Correspondence

The Twisted N-Cube with Application to Multiprocessing

Abdol-Hossein Esfahanian, Lionel M. Ni, and Bruce E. Sagan

Abstract—**We show that by exchanging any two *independent edges* in any shortest *cycle* of the n-cube ($n \geq 3$), its *diameter* decreases by one unit. This leads us to define a new class of *n-regular* graphs, denoted TQ_n, with 2^n vertices and diameter $n - 1$, which has the ($n - 1$)-cube as subgraph. Other properties of TQ_n such as *connectivity* and the lengths of the disjoints paths are also investigated. Moreover, we show that the *complete binary tree* on $2^n - 1$ vertices, which is not a subgraph of the n-cube, is a subgraph of TQ_n. Finally, we discuss how these results can be used to enhance existing hypercube multiprocessors.**

Index Terms—**Binary tree, graph theory, hypercube multiprocessors, message routing, N-cube, pattern embedding.**

I. Introduction

The possibility of interconnecting a number of processors together to solve very large problems in scientific computations has been extensively considered in the past [19]. *Distributed-memory multiprocessor systems* (or *multicomputers*) have proven to be one of the most straightforward and the least expensive methods to build such arrays with hundreds or even thousands of processors [28]. In such networks, each processor has its own memory, and message passing is the means of information exchange between processors. It is well known that the topology of the interconnection network plays a significant role in system performance [10], [20], [24], [27].

The *hypercube* interconnection scheme is the most popular topology being used in distributed-memory multiprocessors. Since the successful design of the first working hypercube computer, the 64-node Cosmic Cube, at Caltech [28], a number of hypercube multiprocessors have become commercially available. Among them, the NCUBE's NCUBE/10 can have up to 1024 processors, the Intel's iPSC series can have up to 128 processors, the Ametek's S-14 can have up to 256 processors, and the FPS's T series can have up to 2^{12} processors [12], [18], [26]. All these first generation hypercube multiprocessors adopt a packet switched store-and-forward mechanism for handling information exchange.

An *n-dimensional* hypercube multiprocessor consists of $N = 2^n$ processors interconnected as follows. Each processor is labeled by a different n-bit binary number $(b_{n-1}b_{n-2}\cdots b_1b_0)$. Two processors are connected by a full duplex link if and only if their binary labels differ in exactly one bit position. The popularity of hypercube multiprocessors is due to their underlying topology which is known as the *n-cube graph* Q_n. The n-cube graph has been the subject of many research projects in recent years, mainly because of the availability of hypercube multiprocessors [5], [9], [22], [27], [29]. As a result, many properties of the n-cube have been discovered [2], [4], [8], [11], [16], [21], [23].

The rest of this paper is organized as follows. Our notation and terminology are given in the next section. A new interconnection topology, denoted TQ_n, which is based on a simple modification of the n-cube is given in Section III. We will show in Sections IV–VII that TQ_n has certain topological advantages over Q_n. In particular, it is shown that the *diameter* of TQ_n is one less than that of Q_n, and its *vertex-connectivity* is the same as that of Q_n. It is known that the *complete binary tree* on $2^n - 1$ vertices, T_n, is not a subgraph of Q_n [27]. However, T_{n-1} is contained in Q_n [4]. We prove that TQ_n has the complete binary tree T_n as subgraph. Other subgraphs of TQ_n are also identified. Routing in TQ_n is described in Section VIII. Finally, practical implications of our results are given in Section IX.

Manuscript received December 11, 1987; revised November 15, 1988. L. Ni was supported in part by the DARPA ACMP project.

A.-H. Esfahanian and L. M. Ni are with the Department of Computer Science, Michigan State University, East Lansing, MI 48824.

B. E. Sagan is with the Department of Mathematics, Michigan State University, East Lansing, MI 48824.

IEEE Log Number 9037199.

II. Notation and Terminology

We will closely follow the graph theoretical terminology and notation of [15]; terms not defined here can be found in that book. Let $G(V, E)$ represent a *graph* with *point* or *vertex set* $V(G) = V$ and *edge set* $E(G) = E$. If an edge $e = uv \in E$, then vertices u and v are said to be *adjacent,* and the edge e is said to be *incident* to these vertices, and u and v are the *endpoints* of edge e. Two edges are said to be *independent* if they do not share an endpoint. For a vertex $v \in V, I(v)$ represents the set of all edges incident to v in G, and its cardinality $|I(v)|$ is the *degree deg*(v) of vertex v. We denote by $\delta(G)$ and $\Delta(G)$ the *minimum* and *maximum degrees,* respectively, of vertices of G. If $\delta(G) = \Delta(G) = k$, then G is said to be *k-regular.* For a set $X \subset E$ (or $X \subset V$), the notation $G - X$ represents the graph obtained by removing the edges (vertices) in X from G. The *vertex-connectivity,* $\kappa(G)$, of a graph G is the least cardinality $|X|$ of a set $X \subset V(G)$ such that $G - X$ is either *disconnected* or consists of a single vertex. Furthermore, a p-cycle is defined as a cycle on p vertices.

The *distance* $d(u, v)$ between two distinct vertices u and v is the length (in number of edges) of a shortest path between these vertices. The *diameter* $d(G)$ of graph G is then defined to be $d(G) = \max\{d(u, v) | u, v \in V\}$. If H and G are graphs, then H is *isomorphic* to a *subgraph* of G if there is a one-to-one function $f : V(H) \to V(G)$ such that each edge $uv \in E(H)$ is carried to an edge $f(u)f(v) \in E(G)$. By an abuse of language we will often merely say that H is a subgraph of G (where in reality it is $f(H)$ which is a subgraph of G) and will write $H \subseteq G$.

Two specific graphs with which we will be concerned are complete binary trees and n-cubes. As indicated before, T_n will represent the complete binary tree on 2^{n-1} vertices. The *root* of T_n is the unique vertex whose degree is 2. For an n-cube, Q_n, we have $\delta(Q_n) = \Delta(Q_n) = n$, $d(Q_n) = n$ and $\kappa(Q_n) = n$. In other words, Q_n is n-regular. The binary label of a vertex $v \in V(Q_n)$ will be referred to by an n-bit binary number $b(v)$. Also, $\boldsymbol{O}(b(v))$ and $\boldsymbol{Z}(b(v))$ will denote the number of ones and zeros, respectively, in the binary number $b(v)$.

