

FAULT-TOLERANT EMBEDDINGS OF HAMILTONIAN CIRCUITS IN K -ARY N -CUBES*

YAAGOUB A. ASHIR[†] AND IAIN A. STEWART[‡]

Abstract. We consider the fault-tolerant capabilities of networks of processors whose underlying topology is that of the k -ary n -cube Q_n^k , where $k \geq 3$ and $n \geq 2$. In particular, given a copy of Q_n^k where some of the inter-processor links may be faulty but where every processor is incident with at least two healthy links, we show that if the number of faults is at most $4n - 5$ then Q_n^k still contains a Hamiltonian circuit; but that there are situations where the number of faults is $4n - 4$ (and every processor is incident with at least two healthy links) and no Hamiltonian circuit exists. We also remark that given a faulty Q_n^k , the problem of deciding whether there exists a Hamiltonian circuit is NP-complete.

Key words. Hamiltonian circuits, embeddings, fault-tolerance, k -ary n -cubes, NP-completeness

AMS subject classifications. 68R10, 05C45

1. Introduction. The hypercube or, more precisely, the binary n -cube B_n (where $n \geq 2$), is a popular interconnection network for parallel processing as it possesses a number of topological properties which are highly desirable in the context of parallel processing: for example, it contains a Hamiltonian circuit; many other networks can be embedded into a binary n -cube; and its symmetry results in rich communication properties (see, for example, [3, 5, 8, 10, 12] and the references therein).

Fault tolerance in the binary n -cube is an important issue, given that many other networks can be embedded therein, and has been studied in a number of contexts. For example, the ability of the binary n -cube to route and reconfigure itself in spite of faults has been considered (see the references in [8]), as has the embedding of Hamiltonian circuits in binary n -cubes in the presence of faults [8]. In particular, Chan and Lee [8] proved that a binary n -cube where at most $2n - 5$ links are faulty and where every node is incident with at least two healthy links (a natural assumption to make) has a Hamiltonian circuit; but that there exist binary n -cubes with $2n - 4$ faults (and where every node is incident with at least two healthy links) not containing a Hamiltonian circuit. It is with an analogous version of this result that we are concerned in this paper.

One drawback of the binary n -cube is that the number of links incident with each node is logarithmic in the number of nodes, and this causes problems with regard to current VLSI technology when the networks built upon the binary n -cube topology involve a large number of processors. One means proposed to alleviate this problem is to base networks on the topology of the k -ary n -cube Q_n^k (where $k \geq 3$ and $n \geq 2$). A network based on Q_n^k is such that each node is incident with $2n$ links, and consequently k can be increased, in order to incorporate more processors, at the same time keeping n constant. Moreover, “high-dimensional” networks generally cost more and run more slowly than “low-dimensional” networks, and it has also been shown that low-dimensional networks achieve lower latency and better hot-spot throughput

*The first author was financially supported by the University of Bahrain

[†]Department of Mathematics and Computer Science, Leicester University, Leicester LE1 7RH, U.K. (y.ashir@mcs.le.ac.uk).

[‡]Department of Mathematics and Computer Science, Leicester University, Leicester LE1 7RH, U.K. (i.a.stewart@mcs.le.ac.uk).

than their high-dimensional counterparts [9, 11].

The properties of the k -ary n -cube Q_n^k relevant to parallel processing have not been determined to such an extent as those of the binary n -cube: however, some work has been done (see, for example, [1, 2, 4, 6, 7]). In particular, it has been shown that Q_n^k has a Hamiltonian circuit [6].

In this paper, we examine the number of link faults that a k -ary n -cube Q_n^k can tolerate so that there is still a Hamiltonian circuit (of course, we assume that every node is incident with at least 2 healthy links). In particular, we show that a k -ary n -cube Q_n^k where at most $4n - 5$ links are faulty and where every node is incident with at least two healthy links has a Hamiltonian circuit; but that there exist k -ary n -cubes with $4n - 4$ faults (and where every node is incident with at least two healthy links) not containing a Hamiltonian circuit. We also remark that the general problem of deciding whether a faulty k -ary n -cube contains a Hamiltonian circuit is NP-complete, for all (fixed) $k \geq 3$. Our results can be regarded as direct analogues of those in [8] for k -ary n -cubes as opposed to binary n -cubes.

2. Tolerating faults. Throughout this paper, we prefer to use the terminology “nodes” and “links” as opposed to “vertices” and “edges”, for whilst the results in this paper are entirely graph-theoretic, the use of “nodes” and “links” accentuates the motivational source of our research, i.e., the fault-tolerating capabilities of networks of processors when the faults which may occur are the failures of the links between processors in the network.

The binary n -cube, for $n \geq 2$, can be represented as the set of 2^n nodes $\{0, 1\}^n$ where there is a link joining nodes u and v if and only if u and v agree on all components except one. Note that each node has degree n . The k -ary n -cube Q_n^k , for $k \geq 2$ and $n \geq 2$, can be represented as the set of k^n nodes $\{0, 1, \dots, k - 1\}^n$ where there is a link joining nodes u and v if and only if u and v agree on all components except one, and on that component they differ by 1 modulo k . Note that each node has degree $2n$, when $k \geq 3$, and n when $k = 2$. In particular, Q_n^2 is simply B_n .

For each $i \in \{1, 2, \dots, n\}$, we refer to all links whose incident nodes differ in the i th component as lying in *dimension* i . Note that for any $i \in \{1, 2, \dots, n\}$, Q_n^k consists of k disjoint copies of Q_{n-1}^k where corresponding nodes are joined in circuits of length k using links in dimension i . When we consider Q_n^k in this way, with the disjoint copies joined by links lying in dimension i , we say that we have *partitioned* Q_n^k *over dimension* i .

Let us now proceed to the proof of our main theorem. This proof is by induction. We begin by proving the inductive step and then we return to the base cases of the induction.

