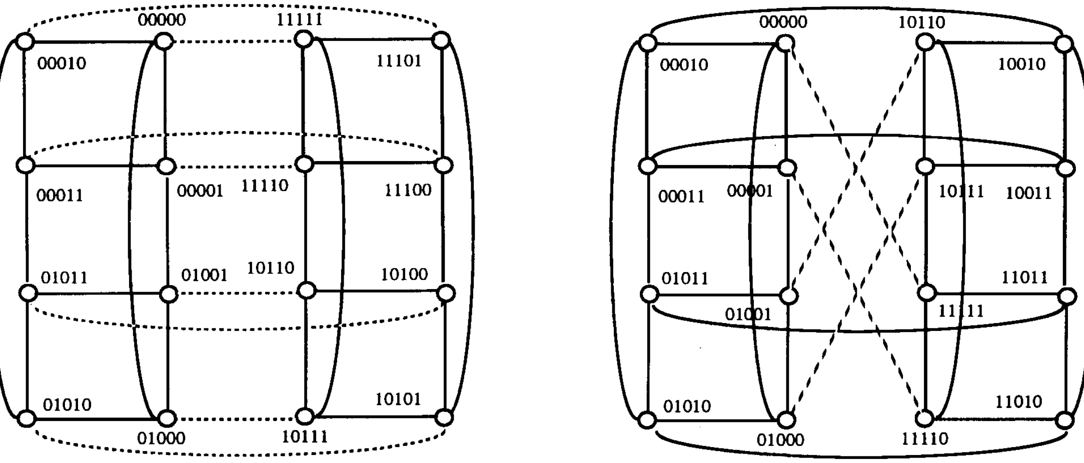

(a) (b)

Fig. 2. Two types of subcube of dimension $n-1$ in an FQ_n.

This short note is organized as follows. Section II introduces the topological properties of the folded hypercube. We analyze the advantages of the folded hypercube over the regular hypercube. The new operation mode is presented in Section III. The advantages of the proposed operation mode are comparatively analyzed in Section IV. A folding method to enhance the availability of healthy subcubes is discussed in Section V. The short note concludes with Section VI.

II. Topological Properties

For completeness, we begin with the list of useful properties of the folded hypercube that have already been reported in the literature.

Lemma 1: The node degree and the diameter of an FQ_n are $n+1$ and $\lceil n/2 \rceil$, respectively [5]. □

Lemma 2: The number of regular m-cubes in an FQ_n is $\binom{n+1}{m}2^{n-m}$, and that in a Q_n is $\binom{n}{m}2^{n-m}$, where $0 \le m \le n$ [6], [7]. □

Lemma 3: The problem of identifying an m-cube in an FQ_n can be reduced to that of identifying an m-cube in a Q_{n+1}, where $0 \le m \le n$ [6]. □

One attractive property of the folded hypercube that has not yet been reported in the literature is the availability of twisted hypercube. As can be seen in Fig. 2, we can observe that there are two different types of subcube in an FQ_n. The first type is a regular subcube. A regular subcube, Q_m, $1 < m < n$, can be formed by combining two neighboring $(m-1)$-cubes or two antidote $(m-1)$-cubes. For the latter case, every node in the subcube includes a link to its antidote node. The second type is the twisted subcube. One-half of the nodes in the twisted subcube include links to their antidote nodes, and the remaining half do not. Since there are many ways of generating twisted subcubes from a regular hypercube by cross-connecting any two nodes [3], the symbol TQ is used to denote the twisted hypercube induced by a folded hypercube. Also, we call the link between antidote nodes a *twisted* link, and those links that are not twisted we call *straight* links.

An n-dimensional twisted hypercube TQ_n (Fig. 2) is constructed by connecting two Q_{n-1}'s. One-half of the nodes in one Q_{n-1} have straight (connection) links to the nodes in the other Q_{n-1}, and the other half have twisted (connection) links. In Fig. 2, the dotted line represents a straight link, and the solid line represents a twisted link. Note that a TQ_n is still a regular structure. For example, if we turn one Q_{n-1} in the figure upside down, the half that used to have twisted links will have straight links, and the half that used to have straight links will have twisted links. The following lemma characterizes the communication diameter and the average message distance of a TQ_n.

Lemma 4: The communication diameter of TQ_n is $n-1$, and its average message distance is $\frac{(4n-1)2^{n-3}}{2^n-1}$.

Proof: Since each node is connected to its antidote node via a twisted link, the maximum distance between a pair of nodes in a TQ_n is less than n. Since there are two regular $(n-1)$-cubes in a TQ_n, the maximum internode distance or the communication diameter cannot be smaller than $n-1$.

