Hamiltonicity and Fault Tolerance
in the k-ary n-cube

By Clifford R. Haithcock

Portland State University
Department of Mathematics and Statistics
2006

In partial fulfillment of the requirements of the degree of
Master's of Science in Mathematics

Advisor: John S. Caughman, IV

Contents

Chapter 1

Introduction

This paper is about Hamiltonian cycles in the k-ary n-cube. While it is known that the k-ary n-cube Q_n^k is Hamiltonian for $(k, n) \neq (2, 1)$, we will consider the Hamiltonicity of Q_n^k in the presence of missing edges. An obvious necessary condition for such a subgraph of Q_n^k to be Hamiltonian is that every vertex have degree at least two. Assuming this minimal condition is met, it is natural to ask how many faults Q_n^k can sustain and still *necessarily* contain a Hamiltonian cycle. In their 2002 paper [1], Ashir and Stewart establish a sharp bound on this number of faults. They prove the following theorem.

Theorem 1 Suppose $k \geq 4$ and $n \geq 2$, or $k = 3$ and $n \geq 3$. If Q_n^k has no more than $4n-5$ faults and every node in Q_n^k is incident to at least two healthy links, then Q_n^k contains a Hamiltonian cycle.

In this paper, we will work through the details of Ashir and Stewart's result. They apply an inductive argument in which the induction step is broken into several cases. We expand upon their explanation and provide a number of diagrams to illuminate the argument.

The motivation for such problems comes from analyzing the fault tolerance of massively parallel computers. A key factor in the architecture of a parallel computer is the communication network. The natural model for this network is a graph with vertices representing the computational units and edges representing the communication links. Fundamental results force us to only consider communication networks based on highly structured graphs that we can more fully analyze. Typical networks include rings, trees, meshes, and hypercubes in addition to the k-ary n-cube. The development of parallel algorithms depends upon the underlying communication network. Porting an existing algorithm to a new architecture involves embedding the original communication network within the new architecture's network (the guest and host networks respectively). In particular, the existence of a Hamiltonian cycle demonstrates the possibility of embedding a ring in Q_n^k

Chapter 2

Preliminaries

2.1 Basic Definitions

- **Graph**

  A graph is an ordered pair of disjoint sets (V, E) such that $E \subset V^{(2)}$, which denotes the set of unordered pairs of V. We refer to the elements of V as vertices and the elements of E as edges. An edge $\{x, y\}$ is said to join the vertices x and y. We use the notation $V(G)$ and $E(G)$ to refer to the sets of vertices and edges of G respectively. We alternatively use nodes and links as synonyms for vertices and edges respectively. These terms come from the application of parallel computing.

- **Deletion ($G - e$)**

  Given a graph G and an edge $e \in E(G)$, we define $G - e$ to be the graph with vertex set $V(G)$ and edge set $E(G) - e$).

- **Graph Instersection ($G \cap H$)**

  Given two graphs G and H, we define the intersection $G \cap H = (V(G) \cap V(H), E(G) \cap E(H))$.

- **Subgraph**

  Given graphs G and H, we say H is a subgraph of G if $V(H) \subset V(G)$ and $E(H) \subset E(G)$.

- **Path**

  A path P is a graph of the form $V(P) = \{x_1, x_2, \cdots, x_n\}$ and $E(P) = \{\{x_1, x_2\}, \{x_2, x_3\}, \cdots, \{x_{n-1}, x_n\}\}$. We refer to the vertices x_1 and x_n as the endpoints of P, and we make use of the notation $P : x_1 \sim x_n$ to indicate that P is a path with endpoints x_1 and x_n.

- **Concatenating Paths**

  If we have two paths $P : x \sim y$ and $R : y \sim z$ where $V(S) \cap V(R) = \{y\}$ then we define $P \cdot R \equiv R \cdot P$ to be the concatenation of the two paths. That is, $V(P \cdot R) = V(R \cdot P) = V(P) \cup V(R)$ and $E(P \cdot R) = E(P) \cup E(R)$. Thus we have $R \cdot P : x \sim z$.

- **Cycle**

  A cycle C is a graph of the form $V(C) = \{x_1, x_1, \cdots, x_n\}, n \geq 3$ and $E(C) = \{\{x_1, x_2\}, \{x_2, x_3\}, \cdots, \{x_{n-1}, x_n\}, \{x_n, x_1\}\}$.

- **n-cycle**

  An n-cycle is a cycle with n vertices.

- **Hamiltonian Path**

  Given a graph G and a path P, we say P is a Hamiltonian path in G if P is a subgraph of G and $V(P) = V(G)$.

- **Hamiltonian Cycle**

  Given a graph G and a cycle C, we say C is a Hamiltonian cycle in G if C is a subgraph of G and $V(C) = V(G)$.

- **Incidence**

  We say two edges are incident if they share an endpoint.

  Furthermore, we say an edge e and a vertex x are incident to one another if x is an endpoint of e.