THEOREM 2.1. *Let $k \geq 4$ and $n \geq 2$, or let $k = 3$ and $n \geq 3$. If Q_n^k has at most $4n - 5$ faulty links and is such that every node is incident with at least 2 healthy links then Q_n^k has a Hamiltonian circuit.*

Proof. The proof proceeds by induction on n . We handle the base cases, when $n = 2$ and $k \geq 4$ and when $n = 3$ and $k = 3$, later. As our induction hypothesis, assume that the result holds for Q_n^k , for some $n \geq 2$ and for all $k \geq 4$, or for some $n \geq 3$ and $k = 3$. Let Q_{n+1}^k have $4n - 1$ faults and be such that every node is incident with at least 2 healthy links. Then there exists some dimension, say dimension 1, which contains at least 3 faults. We can partition Q_{n+1}^k over dimension 1 and consider Q_{n+1}^k to consist of k disjoint copies Q_1, Q_2, \dots, Q_k of Q_n^k with corresponding nodes joined in circuits of length k , where the faults contained in Q_1, Q_2, \dots, Q_k total at most $4n - 4$ (see Fig. 1). Throughout this proof, if u is a node of Q_i , say, then we often denote

it by u_i , and we refer to its corresponding node in Q_j as u_j . Our general aim below is to argue, using induction, that Hamiltonian circuits exist in each of Q_1, Q_2, \dots, Q_k and that we can “join” these circuits together using links in dimension 1 to obtain a Hamiltonian circuit in Q_{n+1}^k (what we mean by “join” will become clear later: also, the general aim of connecting together circuits in Q_1, Q_2, \dots, Q_k has actually to be more sophisticated in some scenarios). Naturally, different scenarios arise according to the distribution of faulty links in Q_1, Q_2, \dots, Q_k and in dimension 1. Another complication is that the chosen Hamiltonian circuit in Q_2 , for example, might depend upon the Hamiltonian circuit chosen in Q_1 .

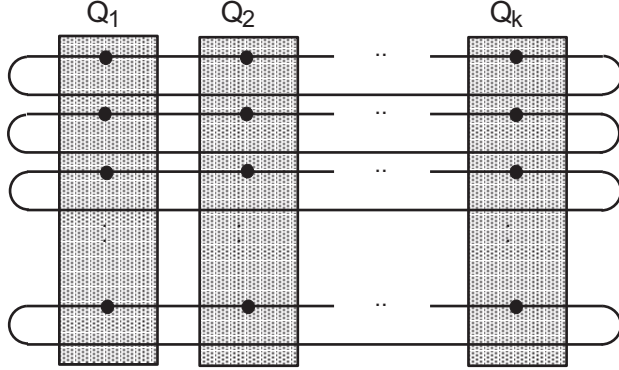


Figure 1. The k copies of Q_n^k .

Case (i) Each Q_i is such that every node is incident with at least 2 healthy links and no Q_i contains $4n - 4$ faults.

W.l.o.g. we may assume that Q_1 has most faults from amongst Q_1, Q_2, \dots, Q_k . Hence, each of Q_2, Q_3, \dots, Q_k has at most $2n - 2$ faults. By the induction hypothesis, Q_1 has a Hamiltonian circuit C_1 . Following our basic strategy, outlined above, we wish to find a Hamiltonian circuit C_k in Q_k or a Hamiltonian circuit C_2 in Q_2 so that we might “join” such a Hamiltonian circuit to C_1 using healthy links in dimension 1. By “join” we mean replace a link (x_1, y_1) of C_1 and the (corresponding) link (x_2, y_2) of C_2 , for example, with the links (x_1, x_2) and (y_1, y_2) in dimension 1. However, we must ensure that two mutually compatible links exist in C_1 and C_2 and also that the relevant dimension 1 links are healthy.

We begin by applying a counting argument to show that there exist links (x_1, y_1) and (y_1, z_1) of C_1 such that either:

- $(x_2, y_2), (y_2, z_2), (x_1, x_2), (y_1, y_2)$ and (z_1, z_2) are all healthy;

or

- $(x_k, y_k), (y_k, z_k), (x_1, x_k), (y_1, y_k)$ and (z_1, z_k) are all healthy.

Suppose that it were otherwise. Then there would exist at least $2\lfloor k^n/3 \rfloor$ faults not in Q_1 (split C_1 into groups of three consecutive vertices and look at the links on either side in dimension 1 and in Q_2 and Q_k). However, when $n \geq 2$ and $k \geq 4$ or when $n \geq 3$ and $k = 3$, we have that $2\lfloor k^n/3 \rfloor > 4n - 1$ which yields a contradiction. Hence, w.l.o.g. we may assume that there exist links (x_1, y_1) and (y_1, z_1) of C_1 such that $(x_2, y_2), (y_2, z_2), (x_1, x_2), (y_1, y_2)$ and (z_1, z_2) are all healthy.

What we need to do now is to show that there is a Hamiltonian circuit C_2 in Q_2 containing either (x_2, y_2) or (y_2, z_2) : we can then join C_1 and C_2 as described above. Suppose that it were otherwise. If necessary, mark some of the links of Q_2 incident with y_2 as faulty (that is, temporarily regard them as faulty) so that y_2 is incident

with at most 3 healthy links in Q_2 , two of which are always (x_2, y_2) and (y_2, z_2) . Consequently, as there were originally at most $2n - 2$ faulty links in Q_2 , there are now at most $4n - 5$ faulty links. However, in order to apply our induction hypothesis (and deduce that this amended Q_2 has a Hamiltonian circuit), we need that every node in (the amended) Q_2 is incident with at least 2 healthy links. Suppose that it were otherwise. Then there is a node w_2 incident with exactly 1 healthy link. This must have been because (y_2, w_2) was a healthy link in the original Q_2 and it was subsequently marked as faulty. Amend the marking of healthy links so that (w_2, y_2) is the third healthy link in the amended Q_2 . Note that in the amended marking every node is incident with at least 2 healthy links (because Q_2 originally had at most $2n - 2$ faults). Now we can apply the induction hypothesis and deduce that Q_2 has a Hamiltonian circuit C_2 containing either (x_2, y_2) or (y_2, z_2) (possibly both). No matter which, we can join C_2 to C_1 (as described above) to obtain a circuit D_2 containing every node of Q_1 and Q_2 (henceforth, we now treat those links of Q_2 which were temporarily marked as faulty as being healthy again).