A TQ_n can be partitioned into four disjoint groups/cubes (either regular or twisted) of dimension $n-2$, labeled A, B, C, and D such that B and D are regular cubes neighboring A, and B and C are antidote pairs (see Fig. 3). Since the TQ_n is a regular structure, without loss of generality, one can derive the average message distance with respect to a node in group A. Specifically, let us assume that the source node is located in the upper half of group A if group A is divided into upper and lower halves. The summation of message distances to other nodes in group A is $(n-2)2^{n-3}$, because group A is a regular $(n-2)$-cube [9]. The summation of message distances to the nodes in B and D is $2((n-1)2^{n-2}-(n-2)2^{n-3})$, because the combination of A and B or the combination of A and D is a Q_{n-1}.

The shortest path from the source node to a node in the upper half of group C is $A \to D \to C$ through straight links. Similarly, the shortest path to a node in the lower half of group C is $A \to B \to C$ through straight and twisted links. The summation of message distances to the nodes in group C is the summation of message distances to the nodes in the upper half of groups B and D with the addition of extra hops from B or D to C. This summation is $2((n-1)2^{n-4}+2^{n-3})$. By adding up all message distances, we get $(4n-1)2^{n-3}$. By dividing it with the total number of nodes, we can get the average message distance as given. □

Lemma 5: A twisted subcube exists in an FQ_n only if the size of this subcube is $n-1$. For other sizes, only regular subcubes can exist.

Proof: Let m be the dimension of a subcube in an FQ_n. To form a twisted cube of size m, two antidote subcubes of dimension $m-2$ connected with twisted links should be combined with two regular $(m-2)$-cubes connected with straight links. Assume that one subcube from the antidote pair is combined with a subcube from the pair of regular subcubes. Then the distance between the other subcube from the antidote pair and the other subcube from the pair of regular

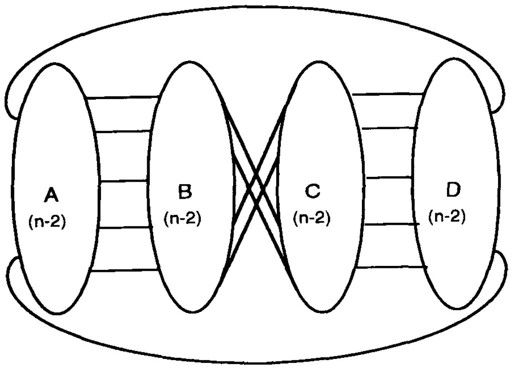

Fig. 3. The structure of a twisted hypercube, TQ_n.

subcubes cannot be less than 2 if the subcube size is less than $n-3$. Hence, there is no twisted cube in the FQ_n if $m < n-1$. □

As discussed in Lemma 4, the communication diameter of TQ_n is $n-1$. The gain in communication diameter for a TQ_n over a Q_n is nominal as the system size, n, increases. Based on the result of Lemma 4, we can also compute the gain in average message distance.

Theorem 1: The gain in average message distance for a TQ_n over a Q_n is given as:

$$\frac{(4n-1)2^{n-3}/2^n - 1}{n2^{n-1}/2^n - 1} = \frac{(4n-1)2^{n-3}}{n2^{n-1}} = \frac{4n-1}{4n}. \quad (2.1)$$

Proof: The average message distance in a Q_n is $n2^{n-1}/2^n - 1$ [9]. The average message distance in a TQ_n follows directly from Lemma 4. □

The result of Theorem 1 is very indicative. Since the communication-distance advantage of a TQ_n is almost negligible as the size of system increases, it is better not to use this advantage if it complicates the operation.

III. Operation Mode

In this section, we present two operation modes of the folded hypercube that can use its subcube fault tolerance capability. First, we summarize the operation mode discussed in [7], in which the system is assumed to *work* or *function* as long as one can find a healthy Q_{n-1} in an injured FQ_n. Next we propose a new operation mode that uses not only regular subcubes but also twisted subcubes.

A. Regular-Subcube Operation Mode

A regular $(n-1)$-cube in an FQ_n can be formed in two different ways. First, two neighboring Q_{n-2}'s form a Q_{n-1}. This is the same as the case of finding an $(n-1)$-cube from a Q_n. Second, two antidote Q_{n-2}'s form an $(n-1)$-cube, which is not possible in regular hypercubes. When the system operates in this mode, the two most important issues are the identification/recognition of a Q_{n-1} and the routing within the identified Q_{n-1}. Because these issues for the folded hypercube are different from those for the regular hypercube, they are reviewed briefly as follows.

Identification of an m-Cube: As stated in Lemma 3, the problem of identifying an m-cube in an FQ_n is the same as that of identifying an m-cube in a Q_{n+1}. A node in an FQ_n can be represented by two nodes in a Q_{n+1}. For example, node 000 in an FQ_3 corresponds to nodes 0000 and 1111 in a Q_4. If node 000 in an FQ_3 failed, then nodes 0000 and 1111 in a Q_4 are marked as failed. If we can find an m-cube in a Q_4, then there is an m-cube in an FQ_3, where $m < 4$.