III. The Twisted N-Cube

Let C be any shortest cycle (i.e., a 4-cycle) in Q_n. Also, let ux and vy be any two independent edges in C. The *twisted n-cube* graph TQ_n is then constructed as follows. Delete edges ux and vy from Q_n. Then, connect, via an edge, vertex u to vertex y, and vertex v to vertex x. That is, $TQ_n = Q_n - \{ux, vy\} + \{uy, vx\}$ Fig. 1 shows a Q_3 and a corresponding TQ_3. Note that by construction, TQ_n is n-regular just as Q_n is. Also, observe that TQ_n has two disjoint Q_{n-1} as subgraphs.

Although the cube can be twisted around any 4-cycle, we will usually use the *canonically twisted* Q_n where vertices u, v, x, and y have the labels $b(u) = 000\cdots0$, $b(v) = 010\cdots0$, $b(x) = 100\cdots0$, and $b(y) = 110\cdots0$. In the subsequent sections, we describe some of the properties of the twisted cube TQ_n.

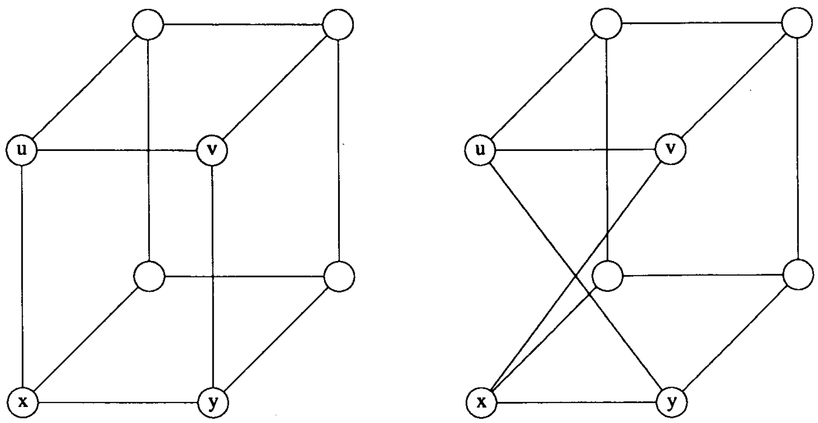

Q_3 TQ_3

Fig. 1. A 3-cube (Q_3) and a corresponding twisted 3-cube (TQ_3).

IV. Diameter of TQ_n

It is well known that the diameter of Q_n, $d(Q_n)$, is n. Also, between any pair of vertices u and v in Q_n there are n disjoint paths, of which $d(u,v)$ are of length $d(u,v)$ and the rest are of length $d(u,v)+2$ [21], [27]. As a result, if $d(u,v) \le n-1$ then there are at least $n-2$ disjoint paths between u and v, each of which is of length at most $n-1$. This property of Q_n will be used shortly.

Theorem 1: $d(TQ_n) = n-1$ for $n \ge 3$.

Proof: Let TQ_n be the canonically twisted cube. The theorem can easily be verified when n equals 3 and 4. Here we consider the case in which $n \ge 5$. Now, let s and t be any two vertices in TQ_n. We will show that in TQ_n we have $d(s,t) \le n-1$ for all s,t with equality for at least one pair. Depending on the value of $d(s,t)$ in Q_n, the following two cases are considered.

Case 1: In Q_n we have $d(s,t) \le n-1$. Then there are at least $n-2 \ge 3$ disjoint paths between s and t in Q_n, each of which is of length at most $n-1$. Thus, removal of edges ux and vy from Q_n can destroy at most two of such paths. This implies that in TQ_n we have $d(s,t) \le n-1$.

Case 2: In Q_n we have $d(s,t) = n$. Let $b(s) = (b_{n-1}b_{n-2}b_{n-3}\cdots b_1b_0)$ so that $b(t) = (\bar{b}_{n-1}\bar{b}_{n-2}\bar{b}_{n-3}\cdots \bar{b}_1\bar{b}_0)$ where $\bar{b}_i$ is the binary complement of b_i. A shortest s–t path in TQ_n can be constructed as follows.

First let us concentrate on the ones of $b(s)$ in positions $n-3$, $n-4,\cdots,0$. We can change these ones to zeros by traveling over a single edge for each exchange. Thus, after traveling $O(b_{n-3}b_{n-4}\cdots b_0)$ edges we will arrive at one of the vertices u, v, x, or y (which one is determined, of course, by the two leading bits $b_{n-1}b_{n-2}$ of s.) Next, we can change $b_{n-1}b_{n-2}$ to $\bar{b}_{n-1}\bar{b}_{n-2}$ by using a single edge of TQ_n. That edge will be uy or vx depending upon which of the four vertices we were led to by the first part of the path. Finally, all the zeros in $(b_{n-3}b_{n-4}\cdots b_0)$ must be turned to ones. Again a single edge is used for each of the $Z(b_{n-3}b_{n-4}\cdots b_0)$ bits involved. Hence, the total number of edges in our s–t path is

$$\begin{aligned} O(b_{n-3}b_{n-4}\cdots b_0) + Z(b_{n-3}b_{n-4}\cdots b_0)1 &= (n-2) + 1 \\ &= n-1. \end{aligned}$$

It is easy to see that there is no shorter s–t path: traveling over any edge of TQ_n changes only one bit with the exception of uy and vx which change two. It can be easily seen that edges uy and vx cannot both appear in any shortest path. Since $b(s)$ and $b(t)$ differ in all n positions, at least $n-1$ edges are needed to transform all bits. It follows that $d(s,t) = n-1$ by the construction above. Combining this fact with Case 1, we see that $d(TQ_n) = n-1$ as desired. □

V. Vertex-Connectivity of TQ_n

It is known that $\kappa(Q_n) = n$ [1], [21]. We next prove that $\kappa(TQ_n) = n$. In fact, we prove a more general connectivity theorem. Let G_1 and G_2 be two connected graphs with the same number p of vertices. Furthermore, let $V(G_1) = \{u_1, u_2, \cdots, u_p\}$ and $V(G_2) = \{v_1, v_2, \cdots, v_p\}$. Then $H = G_1 \odot G_2$ represents the graph obtained by taking G_1 and G_2 and connecting, via a new edge, vertex u_i to vertex v_i, for $1 \le i \le p$. That is,

$$V(H) = V(G_1) \cup V(G_2)$$

and

$$\begin{aligned} E(H) = {} & E(G_1) \cup E(G_2) \\ & \cup \{u_iv_i | u_i \in E(G_1), v_i \in E(G_2), 1 \le i \le p\}. \end{aligned}$$

The u_iv_i edges will be referred to as *cross* edges. Note that operation $\odot$ may generate different H graphs depending on how the vertices in graphs G_1 and G_2 are labeled [17].