All links of D_2 except for (x_1, x_2) and (y_1, y_2) are links in Q_1 or Q_2 . Hence, there is much potential to join D_2 , as above, to a Hamiltonian circuit in Q_3 or Q_k . Similarly to as before (by applying exactly the same counting argument), w.l.o.g. there exist two consecutive links (u_2, v_2) and (v_2, w_2) of $D_2 \cap Q_2$ such that the links (u_3, v_3) , (v_3, w_3) , (u_2, u_3) , (v_2, v_3) and (w_2, w_3) are healthy. Again, by arguing exactly as before, there is a Hamiltonian circuit C_3 in Q_3 containing either the link (u_3, v_3) or the link (v_3, w_3) ; and we can join D_2 to C_3 using links in dimension 1 to obtain a circuit D_3 containing all nodes of Q_1 , Q_2 and Q_3 . Exactly the same arguments apply so that we might extend D_3 to a circuit D_4 , containing all nodes of Q_1 , Q_2 , Q_3 and Q_4 , and so on until we obtain a Hamiltonian circuit in Q_{k+1}^n .

Case (ii) Each Q_i is such that every node is incident with at least 2 healthy links and some Q_j has exactly $4n - 4$ faults.

W.l.o.g. we may assume that $j = 1$. Suppose that there is some fault (x_1, y_1) of Q_1 such that (x_1, x_2) and (y_1, y_2) are healthy. Amend Q_1 so that (x_1, y_1) is temporarily marked as healthy. By the induction hypothesis applied to this amended Q_1 , there is a Hamiltonian circuit C_1 which may or may not contain (x_1, y_1) ; and C_1 is a circuit in the original Q_1 . The circuit C_1 has an isomorphic copy C_i in each Q_i , for $i = 2, 3, \dots, k$. If (x_1, y_1) is in C_1 , the circuit C_1 can be joined to C_2 using the healthy links (x_1, x_2) and (y_1, y_2) . Otherwise, because there are exactly 3 faults in dimension 1 and $\lfloor k^n/2 \rfloor > 3$, there is a link (u_1, v_1) of C_1 such that (u_1, u_2) and (v_1, v_2) are healthy (use a counting argument similar to that used before except split C_1 into groups of two consecutive vertices and look at the pairs of links in dimension 1 joining Q_1 to Q_2). C_1 can now be joined to C_2 using these links to yield a circuit D_2 containing every node of Q_1 and Q_2 . The circuit D_2 contains $k^n - 1$ links of Q_2 . As $\lfloor (k^n - 1)/2 \rfloor > 3$, the same argument yields that there is a link (u_2, v_2) of $D_2 \cap Q_2$ such that the links (w_3, z_3) , (w_2, w_3) and (z_2, z_3) are all healthy. Moreover, (w_3, z_3) lies on the circuit C_3 of Q_3 . Hence, we can join D_2 and C_3 to obtain a circuit D_3 containing every node of Q_1 , Q_2 and Q_3 . Exactly the same arguments apply so that we can extend D_3 to a Hamiltonian circuit of Q_{k+1}^n .

On the other hand, suppose that for every fault (x_1, y_1) of Q_1 , at least one of (x_1, x_2) and (y_1, y_2) , and at least one of (x_1, x_k) and (y_1, y_k) are faulty. Let (x_1, y_1) be some fault of Q_1 . As there are exactly 3 faults in dimension 1, it can not be the case that two faults in Q_1 are not incident with one another. Let us now count the maximum number μ of faults of Q_1 which could be incident with either x_1 or y_1 .

Consider x_1 . The number of faults incident with x_1 , apart from the fault (x_1, y_1) , is at most $2n - 3$. Similarly, the number of faults incident with y_1 , apart from the fault (x_1, y_1) , is at most $2n - 3$. Hence, $\mu \leq (2n - 3) + (2n - 3) + 1 = 4n - 5$. But there are $4n - 4$ faults in Q_1 and so we obtain a contradiction.

Case (iii) There exists some Q_i in which there is a node incident with exactly 1 healthy link in Q_i .

W.l.o.g. we may assume that the node x_1 in Q_1 is incident with exactly 1 healthy link, (x_1, y_1) , in Q_1 . As x_1 is incident with $2n - 1$ faults in Q_1 : each Q_i , for $i = 2, 3, \dots, k$, contains at most $2n - 3$ faults; there is no node in any Q_i , for $i = 2, 3, \dots, k$, which is incident with less than 3 healthy links in that Q_i ; and apart from x_1 , there is no other node in Q_1 which is incident with less than 2 healthy links in Q_1 . Also, as x_1 is incident with at least 2 healthy links in Q_{n+1}^k , we may suppose that (x_1, x_2) is healthy. Consider w_1 , one of the $2n - 1$ potential neighbours of x_1 in Q_1 for which the link (x_1, w_1) is faulty. There are two scenarios.

Case (iii)(a) (w_1, w_2) is a healthy link.

Mark the previously faulty link (x_1, w_1) as temporarily healthy. By the induction hypothesis applied to this amended Q_1 , there is a Hamiltonian path P_1 from x_1 to w_1 . Moreover, this Hamiltonian path P_1 is a Hamiltonian path in the original Q_1 (where the links temporarily marked as faulty resume their healthy status).