Routing in a Subcube [5]: Since the routing method for a regular hypercube is well established and well understood, what we need is to show how to adapt this routing method to the subcubes of a folded hypercube. For example, consider the deadlock-free wormhole routing in a subcube [10]. If the links connecting antidote nodes are the ones in the lowest dimensional direction, one can easily apply e-cube routing [10]. In case the subcube is formed using the link to the antidote node, messages must traverse the antidote link first if they need to traverse it as well as other dimensional links. For the other dimensional directions, messages can follow the rule of e-cube routing.

B. The New Operation Mode

The proposed new operation mode assumes that the system is *operational* as long as one can find any kind of $(n-1)$-cube that is either regular or twisted. The inclusion of twisted subcubes complicates the operation mode, because a twisted subcube has different characteristics as compared to a regular hypercube. For example, regular and twisted hypercubes have different communication diameters and average message distances. However, we do not use the communication-diameter advantage of the twisted hypercube in order to simplify the operation mode. Similarly to the regular-subcube operation mode, we consider only the identification of a given size subcube and the message routing within the subcube.

Identification of an $(n-1)$-Cube: The following lemma and theorem show results on how to form and identify an $(n-1)$-cube in an FQ_n.

Lemma 6: Any two disjoint $(n-2)$-cubes in a Q_n can form an $(n-1)$-cube (regular or twisted) in an FQ_n that is either of the two types of subcube in Fig. 2.

Proof: Choose any two disjoint $(n-2)$-cubes in a Q_n. There are only three possibilities of connecting them. First, these two $(n-2)$-cubes are neighbors of each other, so that they naturally form a regular $(n-1)$-cube. Second, these two cubes are an antidote pair, so that they can form a Q_{n-1} in an FQ_n. Last, these two cubes are partially connected in a Q_n. Let A and B denote two disjoint $(n-2)$-cubes. When the addresses of cubes A and B are compared, there should be three bit positions that do not match. At one of these three positions, one cube should have 0, and the other 1, in order to be disjoint with each other. At the other two positions, one cube has X and $\overline{X}$, and the other has $\overline{X}$ and X. One can easily see that in a Q_n, one-half of the nodes in A have connections to one-half of the nodes in B, and each node in the other half of A has an antidote node in the other half of B. Since there is a link between each antidote pair in an FQ_n, the subcube of the last case becomes a TQ_{n-1}. □

Theorem 2: The problem of identifying an $(n-1)$-cube in an FQ_n can be reduced to that of identifying two disjoint $(n-2)$-cubes in a Q_n. □

The problem of identifying two disjoint $(n-2)$-cubes in a Q_n can easily be solved using the information on faulty nodes. Fig. 4 shows an algorithm that identifies $(n-1)$-cubes. The R and a[i] in the algorithm represent the set of all faulty nodes and the ith bit value of element a in the set R, respectively.

This algorithm finds an $(n-1)$-cube in two different ways. First, the set R is searched for a Q_{n-1} that contains all faulty nodes. If such a Q_{n-1} does not exist, then the faulty nodes are divided into two groups, depending on their ith bit value in the next step. (These divided groups are denoted as U and L in the algorithm.) If we can find an $(n-2)$-cube from each group, then there exists a twisted or regular $(n-1)$-cube in an FQ_n. The search must be performed for all values of $i \leq n$.

Routing in a Subcube: The communication diameter of a TQ_n is $n-1$. However, this reduction of the communication diameter makes

```
Algorithm : Subcube-identification
begin
  /* check if there is a healthy (n-1)-cube */
  for any i such that (0 <= i < n)
      if a[i] = b[i]  for all elements a and b in R then
         there exists a healthy (n-1)-cube.  Exit.
  for any i such that (0 <= i < n)
      divide the set R into two groups, U and L,
      according to the bit value at position i.
      If   a[j] = b[j]   for any j such that i < j < n
                          and for all elements a and b in U
      and  c[k] = d[k]   for any k such that i < k < n
                          and for all elements c and d in L
         there exists a healthy (n-1)-cube.
end; /* of algorithm */
```

Fig. 4. The algorithm for identifying an $(n-1)$-cube.

```
Algorithm :  Message-routing
begin
  /* S is the source node and D is the destination node */
  if S and D in the same subcube (i.e., U or L)
    send a message to D using the e-cube routing
  else
     find the antidote node of D and let it be D'
     if S and D' in the same subcube
       send a message to S' through the antidote link
          and send it to D via S' using the e-cube routing
     else
        find the neighboring node of S in the opposite subcube
              and let it be S"
        send a message to S" through the neighboring link
          and send it to D via S" using the e-cube routing
end; /* of algorithm */
```

Fig. 5. The message routing algorithm.

only a negligible improvement in the average message distance, as shown in Theorem 1. That is, the average message distance gain of a TQ_n over a Q_n converges to 1 as n gets large. Moreover, we cannot implement e-cube routing in a TQ_n if every message takes a minimum-hop path. Hence, it is preferable not to take advantage of the shorter communication diameter of a TQ_n. One can then use e-cube routing within a twisted subcube.