Theorem 2: Let G_1 and G_2 be connected graphs defined as above, and let $H = G_1 \odot G_2$. Then $\kappa(H) \ge 1 + \min(\kappa(G_1), \kappa(G_2))$.

Proof: Let $k = \min(\kappa(G_1), \kappa(G_2))$, and let X be an arbitrary subset of $V(H)$ such that $|X| = k$. We prove the theorem by showing that $H - X$ is connected. Observe that H contains at least $k+1$ cross edges since k must be smaller than the number of vertices in each of the graphs G_1 and G_2. Therefore, removal of k vertices from H cannot cause deletion of all cross edges. Now if $X \cap V(G_1) = \varnothing$ (respectively, $X \cap V(G_2) = \varnothing$ then G_1 (respectively, G_2) is a connected subgraph of $H - X$. Furthermore, every remaining vertex of G_2 (respectively, G_1) is connected to this connected subgraph. Hence, $H - X$ is connected.

Now suppose $X \cap V(G_1) = X_1 \ne \varnothing$ and $X \cap V(G_2) = X_2 \ne \varnothing$. We must then have $1 \le |X_1| \le k-1$ and $1 \le |X_2| \le k-1$. This implies that both $G_1 - X_1$ and $G_2 - X_2$ are connected by definition of k. Since there is at least one cross edge, say e, in $H - X$, the endpoints of e lie in $G_1 - X_1$ and $G_2 - X_2$, and therefore $H - X$ must be connected. □

Theorem 3: $\kappa(TQ_n) = n$.

Proof: Clearly it is possible to take two copies of Q_{n-1} and label their vertices such that $TQ_n = Q_{n-1} \odot Q_{n-1}$. Since $\kappa(Q_{n-1}) = n-1$, Theorem 2 implies that $\kappa(TQ_n) \ge n$. Also for any n-regular graph G, $\kappa(G) \le n$, hence the desired result. □

VI. Lengths of Disjoint Paths in TQ_n

It is well known that if G is a graph with $\kappa(G) = n$, then given any two distinct vertices $s, t \in V(G)$ we can find n disjoint s–t paths

TABLE I
LENGTHS OF DISJOINT PATHS IN TQ_n, THE EXCEPTIONAL CASES

Exception	Possible Path Lengths				
	$d-1$	d	$d+1$	$d+2$	$d+3$
1. There is one fixed 1 in *suf* (s) and					
(a) $b_{n-1}b_{n-2} = \bar{c}_{n-1}\bar{c}_{n-2}$, or		d	1	$n-d-1$	
(b) either s, t are adjacent to u, x or s, t are adjacent to v, y		d		$n-d-1$	1
2. There is no fixed 1 in *suf* (s) and					
(a) $b_{n-1}b_{n-2} = \bar{c}_{n-1}\bar{c}_{n-2}$, or	1	$d-1$		$n-d$	
(b) $b_{n-1}b_{n-2} = c_{n-1}c_{n-2}$ with exactly one of s, t equal to u, v, x, y, or		d		$n-d-1$	1
(c) $b_{n-1}b_{n-2} = c_{n-1}\bar{c}_{n-2}$ with exactly one of s, t equal to u, v, x, y, or		d	1	$n-d-1$	
(d) $b_{n-1}b_{n-2} = \bar{c}_{n-1}c_{n-2}$ with s or t equal to u, v, x, y		$d-1$	2	$n-d-1$	

in G [15]. Below we will give an explicit description of such paths in TQ_n, but first it will be convenient to review the construction in Q_n.

If $s,t \in V(Q_n)$ with $b(s) = b_{n-1}b_{n-2}\cdots b_0$ and $b(t) = c_{n-1}c_{n-2}\cdots c_0$, then position i will be called *fixed* (respectively, *variable*) with respect to s and t if $b_i = c_i$ (respectively, $b_i \neq c_i$). A *fixed one* is a fixed position, i, with $b_i = c_i = 1$. A *fixed zero* is similarly defined.

Now if $d(s,t) = d$ then s and t determine a d-dimensional subcube $Q^0 = Q^0(s,t) \subseteq Q_n$ where $V(Q^0) = \{v \in V(Q_n) \,|\, b(v)$ agrees with $b(s)$ and $b(t)$ in all $n-d$ fixed positions$\}$. The vertices s and t are called *diametrically* opposite points of this subcube. In fact, there are 2^{d-1} pairs of points that are diametrically opposite in the same Q^0.

We can construct d disjoint s–t paths of length d inside Q^0 as follows. Let $i_0, i_1, \cdots, i_{d-1}$ be any permutation of the variable positions. Then the jth path is obtained by changing the bit in positions $i_j, i_{j+1}, \cdots, i_{d-1}, i_0, \cdots, i_{j-1}$ in that order. The permutation $i_0, i_1, \cdots, i_{d-1}$ will be called a *changing order*. Although the changing order can be arbitrary in Q_n, the asymmetry of TQ_n will make certain changing orders more efficient.

To construct the $n-d$ disjoint s–t paths of length $d+2$ in Q_n, we travel to $n-d$ *parallel d-subcubes* $Q^1, Q^2, \cdots, Q^{n-d}$ where $V(Q^j) = \{v \in V(Q_n) \,|\, b(v)$ agrees with $b(s)$ and $b(t)$ in all but the jth fixed position where $b(v)$ has the complementary bit$\}$. These paths through Q^j start by changing the jth fixed position bit of s to its complement so that we are now at a vertex $s_1 \in V(Q^j)$. Next travel from s_1 to its diametrically opposite point $t_1 \in V(Q^j)$. Finally change the jth fixed position back to its original value to reach t.