Mark some of the previously healthy links in Q_2 that are incident with x_2 as temporarily faulty and mark the link (x_2, w_2) as temporarily healthy (if necessary) so as to ensure that x_2 is incident with exactly 2 healthy links in this amended Q_2 (one of which is (x_2, w_2)). Note that in order to build this amended Q_2 we have introduced at most $2n - 2$ temporary faults; and so this amended Q_2 has at most $4n - 5$ faults and every node is incident with at least 2 healthy links. Hence, by the induction hypothesis, there exists a Hamiltonian path P_2 in this amended Q_2 from x_2 to w_2 . Moreover, this Hamiltonian path P_2 is a Hamiltonian path in the original Q_2 . Join P_1 and P_2 using the healthy links (x_1, x_2) and (w_1, w_2) to form a circuit D_2 which contains all nodes of Q_1 and Q_2 .

Applying a counting argument similar to that used in Case (ii), along with the fact that $\lfloor (k^n - 1)/2 \rfloor > 2n$ (note that the total number of faults in Q_{n+1}^k not contained in Q_1 is at most $2n$), there exists a link (u_2, v_2) of $D_2 \cap Q_2$ such that the links (u_3, v_3) , (u_2, u_3) and (v_2, v_3) are healthy. Temporarily mark healthy links in Q_3 incident with u_3 as faulty so that in this amended Q_3 , u_3 is incident with exactly 2 healthy links one of which is (u_3, v_3) . In order to build this amended Q_3 we have introduced at most $2n - 2$ temporary faults; and so this amended Q_3 has at most $4n - 5$ faults and every node is incident with at least 2 healthy links. By the induction hypothesis, there is a Hamiltonian circuit C_3 in the original Q_3 containing the link (u_3, v_3) . We can join D_2 and C_3 , using the healthy links (u_2, u_3) and (v_2, v_3) , to obtain a circuit D_3 containing every node of Q_1 , Q_2 and Q_3 . Exactly the same argument can be applied to extend D_3 to a circuit D_4 , and so on until we have a Hamiltonian circuit of Q_{n+1}^k .

Case (iii)(b) All links from every such w_1 to its corresponding node w_2 in Q_2 are faulty.

This accounts for another $2n - 1$ faults in Q_{n+1}^k . Also, if (x_1, x_k) is healthy then by symmetry we are in Case (iii)(a) (as all but at most 1 link of the form (w_1, w_k) is healthy). Hence, we may assume that (x_1, x_k) is faulty and this accounts for all the faults in Q_{n+1}^k .

Consequently, (y_1, y_2) and (y_1, y_k) are both healthy links (recall that (x_1, y_1) is the only healthy link of Q_1 incident with x_1). Let w_1 be some potential neighbour of x_1 in Q_1 for which the link (x_1, w_1) is faulty. Amend Q_1 by marking the link (x_1, w_1) as temporarily healthy. By the induction hypothesis applied to this amended Q_1 , there is a Hamiltonian path P_1 in the original Q_1 from x_1 to w_1 . Rename the nodes of P_1 as $x_{1,1} = x_1, x_{1,2} = y_1, x_{1,3}, \dots, x_{1,k^n} = w_1$, and note that in each Q_i , $i \geq 2$, there is a corresponding Hamiltonian path P_i which can be extended to a Hamiltonian circuit C_i of Q_i (as (x_i, w_i) is healthy in Q_i). Rename the nodes of C_i as $x_{i,1} = x_i, x_{i,2} = y_i, x_{i,3}, \dots, x_{i,k^n} = w_i$, for each $i \geq 2$.

For ease of notation, denote k^n by m . Suppose k is even. Then the following is a Hamiltonian circuit in Q_{n+1}^k :

$$\begin{aligned} & (x_{1,1}, x_{2,1}, \dots, x_{k,1}, x_{k,2}, x_{k,3}, x_{1,3}, x_{1,4}, \dots, x_{1,m}, x_{k,m}, x_{k-1,m}, \dots, x_{2,m}, \\ & x_{2,m-1}, x_{3,m-1}, \dots, x_{k,m-1}, x_{k,m-2}, x_{k-1,m-2}, \dots, x_{2,m-2}, x_{2,m-3}, \\ & x_{3,m-3}, \dots, x_{k,m-3}, x_{k,m-4}, \dots, x_{k,4}, x_{k-1,4}, \dots, x_{2,4}, x_{2,3}, x_{3,3}, \dots, \\ & x_{k-1,3}, x_{k-1,2}, x_{k-2,2}, \dots, x_{2,2}, x_{1,2}, x_{1,1}) \end{aligned}$$

(see Fig. 2 where some of the healthy links between the Q_i 's are shown and bold links denote the links of the Hamiltonian circuit). If k is odd then the following is a Hamiltonian circuit in Q_{n+1}^k :

$$\begin{aligned} & (x_{1,1}, x_{2,1}, \dots, x_{k,1}, x_{k,2}, x_{k-1,2}, \dots, x_{2,2}, x_{2,3}, x_{3,3}, \dots, x_{k,3}, x_{k,4}, x_{k-1,4}, \dots, \\ & x_{2,4}, x_{2,5}, \dots, x_{2,m}, x_{3,m}, \dots, x_{2,m}, x_{k,m}, \dots, x_{1,m}, x_{1,m-1}, \dots, x_{1,2}, x_{1,1}) \end{aligned}$$

(see Fig. 3).