Fig. 5 shows the routing algorithm that can be applied to both regular and twisted subcubes. The notation U and L represent two disjoint Q_{n-2}'s in an FQ_n. The algorithm is developed based on the fact that any node in one Q_{n-2} has only either the antidote node or the neighboring node in the other Q_{n-2}, but not both. Thus, we classify links between two disjoint Q_{n-2}'s into *antidote* and *neighboring* links. Antidote links are the ones between antidote nodes, and the other links are called neighboring links.

IV. Performance Analysis of the Proposed Operation Mode

The proposed operation mode is comparatively analyzed in terms of number of distinct subcubes, robustness, and dependability.

A. Number of Distinct Subcubes

Lemma 2 shows the number of regular subcubes in an FQ_n and that in a Q_n. From Lemma 6, we know that if there are two disjoint $(n-2)$-cubes in a Q_n, then we can form a subcube of dimension $(n-1)$ in an FQ_n. The resulting subcube could be either regular or twisted. The next lemma shows how many different pairs of disjoint $(n-2)$-cubes exist in a Q_n.

Lemma 7: The total number of both regular and twisted $(n-1)$-cubes in an FQ_n is $n(4n^2-11n+9)$.

Proof: The number of $(n-2)$-cubes in a Q_n is $\binom{n}{n-2}2^2$. Without loss of generality, we can choose one $(n-2)$-cube and assume its address to be $00X\cdots X$. We can find another disjoint $(n-2)$-cube from $4n-5$ different subcube locations, because there are $2(n-1)$ subcube locations in either $1XX\cdots X$ or $X1X\cdots X$, and one $(n-2)$-cube, whose address is $11X\cdots X$, is counted in both $(n-1)$-cubes. We have to divide the product of $\binom{n}{n-2}2^2$ and $4n-5$ by 2, because positions of two disjoint $(n-2)$-cubes are interchangeable. Also, we have to subtract the redundant count from the result of the above computation. Since a regular $(n-1)$-cube can be formed by combining two $(n-2)$-cubes in $n-1$ different ways, $n-2$ is the redundant count for each regular $(n-1)$-cube. Hence, the final equation becomes as follows:

$$\binom{n}{n-2}2^2[4(n-1)-1]/2-\left[\binom{n-1}{n-2}-1\right]\binom{n}{n-1}2. \quad \square \tag{4.1}$$

The above lemma shows how many $(n-1)$-cubes exist in an FQ_n. We also know that no twisted subcubes exist if the dimension of subcube gets smaller than $n-1$ by Lemma 5. Now we want to determine the gain of the folded hypercube over the regular hypercube in terms of the number of different subcubes available.

Theorem 3: Let $\mathrm{NQ}_{n,m}$ and $\mathrm{NFQ}_{n,m}$ be the number of different m-cubes in a regular as well as a folded n-cube, respectively. Then the gain in the number of m-cubes available is given by the following:

$$\begin{aligned}\frac{\mathrm{NFQ}_{n,m}}{\mathrm{NQ}_{n,m}} &= \frac{n+1}{n-m+1} \quad \text{if } m \le n-1 \\ &\qquad\qquad \text{(regular cubes only)} \\ &= \frac{4n^2-11n+9}{2} \quad \text{if } m = n-1 \\ &\qquad\qquad \text{(both regular and twisted cubes).}\end{aligned} \tag{4.2}$$

Proof: This follows from Lemmas 2 and 7. $\square$

B. Robustness

One of the major advantages of the hypercube is that each of its subcubes possesses the same properties of the original hypercube, except for the reduced cube size. Hence, a task that consists of several modules is usually mapped onto a subcube, rather than a set of nodes that do not form a subcube. The ability of finding a given size subcube in the presence of faulty nodes is thus an important measure for the fault tolerance of the hypercube and variations thereof. So, we want to analyze the availability of a subcube of dimension $m < n$ in both a FQ_n and a Q_n in the presence of faulty nodes as a measure of the system's fault tolerance. Note that the availability of an m-cube in an injured Q_n with x faulty nodes was studied in [11], but that in an injured FQ_n, to our best knowledge, has not been reported elsewhere.