To carry over our discussion to the canonically twisted n-cube we will again be concerned with the last $n-2$ bits. Define the *suffix* of vertex $v \in V(Q_n)$ with $b(v) = b_{n-1}b_{n-2}\cdots b_0$ to be *suf*$(v) = b_{n-3}b_{n-4}\cdots b_0$. We will first start with the special case where s and t are diametrically opposite (since the definition of such points only depends on the bit vector, it still makes sense in TQ_n). The proof will be a refinement of the argument used for Theorem 1.

Theorem 4: For $n \geq 3$ let TQ_n be the canonical twisted n-cube and consider $s,t \in V(TQ_n)$ that are diametrically opposite. Then there are n disjoint s–t paths in TQ_n consisting of one path of length $n-1$ and $n-1$ paths of length n.

Proof: Since $n \geq 3$ and $b(s), b(t)$ are complements we can assume (without loss of generality) that *suf*(s) contains a 1 bit. Use the changing order $i_0, i_1, \cdots, i_k, \cdots, i_{n-2} = n-2, i_{n-1} = n-1$ with k chosen so that $i_0, i_1, \cdots, i_{k-1}$ are the positions of *suf*(s) containing zeros and $i_k, \cdots, i_{n-3}$ are the positions of *suf*(s) containing ones. Now attempt to construct the jth path as done in Q_n.

Since TQ_n is canonical, there will be no difficulties until we reach a vertex w where the leading bit must be altered. But by the order in which the bits are changed we see that if $j \neq k$ then *suf*(w) still contains a 1 and the $(n-1)$st bit can be complemented as usual. Thus, for $n-1$ value of j we have paths of length n as before.

When $j = k$ the order that we have chosen guarantees that we still travel to w such that $b(w) = b_{n-1}b_{n-2}00\cdots 0 \in \{u,v,x,y\}$ and can change both the leading bits using one edge. Ending the path by altering positions $i_0, \cdots, i_{k-1}$ gets us from s to t in only $n-1$ steps. □

Another theorem characterizing disjoint paths in the n-cube graph with an edge deleted will be useful.

Theorem 5: Let s and t be diametrically opposite in Q_n where $n \geq 2$. Consider $G = Q_n - \{e\}$ where $e \in E(Q_n)$ is arbitrary. Then there are still n disjoint s–t paths of length n unless s or t is an endpoint of e. In the latter case, there are only $n-1$ disjoint s–t paths (but still there is at least one path since $n \geq 2$).

Proof: Suppose for the sake of definiteness that $e = ux$ where $b(u) = 00\cdots 0$, and $b(x) = 10\cdots 0$. If $s,t \notin \{u,x\}$ then some nonleading digit of s must be a 1 (since $s \neq u,x$) and some other nonleading digit of s must be a 0 (since $t \neq u,x$ and $b(s) = \overline{b(t)}$). Use the changing order $i_0, i_1, \cdots, i_k, \cdots, i_{n-1} = n-1$ with k chosen so that $i_0, i_1, \cdots, i_{k-1}$ are the nonleading positions of s containing ones and $i_k, i_{k+1}, \cdots, i_{n-2}$ are then nonleading positions of s containing zeros.

We must check that any path in Q_n which contains u or x is still intact in G. But by the changing order and the fact that the segment $i_k, \cdots, i_{n-2}$ is nonempty, we see that the only such path is the one starting with $i_{n-1} = n-1$. Thus, the leading bit is complemented in the very first step (which is still possible since $i_0, \cdots, i_{k-1}$ is nonempty). Hence, when we reach u or x, a bit other than the leading one is to be altered and no complications ensue.

If, on the other hand, s is an endpoint of e (the case where t is an endpoint is similar) then $deg(s) = n-1$ in G. It follows that there can be at most $n-1$ disjoint s–t paths in G. But removal of an edge of Q_n destroys at most one of the n disjoint s–t paths in the cube. Hence, exactly $n-1$ disjoint s–t paths are left in G. □

Finally we are in a position to prove the main result of this section.

Theorem 6: Let TQ_n be the canonically twisted n-cube and consider $s,t \in V(TQ_n)$ with $b(s) = b_{n-1}b_{n-2}, \cdots, b_0$ and $b(t) = c_{n-1}c_{n-2}, \cdots, c_0$. If $d(s,t) = d$ in Q_n then a set of n disjoint paths consisting of d of length d and $n-d$ of length $d+2$ continues to exist in TQ_n with the exception of the case noted in Table I. If the entry in row i and column $d+j$ is k this means that there are k disjoint s–t paths of length $d+j$ for exception i. A blank indicates no such paths. All paths for a given row can be taken to be disjoint.

Proof: First consider the case where the number of fixed 1's in *suf*(s) is at least two. Then all of the cubes $Q^0, Q^1, \cdots, Q^{n-d}$ remain intact in TQ_n since every $w \in V(Q^j)$ has at least one 1 in *suf*(w). Hence, the canonically s–t paths in Q_n can still be used and those pairs are not exceptional.

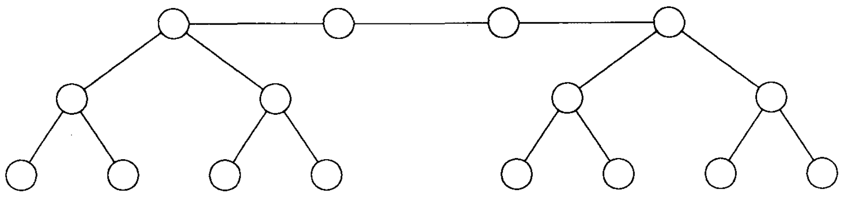

S_4

Fig. 2. An S_4 is formed by connecting two trees, T_3's, with their roots being three hops apart.