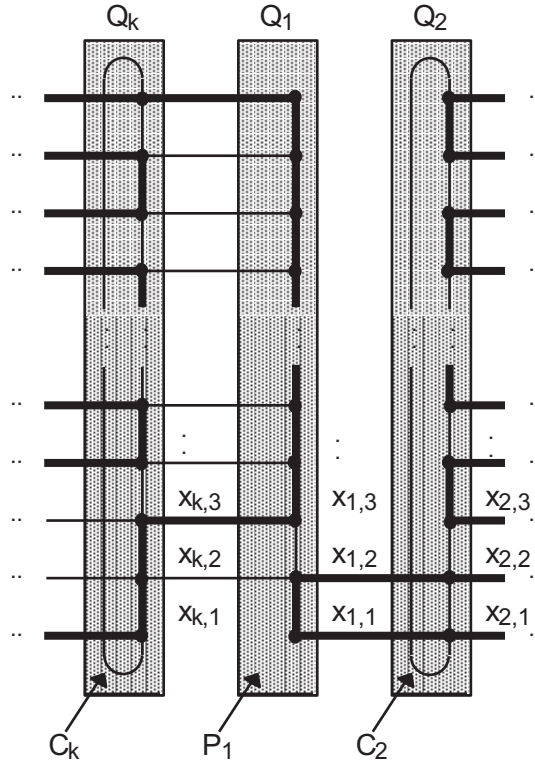


Figure 2. The Hamiltonian circuit when k is even.

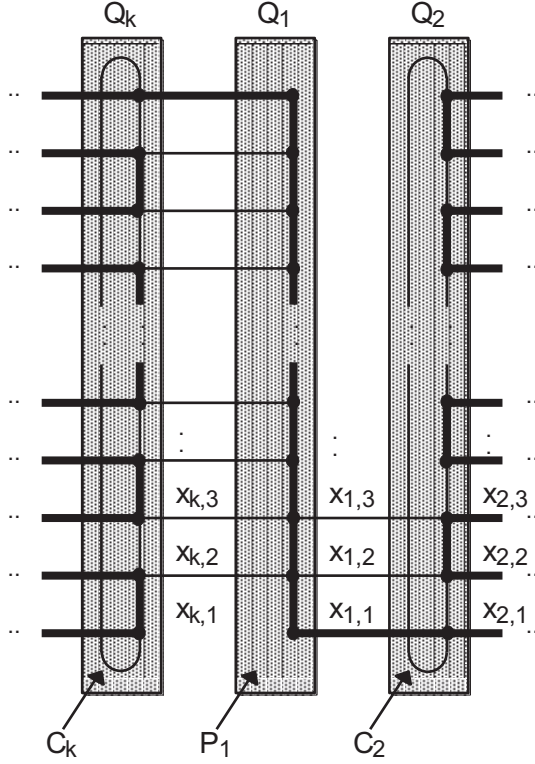


Figure 3. The Hamiltonian circuit when k is odd.

Case (iv) There exists some Q_i in which there is a node incident with no healthy links in Q_i .

W.l.o.g. we may assume that x_1 is incident with no healthy links in Q_1 . As x_1 is incident with at least 2 healthy links in Q_{n+1}^k , the links (x_1, x_2) and (x_1, x_k) must be healthy. There are at least $2n$ faults in Q_1 , and so there must be at most $2n - 4$ faults distributed amongst Q_2, Q_3, \dots, Q_k . Hence, apart from x_1 , there are no nodes which are incident with less than 4 healthy links in their respective copy of Q_n^k .

The node x_1 has $2n$ potential neighbours in Q_1 . Each of these potential neighbours is incident with a potential dimension 1 link to Q_1 and a potential dimension 1 link to Q_k (these dimension 1 links might be faulty). As there are at most $2n - 1$ faults in dimension 1, there must exist potential neighbours y_1 and z_1 of x_1 such that the links (y_1, y_2) and (z_1, z_k) are healthy (partition the potential neighbours into n pairs $\{y_1, z_1\}$ and look at the pairs of dimension 1 links $\{(y_1, y_2), (z_1, z_k)\}$ and $\{(y_1, y_k), (z_1, z_2)\}$). Mark the faulty links (x_1, y_1) and (x_1, z_1) as temporarily healthy in Q_1 . Applying the induction hypothesis to this amended Q_1 , we obtain a path P_1 in the original Q_1 from y_1 to z_1 upon which every node of Q_1 appears exactly once, except for x_1 which does not appear at all.

By marking previously healthy links in Q_2 that are incident with x_2 as temporarily faulty, and by marking the link (x_2, y_2) as temporarily healthy (if necessary), ensure that x_2 is incident with exactly 2 healthy links in this amended Q_2 , one of which is (x_2, y_2) . This involves introducing at most $2n - 2$ temporary faults into Q_2 ; and so the amended Q_2 has at most $4n - 6$ faults and every node is incident with at least two healthy links. The induction hypothesis yields that there is a Hamiltonian path from x_2 to y_2 in the original Q_2 . Likewise, there is a Hamiltonian path from x_k to

z_k in Q_k . Hence, let D_2 be the circuit obtained by joining P_1 , P_2 and P_k using the healthy links (x_1, x_2) , (y_1, y_2) , (x_1, x_k) and (z_1, z_k) .

Applying a counting argument similar to that used in Case (ii), along with the fact that $\lfloor (k^n - 1)/2 \rfloor > 2n - 1$ (note that the total number of faults in Q_{n+1}^k not contained in Q_1 is at most $2n - 1$), there exists a link (u_2, v_2) of $D_2 \cap Q_2$ such that the links (u_3, v_3) , (u_2, u_3) and (v_2, v_3) are healthy. Temporarily mark healthy links in Q_3 incident with u_3 as faulty so that in this amended Q_3 , u_3 is incident with exactly two healthy links one of which is (u_3, v_3) . In order to build this amended Q_3 we have introduced at most $2n - 2$ temporary faults; and so this amended Q_3 has at most $4n - 6$ faults and every node is incident with at least 2 healthy links. By the induction hypothesis, there is a Hamiltonian circuit C_3 in the original Q_3 containing the link (u_3, v_3) . We can join D_2 and C_3 , using the healthy links (u_2, u_3) and (v_2, v_3) , to obtain a circuit D_3 containing every node of Q_k , Q_1 , Q_2 and Q_3 . Exactly the same argument can be applied to extend D_3 to a circuit D_4 , and so on until we have a Hamiltonian circuit of Q_{n+1}^k .