Let $\Pr\{C_n = m \mid X_n = x\}$ denote the probability (called the *availability* of an m-cube) that there exists a healthy m-cube in an injured Q_n with x faulty nodes, where C_n and X_n are random variables representing the size of subcube and the number of faulty nodes in a Q_n, respectively. If we subtract the probability that there is only one fault-free Q_{n-2} from the probability that there exists at least one fault-free Q_{n-2} in a Q_n, then we get the probability that there are at least two disjoint fault-free Q_{n-2}'s in a Q_n. Thus, the probability that there exists an $(n-1)$-cube (either Q_{n-1} or TQ_{n-1}) in an FQ_n after the failure of k nodes is as follows:

$$\begin{aligned}&\Pr\{C_n \ge n-2 \mid X_n = k\} \\ &\quad - \Pr\{C_n = n-2; \mathrm{NC}_n = 1 \mid X_n = k\},\end{aligned} \tag{4.3}$$

where NC_n is a random variable representing the number of *disjoint* fault-free Q_{n-2} subcubes in a Q_n. The computation of the first term

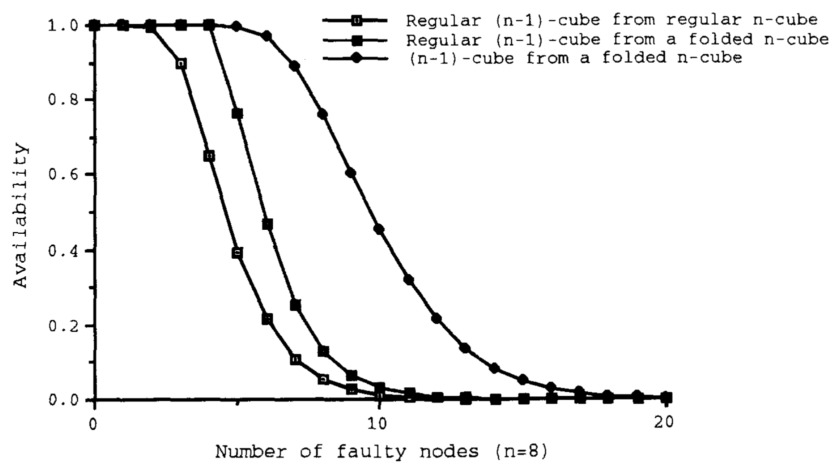

Fig. 6. Subcube availability vs. the number of faulty nodes.

was reported in [11]. The computation of the second term in the above expression is much more involved, so it is given in the appendix (for better readability).

Theorem 4: The gain in the availability of an $(n-1)$-cube (either a Q_{n-1} or a TQ_{n-1}) in an FQ_n over a Q_n in the presence of k faulty nodes is given in (4.4) at the bottom of this page. □

In Fig. 6, we have plotted the subcube availability as a function of the number of faulty nodes. The availability of a Q_{n-1} in an FQ_n is computed by using the result in [7]. As can be seen from Fig. 6, the availability of an $(n-1)$-cube (either Q_{n-1} or TQ_{n-1}) in an FQ_n is much larger than that in a Q_n. This gain in subcube availability/fault tolerance is very important to critical applications, because, as mentioned earlier, each application task usually needs to be assigned to a subcube, rather than to a set of nodes that do not form a subcube.

We define the robustness of a system as the average number of node failures that prevent us from finding an operational subcube, i.e., finding an $(n-1)$-cube from an n-cube system. The average number of faulty nodes that prevent us from finding a healthy Q_{n-1} in an FQ_n, denoted as $S_{avg}(n-1)$, is the summation of probabilities that all possible new faulty nodes do not block us from finding a healthy Q_{n-1}. This can be represented as follows:

$$S_{avg}(n-1) = \sum_{k=1}^{2^{n-1}} \sum_{i=0}^{n-1} \Pr\{F_n < n; k \mid F_n = i; k-1\} \times \Pr\{F_n = i \mid C_n = k-1\}, \quad (4.5)$$

where F_n is the random variable representing the smallest subcube that can include all faulty nodes in an FQ_n. The first term is the conditional probability that a new faulty node does not prevent us from finding a maximal healthy subcube, Q_{n-1}. The final equation becomes:

$$S_{\mathrm{avg}}(n-1) = \sum_{k=1}^{2^{n-1}} (1 - \Pr\{F_n = n \mid X_n = k\}). \quad (4.6)$$

The probability term $\Pr\{F_n = n \mid X_n = k\}$ was studied in [12].

TABLE I
COMPARISON OF ROBUSTNESS

n	$R_{avg}(n-1)$	$S_{avg}(n-1)$	$T_{avg}(n-1)$
6	4.85819	6.00272	8.01956
7	5.14338	6.37240	8.88077
8	5.36843	6.73785	9.51095
9	5.55327	6.93313	9.98389

Let $T_{\mathrm{avg}}(n-1)$ be the average number of node failures to block us from finding any $(n-1)$-cube (either Q_{n-1} or TQ_{n-1}) in an FQ_n. Then $T_{\mathrm{avg}}(n-1)$ is as follows:

$$T_{\mathrm{avg}}(n-1) = \sum_{k=1}^{2^{n-1}} (\Pr\{C_n \geq n-2 \mid X_n = k\} - \Pr\{C_n = n-2; NC_n = 1 \mid X_n = k\}). \quad (4.7)$$

Table I shows the robustness of both operation modes. $R_{\mathrm{avg}}(n-1)$ represents the average number of node failures that block us from finding an operational $(n-1)$-cube in a Q_n. Again, the proposed operation mode shows a *significant* improvement in robustness over the regular-subcube operation mode.