Now suppose that there is exactly one fixed 1 in *suf*(*s*), say in position i_j. Then $Q^0, Q^1, \cdots, Q^{j-1}, Q^{j+1}, \cdots, Q^{n-d}$ are still subgraphs of TQ_n. Thus, the d paths of length d along with $n-d-1$ paths of length $d+2$ are still usable. To determine how to modify the path through Q^j let s_1, t_1 be the vertices of Q^j adjacent to s, t, respectively, and consider the first two bits of $b(s)$ and $b(t)$.

If $b_{n-1} = c_{n-1}$ then Q^j is once again contained in TQ_n so this does not qualify as an exceptional case.

If $b_{n-1}b_{n-2} = \bar{c}_{n-1}\bar{c}_{n-2}$ then Q^j becomes a subgraph isomorphic to TQ_d in TQ_n. Thus, by Theorem 4 there is an s–t_1 path of length $d-1$. It follows that there is an s–t path of length $(d-1)+2 = d+1$ through Q^j. This completes case 1(a).

If $b_{n-1}b_{n-2} = \bar{c}_{n-1}c_{n-2}$ then Q^j is missing an edge in TQ_n. But Theorem 5 assures us that there is still at least one s_1–t_1 path of length d, and hence an s–t path of length $d+2$, unless $d=1$. This implies that we are in case 1(b) with $\{s_1, t_1\} = \{u, x\}$ or $\{s_1, t_1\} = \{v, y\}$. In either case, we can get from s_1 to t_1 in two steps via one of the four vertices u, v, x, y. Hence, this s–t path has length $1+2+1 = d+3$.

The time has come to see what happens when *suf*(*s*) contains no fixed 1's. As usual, the initial pairs of bits play a critical role.

If $b_{n-1}b_{n-2} = \bar{c}_{n-1}\bar{c}_{n-2}$ then Q^0 becomes a modified TQ_d inside TQ_n. Theorem 4 applies immediately to yield one path of length $d-1$ and $d-1$ paths of length d. Furthermore, $Q^1, \cdots, Q^{n-d}$ are all normal d-cubes so the $n-d$ paths of length $d+2$ are undisturbed. Case 2(a) follows.

Now suppose $b_{n-1}b_{n-2} = c_{n-1}c_{n-2}$ so that $Q^0, Q^1, \cdots, Q^{n-d}$ all remain intact in TQ_n. The only difficulty that could arise is that it may not be possible to travel from s (or t) to Q^j, $j \geq 1$ via a single edge. This situation only occurs if s (or t) is one of u, v, x, y and Q^j is the subgraph determined by changing the leading bit. For the sake of definiteness, let $s \in \{u, v, x, y\}$ (in which case t is not one of these four vertices since $b_{n-1}b_{n-2} = c_{n-1}c_{n-2}$). Every path of length $d+2$ is usable except the one through Q^j. But we can pass from s to $s_2 = \bar{b}_{n-1}\bar{b}_{n-2}b_{n-3}\cdots b_0$ to $s_1 = \bar{b}_{n-1}b_{n-2}\cdots b_0$ using two edges, then go from s_1 to its diametrically opposite point t_1 using d edges, and finally travel from t_1 to t using a single edge. Note that since $s_2 \notin V(Q^j)$ for all j, this path is disjoint from all the others and has length $2+d+1 = d+3$. Case 2(b) is now completed.

If $b_{n-1}b_{n-2} = c_{n-1}\bar{c}_{n-2}$ then the reasoning of the previous paragraph shows again that the only problems which can occur are when s or $t \in \{u, v, x, y\}$ and Q^j is the subcube for changing the leading bit. Here we cannot exclude the possibility that both s and t are in the set of special vertices. But in that case $\{s, t\} \cup V(Q^j) = \{u, v, x, y\}$ so there is an s–t path through Q^j of length $3 = d+2$ which is not exceptional.

Thus, we are reduced to case 2(c). Assume, without loss of generality, that $s \in \{u, v, x, y\}$ and travel to $s_2 = \bar{b}_{n-1}\bar{b}_{n-2}b_{n-3}\cdots b_0 \in V(Q^j)$ using a single edge. But s_2 is one step closer to $t_1 = \bar{c}_{n-1}c_{n-2}\cdots c_0$ than is s_1. Hence, we can reach t in only $(d+2)-1 = d+1$ steps.

Finally, consider the case where $b_{n-1}b_{n-2} = \bar{c}_{n-1}c_{n-2}$. Thus, Q^0 is either a d-cube or a d-cube missing an edge e. Assume for the moment that neither s nor t is among u, v, x, y. Then even if Q^0 has e deleted, the edge does not contain s or t. Thus, (by Theorem 5 if necessary) Q^0 will have d disjoint s–t paths of length d. The $n-d-1$ parallel subcubes obtained by changing a 0 bit in $suf(s)$ are all full subcubes in TQ_n and so we always have $n-d-1$ disjoint s–t paths of length $d+2$. Now let Q^j be the graph obtained by complementing the second leading bit with s_1, t_1 the diametrically opposite vertices corresponding to s, t. Q^j may be full or missing an edge e_1, but in any event neither s_1 nor t_1 are endpoints of e_1 (since the same is true of s, t, and e). Thus, an s–t path of length $d+2$ still exists through Q^j. Hence, if $\{s, t\} \cap \{u, v, x, y\} = \emptyset$ we do not have an exceptional case.

The only other possibility is that 2(d) holds, so assume $s \in \{u, v, x, y\}$ (the case $t \in \{u, v, x, y\}$ being similar). Now Q^0 is missing an edge containing s, so we are guaranteed only $d-1$ paths of length d in Theorem 5. Also a path of length $d+1$ can be constructed by traveling to $s_2 = b_{n-1}\bar{b}_{n-2}00\cdots 0$ using one edge, then to $s_1 = \overline{b_{n-1}}b_{n-2}00\cdots 0 \in V(Q^0)$ using another, and finally continuing through Q^0 to t using $d-1$ edges. As in the preceding paragraph, there are still $n-d-1$ disjoint s–t paths of length $d+2$ corresponding to the fixed zeros in *suf*(*s*). The final path through Q^j is started by moving from s to $s_3 = \bar{b}_{n-1}\bar{b}_{n-2}00\cdots 0$. This is one step closer to t than usual, so only $d+1$ edges are needed for the trip. Thus, we have a second s–t path of length $d+1$. This completes the last case and hence the proof of theorem. □

Note that in all exceptional cases but two [specifically 1(b) and 2(a)] the average length of the n paths in Table I is at least as short as the average length between the same two points in Q_n. In fact, for some of the cases above the paths from Q_n still exist in TQ_n but the listed set of the paths will be shorter.