It remains to show that the result holds for the base cases of the induction; namely, when $n = 2$ and $k \geq 4$, and when $n = 3$ and $k = 3$.

LEMMA 2.2. *If Q_2^k , where $k \geq 4$, has 3 faulty links and is such that every node is incident with at least 2 healthy links then Q_2^k has a Hamiltonian circuit.*

Proof. There exists some dimension, say dimension 1, that contains at least 2 faults. Partition Q_2^k over dimension 1 to obtain k copies of Q_1^k , namely Q_1, Q_2, \dots, Q_k .

Case (i) All faults are in dimension 1.

Consider the circuit Q_1 of length k . As there are 3 faults in dimension 1, w.l.o.g. there exists an edge (x_1, y_1) of Q_1 such that the links (x_1, x_2) and (y_1, y_2) are both healthy (apply our usual counting argument). Join Q_1 and Q_2 using these links to obtain a circuit D_2 containing every node of Q_1 and Q_2 . By proceeding as we have done throughout, the same argument can be used to extend D_2 to (w.l.o.g.) a circuit D_3 , and so on until we obtain a Hamiltonian circuit of Q_2^k .

Case (ii) Dimension 1 has exactly two faults.

W.l.o.g. the only fault not in dimension 1 may be assumed to be (x_1, y_1) in Q_1 . If the links (x_1, x_2) and (y_1, y_2) are both healthy or the links (x_1, x_k) and (y_1, y_k) are both healthy then we can join Q_1 with Q_2 or Q_k , respectively, as in Case (i), and extend this circuit to a Hamiltonian circuit of Q_2^k .

Hence, w.l.o.g. we may assume that the links (x_1, x_2) and (y_1, y_k) are both faulty. If k is even then there exists a Hamiltonian circuit in Q_2^k as pictured in Fig. 2 (in that picture, $x_{1,3}$, $x_{1,2}$, $x_{2,3}$ and $x_{k,2}$ play the roles of x_1 , y_1 , x_2 and y_k , respectively). If k is odd then there exists a Hamiltonian circuit in Q_2^k as pictured in Fig. 3 (in that picture, $x_{1,m}$, $x_{1,1}$, $x_{2,m}$ and $x_{k,1}$ play the roles of x_1 , y_1 , x_2 and y_k , respectively). \square

LEMMA 2.3. *If Q_2^3 has 3 faulty links and is such that every node is incident with at least 2 healthy links then Q_2^3 has a Hamiltonian circuit unless these 3 faulty links form a circuit of length 3.*

Proof. There exists some dimension, say dimension 1, that contains at least 2 faults. Partition Q_2^3 over dimension 1 to obtain 3 copies of Q_1^3 , namely Q_1 , Q_2 and Q_3 . We may assume that either Q_1 contains 1 fault or all faults are in dimension 1. Denote the nodes of Q_i by x_i , y_i and z_i , for $i = 1, 2, 3$.

Case (i) Q_1 contains 1 fault.

W.l.o.g. we may assume that the fault in Q_1 is (x_1, y_1) .

Case (i)(a) The links (x_1, x_2) and (y_1, y_2) are healthy.

Form the circuit $C = (x_1, z_1, y_1, y_2, z_2, x_2, x_1)$ in Q_2^3 . There are 2 possibilities: either one of the sets of pairs

$$\{(x_1, x_3), (z_1, z_3)\}, \{(y_1, y_3), (z_1, z_3)\}, \{(x_2, x_3), (z_2, z_3)\}, \{(y_2, y_3), (z_2, z_3)\}$$

consists of 2 healthy links or the faulty links in dimension 1 are (z_1, z_3) and (z_2, z_3) . In the former case, the circuit C can be joined to the circuit (x_3, y_3, z_3, x_3) using the pair of healthy links to obtain a Hamiltonian circuit in Q_2^3 : in the latter case, we can define our Hamiltonian circuit in Q_2^3 to be $(x_1, z_1, z_2, y_2, y_1, y_3, z_3, x_3, x_2, x_1)$.

Case (i)(b) At least one of the links (x_1, x_2) and (y_1, y_2) is faulty.

By symmetry, we may also assume that at least one of (x_1, x_3) and (y_1, y_3) is faulty (as otherwise we are in Case (i)(a)); so this accounts for all faults in Q_2^3 . The only configuration possible, up to isomorphism, is that in Fig. 4(a), and so there is a Hamiltonian circuit as depicted in that figure. (In Fig. 4(a), the nodes x_1, y_1 and z_1 of Q_1 form the central column, with the other 2 columns similarly depicting the nodes of Q_2 and Q_3 . Faults are denoted by missing links and links of the Hamiltonian circuit are drawn in bold.)