C. Dependability

Recall that the n-cube system is said to be *operational* if one can find a regular or twisted $(n-1)$-cube in the system. The most suitable measure for the dependability of the proposed operation mode and the regular-subcube operation mode is the mean time to failure (MTTF), because we are interested in the average operation period before failure. The MTTF of the regular-subcube operation mode is reported in [7]. To compute the MTTF of the proposed operation mode, we use the following reliability equation that Kim and Das

$$\frac{\Pr\{C_n \geq n-2 \mid X_n = k\} - \Pr\{C_n = n-2; NC_n = 1 \mid X_n = k\}}{\Pr\{C_n \geq n-1 \mid X_n = k\}}. \quad (4.4)$$

TABLE II
COMPARISON OF MTTF WITH NODE FAILURE RATE, $\lambda = 10^{-5}$ (in hr)

n	MTTF(Q_n)	MTTF(FQ_n) of R–mode	MTTF(FQ_n) T–mode
6	7861	9592	15139
7	4094	5091	8023
8	2117	2663	4191
9	1090	1363	2168

[11] developed:

$$\begin{aligned} R(t) = \sum_{k=0}^{2\cdot 2^{n-2}} & [\Pr\{C_n \geq n-2 \mid X_n = k\} \\ & - \Pr\{C_n = n-2; \mathrm{NC}_n = 1 \mid X_n = k\}] \\ & \times R_{xc}(t;k), \end{aligned} \tag{4.8}$$

where $R_{xc}(t;k)$ represents the reliability of completely connected machines at time t, when k out of 2^n nodes failed. The first term comes from (4.3). By integrating $R(t)$, we can get the MTTF of the proposed operation mode.

Table II shows MTTF's of the two operation modes. The first column shows the operation mode of the regular hypercube, and the second and third columns show the two operation modes of the folded hypercube. R-mode and T-mode represent the regular-subcube operation mode and the proposed operation mode, respectively. This table shows the superiority of T-mode over R-mode via its significant improvements in dependability. The comparison is made using a node failure rate of 10^{-5}. The data for R-mode is obtained from [7].

V. GENERALIZATION TO ENHANCED HYPERCUBE

Thus far, we have discussed the advantages of the folded hypercube and a new operation mode that includes the twisted hypercube. One disadvantage of the new operation mode is that it cannot be extended to subcubes of dimension smaller than $n-1$. When the system is allowed to degrade further (lower than an $(n-1)$-cube) or the dimension of subcube preferred by applications is lower than $n-1$, this limitation may negate the advantages of the proposed operation mode. We alleviate this problem by proposing a general form of the folded hypercube.

The enhanced hypercube, or a general form of folded hypercube (denoted as EQ_n), was introduced first by Tzeng and Wei [13]. They generated hypercube variations by adding an extra link between each pair of nodes that are j hops apart, where $j \leq n$. Let EQ_n^j denote the n-dimensional enhanced hypercube resulting from the addition of an extra link between each pair of nodes that are j hops apart. These hypercube variations are then compared with one another in terms of average message distance, traffic density, and diameter. The authors of [13] concluded that the hypercube with an extra link between each pair of nodes that are farthest apart (i.e., the folded hypercube) offers the best performance improvement. However, they did not consider the fault tolerance aspect of these enhanced hypercubes.

When fault-tolerance is important and the degradation below an $(n-1)$-cube is allowed, we want to determine which enhanced hypercube offers the best fault tolerance for a given subcube dimension m. This result can be used for the design of an enhanced hypercube for those critical applications that require a high degree of fault tolerance. The following lemma states how many regular m-cubes can be found in an enhanced hypercube, EQ_n.

Lemma 8: All enhanced hypercubes have the same number of regular subcubes.

Proof: Let the links of a node in different dimensional directions be labeled from 0 to $n-1$, and let the extra link be the nth dimensional link. Without loss of generality, we can assume that an extra link is added to connect each pair of nodes that are k hops apart. In other words, an EQ_n^k is formed by replacing each node of a Q_{n-k} with an FQ_k. Suppose a Q_m is selected from the EQ_n^k. Then one can get a Q_m by selecting an s-cube from the Q_{n-k} and the same regular $(m-s)$-cube from each FQ_k that corresponds to a node of the selected s-cube. Since k and s are arbitrarily chosen, selection of m different dimensional links from EQ_n can always guarantee the formation of a Q_m. Hence, the number of Q_m's generated by the addition of extra links does not change with the dimension of their placement. □

By adding an extra link between each pair of nodes that are $m+1$ hops apart to form an enhanced hypercube, we can maximize the ability of finding m-cubes (either regular or twisted) in this enhanced hypercube. The combined number of Q_m's and TQ_m's in this enhanced hypercube is given in the following lemma.