VII. Some Subgraphs of TQ_n

In this section, we will identify some of the subgraphs of TQ_n. By construction, Q_{n-1} is a subgraph of TQ_n and thus all its subgraphs are contained in TQ_n. In fact, given any subgraph H of Q_n, if there exists some 4-cycle, say C, in Q_n, such that $E(H) \cap E(C)$ is either empty or a set containing two independent edges, then H is also a subgraph of TQ_n. This implies that TQ_n contains a 2^n cycle and any two-dimensional *mesh* which is a subgraph of Q_n. While Q_n contains only even cycles, TQ_n contains odd cycles as well.

In what follows, we will show that the complete binary tree on $2^n - 1$ vertices, T_n, is a subgraph of TQ_n. It is known that T_n is not a subgraph of Q_n [27]. However, T_{n-1} is contained in Q_n [4]. To present our result, we need to show that two disjoint copies of T_{n-1} can be found in Q_n. This was first demonstrated by Praha, rediscovered independently by Bhatt and Ipsen, and then rerediscovered by us [25], [3]. We include the proof for the sake of completeness.

Let S_n denote the graph obtained by taking two disjoint complete binary tree T_{n-1} and connecting their roots by a path of length 3. A picture of S_4 is given in Fig. 2.

Theorem 7: For $n \geq 2$, S_n is a subgraph of Q_n. Furthermore, for $n \geq 3$ the roots r and u of the two copies of T_{n-1} can be labeled so that $b(r)$ and $b(u)$ differ in exactly three positions.

Proof: Fig. 3 gives labelings which embed S_n in Q_n for $n = 2, 3$. Note that in the latter case the labels along the path r–s–t–u of

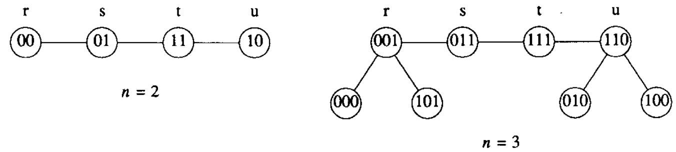

Fig. 3. A labeling scheme to embed S_n in Q_n for $n = 2, 3$.

length 3 are 001, 011, 111, and 110, respectively. By induction we may assume that S_{n-1} is isomorphic to a subgraph of Q_{n-1} with $b(r) = 001\cdots1, b(s) = 011\cdots1, b(t) = 111\cdots1$, and $b(u) = 11\cdots10$.

Now we can find two disjoint subgraphs isomorphic to S_{n-1}, call them S^0_{n-1} and S^1_{n-1}, in Q_n as follows. S^0_{n-1} is obtained by prefixing every label of S_{n-1} with a 0. Thus, the labels of the corresponding path of length 3 are $b(r^0) = 0001\cdots1$, $b(s^0) = 0011\cdots1$, $b(t^0) = 0111\cdots1$ and $b(u^0) = 011\cdots10$. If $v \in S_{n-1}$ is labeled $b(v) = (b_{n-2}b_{n-3}\cdots b_0)$ in Q_{n-1} then in S^1_{n-1} we let $b(v^1) = (1b_0b_1\cdots b_{n-2})$. In particular, $b(r^1) = 11\cdots100$, $b(s^1) = 11\cdots110$, $b(t^1) = 11\cdots111$, and $b(u^1) = 101\cdots11$. A schematic drawing of these subgraphs if displayed in Fig. 4(a). Now, the graph S_n is created by letting

$$S_n = S^0_{n-1} \cup S^1_{n-1} + \{s^0u^1, t^0t^1, u^0s^1\} - \{t^0u^0, t^1u^1\}$$

as in Fig. 4(b). Finally the new roots are s^0 and s^1 with labels $001\cdots1$ and $11\cdots10$, respectively, which differ in exactly three positions. □

Theorem 8: T_n is subgraph of TQ_n.

Proof: Find a subgraph of Q_n which is isomorphic to S_n with the path r–s–t–u labeled as in Theorem 7, that is $b(r) = 001\cdots1$, $b(s) = 011\cdots1$, $b(t) = 111\cdots1$, and $b(u) = 11\cdots10$. Let $v \in V(Q_n)$ be the vertex that $b(v) = 101\cdots1$. Note that the vertices r, s, t, and v induce a 4-cycle. Since TQ_n can be constructed by exchanging any two independent edges in any 4-cycle of Q_n, we will use the above 4-cycle and construct TQ_n. In particular,

$$TQ_n = Q_n - \{rs, tv\} + \{rt, sv\}.$$

Clearly, $T_n = S_n + \{rt\} - \{s\}$ is a subgraph of TQ_n (note that $tv \notin E(S_n)$ so that the removal of this edge from Q_n causes no difficulties). □

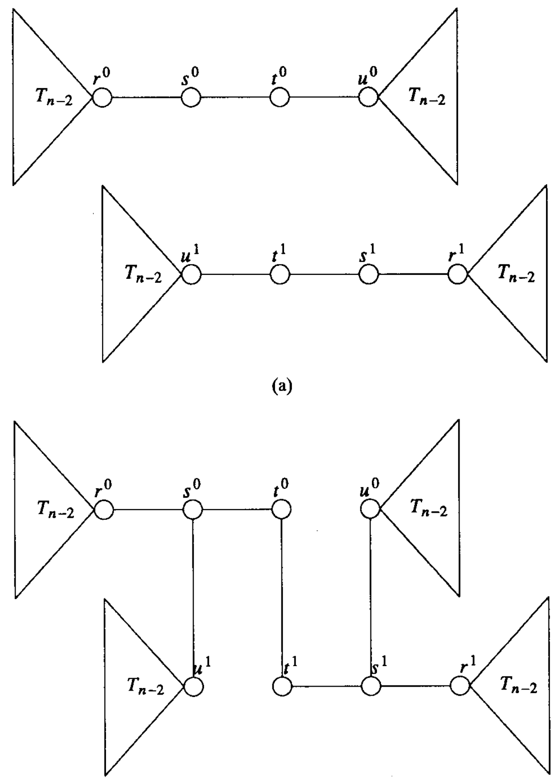

Fig. 4. S_n is constructed using two copies of S_{n-1}. (a) Two copies of S_{n-1}. (b) Constructing S_n using two copies of S_{n-1}.