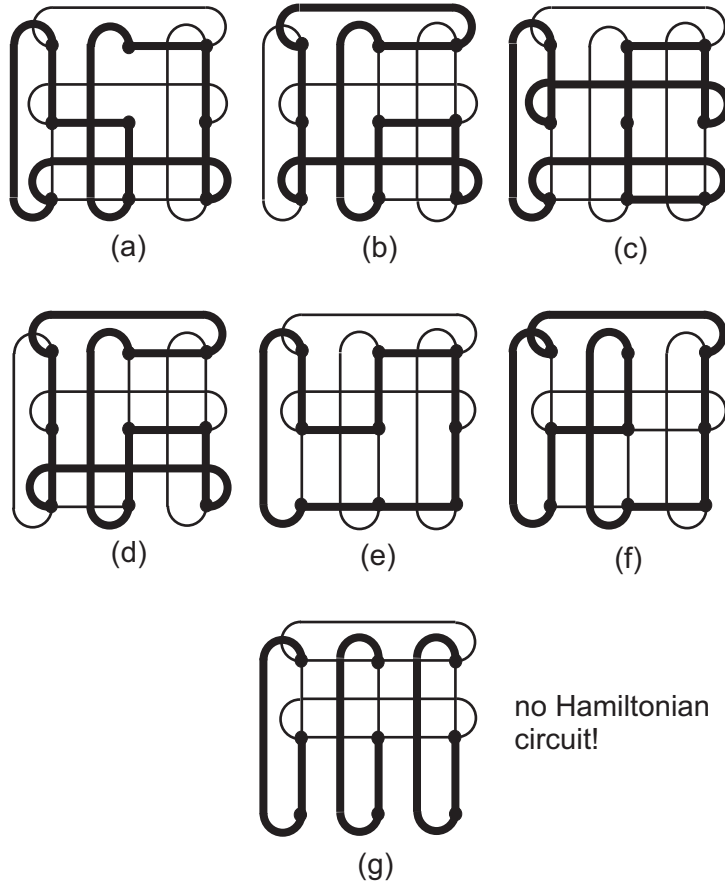


Figure 4. The different configurations for Q_2^3 .

Case (ii) All faults are in dimension 1.

Up to isomorphism, there are 6 different configurations possible, shown in Fig. 4(b–g), with Hamiltonian circuits as depicted except for Fig. 4(g) where no such Hamiltonian circuit exists (in Fig. 4(g), w.l.o.g. the bold links are necessarily in any Hamiltonian circuit, if there were to exist one; and one can immediately see that there is no extension of these bold links to a Hamiltonian circuit). \square

LEMMA 2.4. *If Q_3^3 has 7 faulty links and is such that every node is incident with at least 2 healthy links then Q_3^3 has a Hamiltonian circuit.*

Proof. Case (i) Q_3^3 contains faults forming a circuit C of length 3.

All of the faults in C must appear in the same dimension, dimension 1 say. Partition Q_3^3 across dimension 1 to obtain 3 copies of Q_2^3 , namely Q_1 , Q_2 and Q_3 , and let the faulty links in C be (x_1, x_2) , (x_2, x_3) and (x_3, x_1) . We may assume that Q_1 contains the most faults amongst these copies, then Q_2 and then Q_3 .

Case (i)(a) Q_1 contains faults forming a circuit D of length 3.

Let y_1 and z_1 be nodes of D different from x_1 (x_1 may or may not be on D) so that the number of faults incident with y_1 is no greater than the number of faults incident with any node of D different from x_1 (note that x_1 is incident with at most 2 faults in Q_1). If y_1 is incident with 1 healthy link in Q_1 then every other node of Q_1 is incident with at least 2 healthy links in Q_1 (as Q_3^3 has 7 faults, y_1 must be incident with at least 1 healthy link in Q_1). In this case, temporarily mark the link (y_1, z_1) as healthy so that there are at most 3 faults in the amended Q_1 (and these faults do not form a circuit). Lemma 2.3 yields that there is a Hamiltonian path in the original Q_1 from y_1 to z_1 .

If y_1 is incident with 2 healthy links in Q_1 then every node in Q_1 is incident with at least 2 healthy links in Q_1 . Mark the link (y_1, z_1) as temporarily healthy and a healthy link of Q_1 incident with y_1 as temporarily faulty. Every node in the amended Q_1 is incident with at least 2 healthy links and there are at most 3 faults (and these faults do not form a circuit). Lemma 2.3 yields that there is a Hamiltonian path in the original Q_1 from y_1 to z_1 .

Whichever of the above scenarios applies, denote the Hamiltonian path in Q_1 from y_1 to z_1 by P_1 . The faults in Q_1 and the faults (x_1, x_2) , (x_2, x_3) and (x_3, x_1) account for at least 6 of the 7 faults in Q_3^3 . Hence, w.l.o.g. we may assume that the links (x_1, x_2) and (y_1, y_2) are healthy. There is at most 1 fault in Q_2 . By marking healthy links of Q_2 as temporarily faulty (if necessary), ensure that (y_2, z_2) is healthy and y_2 is incident with exactly 2 healthy links. Applying Lemma 2.3 to this amended Q_2 yields that there is a Hamiltonian circuit C_2 (that is also a Hamiltonian circuit in the original Q_2) including the link (y_2, z_2) . Join P_1 and C_2 using the healthy links (y_1, y_2) and (z_1, z_2) to obtain a circuit D_2 containing every node of Q_1 and Q_2 .

Q_3 has an isomorphic copy C_3 of C_2 and there are no faults in Q_3 . As C_3 has length 9 and there are at most 4 faults in dimension 1, by applying our counting argument as we have done throughout, we can join D_2 and C_3 using appropriate dimension 1 links to obtain a Hamiltonian circuit in Q_3^3 .

Case (i)(b) Q_1 does not contain faults forming a circuit D of length 3.

Note that the proofs of Cases (i), (ii), (iii) and (iv) of the main theorem hold for Q_3^3 except that: throughout, instead of appealing to an inductive hypothesis, we use Lemma 2.3; in Case (i), we assume that dimension 1 contains at most 5 faults; and in Case (iii(a)), when amending Q_2 we must ensure that we do not introduce a circuit of faults of length 3 (this can be done as Q_2 has at most 1 fault). Consequently, we are left with one scenario to consider: the sub-case of Case (i) when each Q_i is such that

every node is incident with at least 2 healthy links and when dimension 1 contains 6 or 7 faults.