Lemma 9: Let $\mathrm{NE}_{n,m}^{m+1}$ be the combined number of Q_m's and TQ_m's in an EQ_n^{m+1}. Then we have:

$$\begin{aligned} \mathrm{NE}_{n,m}^{m+1} &= 2^{n-m-1}\left[\mathrm{NFQ}_{m+1,m} - 2\binom{m+2}{m}\right] \\ &\quad + \binom{n+1}{m} 2^{n-m}. \end{aligned} \tag{5.1}$$

Proof: An enhanced hypercube EQ_n^{m+1} can be formed by replacing each node of an $(n-m-1)$-cube with an FQ_{m+1}. The term 2^{n-m-1} represents the number of disjoint FQ_{m+1}'s in an EQ_n^{m+1}. The term $\mathrm{NFQ}_{m+1,m}$ represents the combined number of Q_m's and TQ_m's in an FQ_{m+1}. The second term in the square bracket represents the number of Q_m's in an FQ_{m+1}. The subtraction of second term from the first term in the square bracket results in the number of twisted hypercubes in an FQ_{m+1}. The final term $\binom{n+1}{m}2^{n-m}$ is the total number of Q_m's in an FQ_n. Since the number of Q_m's in an EQ_n^{m+1} is equal to that of an FQ_n, we can use the number of Q_m's in an FQ_n for the number of Q_m's in an EQ_n^{m+1}. Hence, the final result (5.1) is the total combined number of Q_m's and TQ_m's in an EQ_n^{m+1}. □

It is obvious that the combined number of Q_m's and TQ_m's in the enhanced hypercube EQ_n^{m+1} is significantly larger than the number of Q_m's in the EQ_n^{m+1}. When an m-cube is preferred for certain applications, one must consider the enhanced hypercube EQ_n^{m+1} as a candidate structure and must include use of twisted subcubes.

VI. CONCLUSION

In this short note, we proposed a new operation mode of the folded hypercube based on twisted hypercubes. We have found that there is a twisted subcube in an FQ_n and that the size of available twisted subcube is $n-1$. The algorithms to solve the problems of identifying subcubes and routing in a subcube are presented. Also, we have shown that the new operation mode provides significantly better fault tolerance and robustness than the operation of a Q_n and the regular subcube operation of an FQ_n. Finally, we analyzed the enhanced hypercube that is a general form of folded hypercube in terms of the number of distinct subcubes available. The result of this analysis can be used to design or choose application-specific fault-tolerant hypercube variants.

The inclusion of twisted hypercubes in the operation mode enhances the value of the folded hypercube to a great extent. One disadvantage of the proposed operation mode is the added complexity in the routing algorithm. However, the fault tolerance advantage is significant enough to offset this disadvantage when critical applications must be handled by the system.

Appendix
Computation of (4.3)

Let us assume that there are k faulty nodes in an injured FQ_n. To have only one healthy Q_{n-2}, k faulty nodes should be dispersed to prevent the finding of another disjoint Q_{n-2}. Without loss of generality, we can assume that one healthy Q_{n-2} is $\mathtt{00X\cdots X}$. Let A, B, and C denote subcubes $\mathtt{01X\cdots X}$, $\mathtt{10X\cdots X}$, and $\mathtt{11X\cdots X}$, respectively. If k faulty nodes are dispersed into three Q_{n-2}'s and they block us from finding another healthy disjoint Q_{n-2}, then there will be only one Q_{n-2} in the system. Assume that there are i faulty nodes in A, j faulty nodes in B, and $k-i-j$ faulty nodes in C. The probability that no Q_{n-2} can be found from the combination of A, B, and C is as follows:

$$\begin{aligned}\Pr{}_1 = &\Pr\{C_{n-1} < n-2 \mid X_{n-1} = k-i\}\\ &\times \Pr\{C_{n-1} < n-2 \mid X_{n-1} = k-j\}. \quad \text{(A1)}\end{aligned}$$

The first term represents the case that no Q_{n-2} can be found from the combination of B and C. Similarly, the second term represents the case that no Q_{n-2} can be found from the combination of A and C.