VIII. Routing in Twisted N-Cube

A hypercube can be converted to TQ_n by exchanging two of its physical links. Or two extra physical links can be added to a hypercube multiprocessor to obtain a topology which has both Q_n and TQ_n as subgraphs. Note that in the latter case, the graph is no longer n-regular, and four nodes will have degree $n + 1$. In both cases, other components of the system should be modified accordingly. One major component is the *router* at each processing node. In what follows, we address this issue for both cases.

Each processor (vertex) in the hypercube multiprocessor has a router to handle the interprocessor communication [22]. The function of the router may be performed by the processor or by a dedicated router chip. In a hypercube multiprocessor, upon receiving a message, a *routing tag* $(r_{n-1}r_{n-2}r_{n-3}\cdots r_0)$ is obtained by taking a bit-wise Exclusive-OR operation between the router's local address $(c_{n-1}c_{n-2}\cdots c_0)$ and the destination address $(d_{n-1}d_{n-2}\cdots d_0)$ of the message. The message can then be forwarded to one of the neighboring processors through the jth link if $r_j = 1$ for $0 \le j \le n-1$.

To support the TQ_n topology, the function of the routers should be slightly modified. For these routers, the routing tag is computed as above. Suppose TQ_n is the canonically twisted n-cube. Let us first consider the four routers at vertices u, v, x, and y; we will refer to these routers as *twisted* routers. If $r_{n-1}r_{n-2} = 01$ then the message is forwarded through the $(n-2)$nd link, that is, either uv or xy. If $r_{n-1}r_{n-2} = 11$ then the message is forwarded through the $(n-1)$st link, that is, either uy or vx. Note that in this case one routing step is saved compared to that in Q_n. If $r_{n-1}r_{n-2} = 10$ then the message is forwarded through the $(n-2)$nd link if $r_j = 0$ for all $0 \le j \le n-3$; otherwise, the message is forwarded through some jth link with $r_j = 1$ where $0 \le j \le n-3$. Note that in the former case, it will take two routing steps rather than one as required in Q_n. However, this additional routing step may not be necessary if the message is forwarded through other links first as in the latter case.

The function of the remaining $2^n - 4$ routers will also have to be slightly modified in order to take advantage of a possible saving of one routing step. If $r_{n-1}r_{n-2} = 11$ then one routing step can be saved by first forwarding the message to the node $d_{n-1}d_{n-2}00\cdots0$, one of the four twisted routers, and then the message is forwarded to the final destination. If $r_{n-1}r_{n-2} = 10$, then the message has to be forwarded through the $(n-1)$th link if there exists only one j $(0 \le j \le n-3)$ such that $r_j = 1$. This is to avoid having an additional routing step. For all other cases, the message can be forwarded to any jth link so long as $r_j = 1$.

For the case where two edges (i.e., uy and vx) are added to Q_n, the routers are modified as follows. For the four twisted routers, now each with $n+1$ links, if $r_{n-1}r_{n-2} = 11$ then the message should be forwarded through the added link. Thus, one routing step is saved. For all other cases, the normal routing procedure should be followed. For the remaining $2^n - 4$ routers, if $r_{n-1}r_{n-2} = 11$, then one routing step can be saved by first forwarding the message to the node $d_{n-1}d_{n-2}00\cdots 0$, one of the four twisted routers, and then the message is forwarded to the final destination.

Although hypercube multiprocessors have been commercially available for several years and some routing algorithms have been proposed, the study of routing algorithms to the global performance of communication networks has not been investigated until the recent work by [13] and [6]. Performance evaluation of communication networks is difficult because there are many parameters involved. Among them, the communication patterns, which are dependent on the parallel applications, and the communication technologies are two major parameters. Further study is needed to evaluate the global network performance of the twisted n-cube.

IX. Concluding Remarks

The hypercube interconnection topology, due to its powerful topological properties, has been widely adopted in the construction of distributed-memory multiprocessors. In this paper, we have shown that by exchanging any two independent edges in any shortest cycle of the hypercube, an interconnection topology, namely TQ_n, can be achieved that has some nice properties.

In summary, the twisted n-cube, TQ_n, has the following properties as the n-cube Q_n. TQ_n consists of two disjoint Q_{n-1} subgraphs. Even rings and two-dimensional meshes are subgraphs of TQ_n. TQ_n is n-regular and its vertex connectivity remains n. In addition, TQ_n has the following unique properties not possessed by Q_n. Any odd length ring with $2^n - 1$ or fewer vertices is contained in TQ_n. A complete binary tree with $2^n - 1$ vertices, which is a highly demanded topology by many applications, is a subgraph of TQ_n [7]. The worst case number of routing steps is reduced from n to $n-1$. Furthermore, the average number of routing steps is also reduced. This implies improvement on communication delay which is critical to system performance.

Existing hypercube multiprocessors can be modified to take advantage of this new topology in two ways. A hypercube can be converted to TQ_n by exchanging two of its physical links. Second, two extra physical links can be added to a hypercube multiprocessor to obtain a topology which has both Q_n and TQ_n as subgraphs. For example, the FPS T-series can be easily twisted [14].

The results presented in this paper can naturally give rise to the following two graph theoretic questions.

1) What is the minimum number of *edge twistings* (that is, exchanging two independent edges in any 4-cycle) that can be done to Q_n to obtain a graph whose diameter is k less than that of Q_n?
2) What is the minimum number of new edges which can be added to Q_n to obtain a graph whose diameter is k less than that of Q_n?

We have solved problem 1) for $k = 1$. It is not difficult to prove that, for $k = 1$, the answer to problem 2) is 2 (we have already shown that it is at most 2). However, for $k > 1$, no elegant solution is known for either problem.