Let (a new) 3-ary 2-cube Q_2^3 be such that there is a fault (x, y) in Q_2^3 if and only if there is a fault (x_i, y_i) in Q_i , for some $i \in \{1, 2, 3\}$. Then Q_2^3 has at most 2 faults and by Lemma 2.3, it has a Hamiltonian circuit C . For each $i \in \{1, 2, 3\}$, let C_i be the isomorphic copy of C in Q_i (note that each C_i consists entirely of healthy links). Even if dimension 1 (of our original Q_3^3) contains 7 faults, our usual counting argument yields that there exists a pair of healthy links $\{(u_1, u_2), (v_1, v_2)\}$ or $\{(u_1, u_3), (v_1, v_3)\}$, where (u_1, v_1) is a link of C_1 : w.l.o.g. we may assume that these healthy links are (u_1, u_2) and (v_1, v_2) . We can join C_1 and C_2 using these healthy links and then proceed similarly to join the resulting circuit to C_3 and obtain a Hamiltonian circuit of Q_3^3 .

Case (ii) Q_3^3 does not contain faults forming a circuit of length 3.

There exists a dimension, dimension 1 say, containing at least 3 faults. Partition Q_3^3 across dimension 1 to obtain 3 copies of Q_2^3 , namely Q_1, Q_2 and Q_3 . Let Q_1 contain the most faults amongst these copies, then Q_2 and then Q_3 . Proceeding as in Case (i)(b) yields the result. \square

The main theorem now follows by induction. \square

The result in Theorem 2.1 is optimal in the following sense. Let a, b, c and d be 4 nodes in Q_n^k , where $k \geq 4$ and $n \geq 2$, or $k = 3$ and $n \geq 3$, such that there are links $(a, b), (b, c), (c, d)$ and (d, a) . Let the faults of Q_n^k consist of those links incident with a that are different from (a, b) and (a, d) , and those links incident with c that are different from (b, c) and (c, d) . In particular, Q_n^k has $4n - 4$ faults and every node is incident with at least 2 healthy links; but this faulty Q_n^k does not contain a Hamiltonian circuit as any Hamiltonian circuit necessarily contains the links (a, b) and (a, d) , and also the links (c, b) and (c, d) , which yields a contradiction.

3. Conclusions. We have proven that every k -ary n -cube Q_n^k which has at most $4n - 5$ faulty links and is such that every node is incident with at least 2 healthy links has a Hamiltonian circuit. As mentioned earlier, an analogous result for hypercubes was proven by Chan and Lee [8]. In [8], it was also shown that the problem of deciding whether a faulty binary n -cube has a Hamiltonian circuit is NP-complete. Their complexity-theoretic reduction (from the 3-Satisfiability Problem) can easily be adapted to show that the problem of deciding whether a faulty k -ary n -cube has a Hamiltonian circuit is also NP-complete (we leave the proof of this as a simple exercise).

As open problems relating to the research in this paper, we propose the following. The construction of our Hamiltonian circuits in our faulty k -ary n -cubes does not yield efficient parallel distributed algorithms for actually building the Hamiltonian circuits. For example, suppose one had a parallel computer whose underlying interconnection network was a k -ary n -cube and each node, i.e., processor, had local (or even global) knowledge of the faulty links. How could we develop an efficient message-passing algorithm so that upon termination, every node knew its successor and predecessor on a Hamiltonian circuit (without necessarily knowing the Hamiltonian circuit in its entirety)? Such an algorithm would be extremely useful. Also, whilst we provide a precise result as to the threshold value on the number of faulty links occurring in a k -ary n -cube so that there still exists a Hamiltonian circuit (under the assumption that every node is incident with at least 2 healthy links) and we also remark that the general decision problem is NP-complete, it would be useful if ‘safe patterns’ of faults

could be established so that even though there were more than $4n - 5$ faulty links present, one could still be sure of the existence of a Hamiltonian circuit because these faults were arranged in some specific formation. Finally, we have only addressed the problem of finding longest circuits in k -ary n -cubes in the presence of faulty links. It would be interesting to do likewise in the presence of faulty nodes; or even faulty nodes and links.

REFERENCES

- [1] Y.A. Ashir and I.A. Stewart, On embedding cycles in k -ary n -cubes, *Parallel Processing Letters* **7** (1997) 49–55.
- [2] Y.A. Ashir, I.A. Stewart and A. Ahmed, Communication algorithms in k -ary n -cube interconnection networks, *Information Processing Letters* **61** (1997) 43–48.
- [3] D. Bertsekas and J. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*, Prentice-Hall (1989).
- [4] S. Bettayeb, On the k -ary hypercube, *Theoretical Computer Science* **140** (1995) 333–339.
- [5] L. Bhuyan and D. Agrawal, Generalized hypercube and hyperbus structures for a computer network, *IEEE Transactions on Computers* **C-33** (1984) 323–333.
- [6] B. Bose, B. Broeg, Y. Kwon and Y. Ashir, Lee distance and topological properties of k -ary n -cubes, *IEEE Transactions on Computers* **44** (1995) 1021–1030.
- [7] R.B. Broeg, *Topics in Toroidal Interconnection Networks*, Ph.D. Thesis, Oregon State University (1995).
- [8] M.Y. Chan and S.-J. Lee, On the existence of Hamiltonian circuits in faulty hypercubes, *SIAM Journal of Discrete Mathematics* **4** (1991) 511–527.
- [9] W. Dally, Performance analysis of k -ary n -cube interconnection networks, *IEEE Transactions on Computers* **39** (1990) 775–785.
- [10] F.T. Leighton, *Introduction to Parallel Algorithms and Architectures: Arrays. Trees. Hypercubes*, Morgan Kaufmann (1992).
- [11] D. Linder and J. Harden, An adaptive and fault tolerant wormhole routing strategy for k -ary n -cubes, *IEEE Transactions on Computers* **40** (1991) 2–12.
- [12] Y. Saad and M. Schultz, Data communications in hypercubes, *Journal of Parallel and Distributed Computing* **6** (1989) 115–135.