$\Pr_1$ in (A1) still includes some cases of two disjoint Q_{n-2}'s. Let us consider the case that there are healthy Q_{n-3}'s in A and B, and that they are located at antidote positions. For example, the healthy Q_{n-3} in A is $\mathtt{010X\cdots X}$ and the healthy Q_{n-3} in B is $\mathtt{101X\cdots X}$. When these two Q_{n-3}'s are merged with a subcube $\mathtt{00X\cdots X}$, they form two disjoint Q_{n-2}'s, even though such a case is regarded as having a single Q_{n-2} in the Pr_1 if no Q_{n-2} can be found from A, B, and C. To find the probability of this kind of cases, let us assume that the faulty nodes in A can be included in an s-cube and that the faulty nodes in B can be included in a t-cube. If s and t overlap, i.e., if any node in one fault group has an antidote node in the other fault group, then there are no two disjoint Q_{n-2}'s. Mathematically, we have the following relation:

$$\begin{aligned}\Pr{}_2 = &\sum_{u=0}^{\min(n-2-s,t)} \frac{\binom{s}{u}2^{s-u}\binom{n-2-s}{t-u}}{\binom{n-2}{t}2^{n-2-t}} F(n-2,s,i)\\ &\times F(n-2,t,j), \quad \text{(A2)}\end{aligned}$$

where $F(d,p,l)$ is $\Pr\{F_d = p \mid X_d = l\}/\Pr\{F_d \le d-1 \mid X_d = l\}$. The summation represents the size of the overlapping subcube between two faulty subcubes s and t. The combinatorial term represents the probability that the dispersed faulty nodes prevent the finding of two disjoint Q_{n-2}'s. The last two terms provide the probability that the smallest subcube sizes that can include all faulty nodes are s and t in A and B, respectively. $1-\Pr_2$ is the probability that there are two disjoint Q_{n-2}'s. Hence, with the probability that there are healthy Q_{n-3}'s in A nd B, the probability that we can find two disjoint Q_{n-2}'s is as follows:

$$\begin{aligned}\Pr{}_3 = &\sum_{s=0}^{n-3}\sum_{t=0}^{n-3} \Pr\{C_{n-2} = n-3 \mid X_{n-2} = i\}\\ &\times \Pr\{C_{n-2} = n-3 \mid X_{n-2} = j\}(1-Pr_2). \quad \text{(A3)}\end{aligned}$$

The summation represents all possible sizes of the smallest faulty subcubes in A and B.

Now we have to find the probability that i, j, and $k-i-j$ faulty nodes are dispersed in A, B, and C, respectively. Considering this probability, we get the final equation:

$$\begin{aligned}&\Pr\{C_n = n-2; NC_n = 1 \mid X_n = k\}\\ &= \sum_{i=1}^{k-2}\sum_{j=1}^{k-1} \Pr\{C_n \ge n-2 \mid X_n = k\}\\ &\quad\times \frac{\binom{2^{n-2}}{i}\binom{2^{n-2}}{j}\binom{2^{n-2}}{k-i-j}}{\binom{2^n-2^{n-2}}{k}} \Pr_1\left(1-\Pr_3\right). \quad \text{(A4)}\end{aligned}$$

References

[1] Y. Saad and M. H. Schultz, "Topological properties of hypercube," *IEEE Trans. Comput.*, vol. C-37, pp. 867–872, July 1988.
[2] F. P. Preparata and J. Vuillemin, "The cube-connected cycles: A versatile network for parallel computation," *Commun. ACM*, vol. 24, pp. 300–309, May 1981.
[3] A. Esfahanian, L.M. Ni, and B.E. Sagan, "On enhancing hypercube multiprocessors," in *Proc. 1988 ICPP*, 1988, pp. 86–89.
[4] F.J. Meyer and D.K. Pradhan, "Flip-trees: Fault tolerant graphs with wide containers," *IEEE Trans. Comput.* , vol. C-37, pp. 472–478, Apr. 1988.
[5] S. Latifi and A. El-Amawy, "On folded hypercubes," in *Proc. 1989 ICPP*, 1989, pp. 180–187.
[6] S. Latifi, "The efficiency of the folded hypercube in subcube allocation," in *Proc. 1990 ICPP*, vol. I, pp. 218–225, Aug. 1990.
[7] ——, "Fault-tolerant hypercube multiprocessors," *IEEE Trans. Reliability*, vol. 39, pp. 361–367, Aug. 1990.
[8] C.T. Ho, "Full bandwidth communications on folded hypercubes," in *Proc. 1990 ICPP*, vol. I, pp. 276–280, Aug. 1990.
[9] L.N. Bhuyan and D.P. Agrawal, "Generalized hypercube and hyperbus structures for a computer network," *IEEE Trans. Comput.*, 1984, pp. 323–333.
[10] W.J. Dally and C.L. Seitz, "Deadlock-free message routing in multiprocessor interconnection networks," *IEEE Trans. Comput.*, vol. C-36, pp. 547–553, May 1987.
[11] K. Jim and C.R. Das, "On subcube dependability in a hypercube," in *Proc. ACM SIGMETRICS'91*, 1991, pp. 111–119.
[12] S. Abraham and K. Padmanabhan, "Reliability of the hypercube," in *Proc. 1988 ICPP*, vol. I, pp. 90–94, Aug. 1988.
[13] N.F. Tzeng and S. Wei, "Enhanced hypercubes," *IEEE Trans. Comput.*, vol. 40, pp. 284–294, Mar. 1991.