Finally, an interesting generalization of our work was proposed by Prof. Weigeng Shi of the Department of Electrical Engineering at Worcester Polytechnic Institute, which is stated as a conjecture.

Conjecture: It is possible to exchange $2^{\lceil \log_2 n \rceil - 1}$ independent edges in any $(n-1)$-cube of an n-cube, $n \geq 3$, so as to obtain a graph whose diameter is $\lceil \log_2 n \rceil - 1$ less than that of Q_n.

Acknowledgment

We would like to thank the reviewers for their helpful suggestions and constructive comments. We are also indebted to S. Seidel of Michigan Technological University for pointing out that the FPS T-series can be easily *twisted.*

References

[1] J. R. Armstrong and F. G. Gray, "Fault diagnosis in a Boolean n-cube array of microprocessors," *IEEE Trans. Comput.,* vol. C-30, no. 8, pp. 587–590, Aug. 1981.
[2] H.-J. Bandelt and H. M. Mulder, "Infinite median graphs, (0, 2)-graphs, and hypercubes," *J. Graph Theory,* vol. 7, pp. 487–497, 1983.
[3] S. N. Bhatt and C. F. Ipsen, "How to embed trees in hypercubes," Res. Rep. YALEU/DCS/RR-443, Dep. Comput. Sci., Yale Univ., Dec. 1985.
[4] J. E. Brandenburg and D. S. Scott, "Embeddings of communication trees and grids into hypercubes," Tech. Rep., Intel Scientific Computers, 1985.
[5] T. F. Chan and Y. Saad, "Multigrid algorithms on the hypercube multiprocessor," *IEEE Trans. Comput.,* vol. C-35, no. 11, pp. 969–977, Nov. 1986.
[6] W. J. Dally, "Performance analysis of k-ary n-cube interconnection networks," *IEEE Trans. Comput.,* pp. 775–785, June 1990.
[7] E. Dekel and S. Sahni, "Binary trees and parallel scheduling algorithms," *IEEE Trans. Comput.,* pp. 307–312, Mar. 1983.
[8] P. Erdos and J. Spencer, "Evolution of the n-cube," *J. Comput. Math. Appl.,* vol. 5, pp. 33–39, 1979.
[9] A. H. Esfahanian, "Generalized measures of fault-tolerance with application to n-cube networks," *IEEE Trans. Comput.,* to be published.
[10] E. T. Fathi and M. Krieger, "Multiple microprocessor systems: What, why, and when," *IEEE Comput. Mag.,* vol. 16, pp. 23–35, Mar. 1983.
[11] S. Foldes, "A characterization of hypercubes," *J. Discrete Math.,* vol. 17, pp. 155–159, 1977.
[12] D. C. Grunwald and D. A. Reed, "Benchmarking hypercube hardware and software," Tech. Rep. UIUCDCS-R-86-1303, Dep. Comput. Sci., Univ. of Illinois at Urbana–Champaign, 1986.
[13] ——, "Networks for parallel processors: Measurements and prognostications," in *Proc. Third Conf. Hypercube Concurrent Comput. Appl.,* vol. I, 1988, pp. 610–619.
[14] J. L. Gustafson, "Increasing hypercube communication on low-dimensional problems," *Hypercube Multiprocessors 1987,* SIAM, pp. 244–250, 1987.
[15] F. Harary, *Graph Theory.* Reading, MA: Addison-Wesley, 1972.
[16] S. Hart, "A note on the edges of the n-cube," *J. Discrete Math.,* vol. 14, pp. 157–163, 1976.
[17] S. Hedetniemi, "On classes of graphs defined by special cutsets of lines," *The Many Facets of Graph Theory,* Lecture Notes in Mathematics, vol. 110, Springer-Verlag, 1969, pp. 171–189.
[18] J. P. Hayes, T. N. Mudge, Q. F. Stout, S. Colley, and J. Palmer, "Architecture of a hypercube supercomputer," in *Proc. 1986 Int. Conf. Parallel Processing,* Aug. 1986, pp. 653–660.
[19] K. Hwang and F. A. Briggs, *Computer Architecture and Parallel Processing.* New York: McGraw-Hill, 1984.
[20] K. Hwang and J. Ghosh, "Hypernet: A communication-efficient architecture for constructing massively parallel computers," *IEEE Trans. Comput.,* vol. C-36, no. 12, pp. 1450–1467, Dec. 1987.
[21] J. C. Kuhl, "Fault diagnosis in computing networks," Tech. Rep., Dep. Elec. Comput. Eng., Univ. of Iowa, Iowa City, Iowa, 1980.
[22] Y. Lan, A. H. Esfahanian, and L. M. Ni, "Multicast in hypercube multiprocessors," *J. Parallel Distributed Comput.,* pp. 30–41, Jan. 1990.
[23] M. Mulder, "n-cube and median graphs," *J. Graph Theory,* vol. 4, pp. 107–110, 1980.
[24] D. K. Pradhan, "Interconnection topologies for fault-tolerant parallel and distributed architectures," in *Proc. 10th Int. Conf. Parallel Processing,* Aug. 1981, p. 238–242.
[25] L. N. Praha, "On cubes and dichotomic trees," *Casopis Pro Pestovani Matematiky,* roc. 99, 1974.
[26] J. C. Peterson, J. O. Tuazon, D. Lieberman, and M. Pniel, "The Mark III hypercube-ensemble concurrent computer," in *Proc. 1985 Int. Conf. Parallel Processing,* Aug. 1985, pp. 71–73.
[27] Y. Saad and M. H. Schultz, "Topological properties of hypercubes," Tech. Rep., YALEU/DCS/RR-389, Dep. Comput. Sci., Yale Univ., June 1985.
[28] C. Seitz, The Cosmic Cube," *Commun. ACM,* vol. 28, no. 1, pp. 22–23, Jan. 1985.
[29] K. N. Venkataraman, G. Cybenko, and D. W. Krumme, "Hypercube embedding is NP-complete," *Hypercube Multiprocessor 1986,* SIAM, SIAM, Philadelphia, PA, pp. 148–157, 1986.