  It will be clear from context which definition we are applying.

- **Adjacency**

  Given a graph G and vertices x and y, we say x and y are adjacent if $\{x, y\} \in E(G)$.

- **Degree**

  Given a graph G and a vertex x, we say the degree of x, denoted $d(x)$, is the number of edges in $E(G)$ with x as an endpoint.

- **Bridge (2-bridge / 3-bridge)**

  A 2-bridge is a graph G of the form $V(G) = \{x_0, x_1, x_2, x_3\}$ and $E(G) = \{\{x_0, x_1\}, \{x_2, x_3\}, \{x_0, x_2\}, \{x_1, x_3\}\}$.

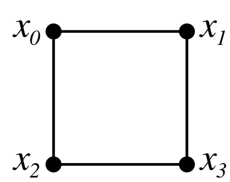

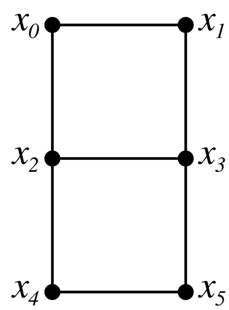

A 3-bridge is a graph G of the form $V(G) = \{x_0, x_1, x_2, x_3, x_4, x_5\}$ and $E(G) = \{\{x_0, x_1\}, \{x_2, x_3\}, \{x_4, x_5\}, \{x_0, x_2\}, \{x_2, x_4\}, \{x_1, x_3\}, \{x_3, x_5\}\}$.

- **Star**

  A star is a graph G of the form $V(G) = \{x_0, x_1, \cdots, x_n\}$ and $E(G) = \{\{x_0, x_1\}, \{x_0, x_2\}, \cdots, \{x_0, x_n\}\}$. We refer to the node x_0 as the center of the star.

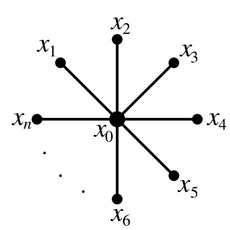

2.2 The Graph Q_n^k

Given positive integers k and n, we define the graph Q_n^k known as the k-ary n-cube. The vertices are n-tuples in which each coordinate is an integer between 0 and $k-1$. Two vertices are adjacent if the two n-tuples differ in exactly one coordinate and that difference equals 1 mod k. Figure 2.1 shows all of the vertices adjacent to $(0, 1, 2)$ in the graph Q_3^3.

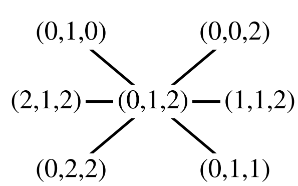

Figure 2.1: Vertices adjacent to $(0, 1, 2)$ in Q_3^3

Note that every vertex in Q_n^k has degree $2n$. If x, y are adjacent, then they differ in exactly one coordinate i. We then say the edge $\{x, y\}$ **lies in dimension** i. We can partition Q_n^k along a dimension i to get k copies of Q_{n-1}^k connected by dimension i links. The vertices of the j^{th} copy are the n-tuples whose i^{th} coordinate equals j.

2.3 Conventions / Notation

- **Faults in a graph**

  The focus of this paper is on the existence of Hamiltonian cycles within Q_n^k in the presence of faulty links. Throughout this work, there are numerous graphs depicted visually. Any given edge may be known to be faulty, known to be not faulty (i.e. healthy), or its state unknown. If an edge is known to be faulty, we will draw that edge with dashed lines. We do not consider faulty edges to be deleted from a graph; we merely note them as faulty. We do not visually distinguish edges of unknown state from healthy edges. When we refer to a graph H as healthy, we understand that all of the edges in H are healthy.

- **Indexing Conventions**

  Furthermore, we will partition the graph Q_n^k along some dimension. We will label the Q_{n-1}^k subgraphs $Q_1, Q_2, \ldots, Q_k$. We then strictly use subscripts in accordance with these labels. Thus by the subscripts it is understood that the nodes x_1, y_2, and w_3 are in Q_1, Q_2, and Q_3 respectively. We also understand the nodes x_1 and x_2 are adjacent via an edge in dimension 1. Similarly, the nodes y_k and y_1 are adjacent via an edge in dimension 1.

- **Joining cycles**

  We will make such frequent use of the following constructions that it is worth covering them as a concept. Our general goal is to establish the existence of Hamiltonian cycles. We will accomplish this by successively joining cycles connected via a 2-bridge to get an ever larger cycle as depicted in Figure 2.3. Joining a path and a cycle or joining two paths is a clear extension of the given example.

2.4 The Theorem of Ashir and Stewart

The theorem of Ashir and Stewart is concerned with guaranteeing the existence of Hamiltonian cycles in Q_n^k under conditions of faulted links. Specifically they give a sharp result on the number of faults Q_n^k can sustain and still maintain the necessity of a Hamiltonian cycle.

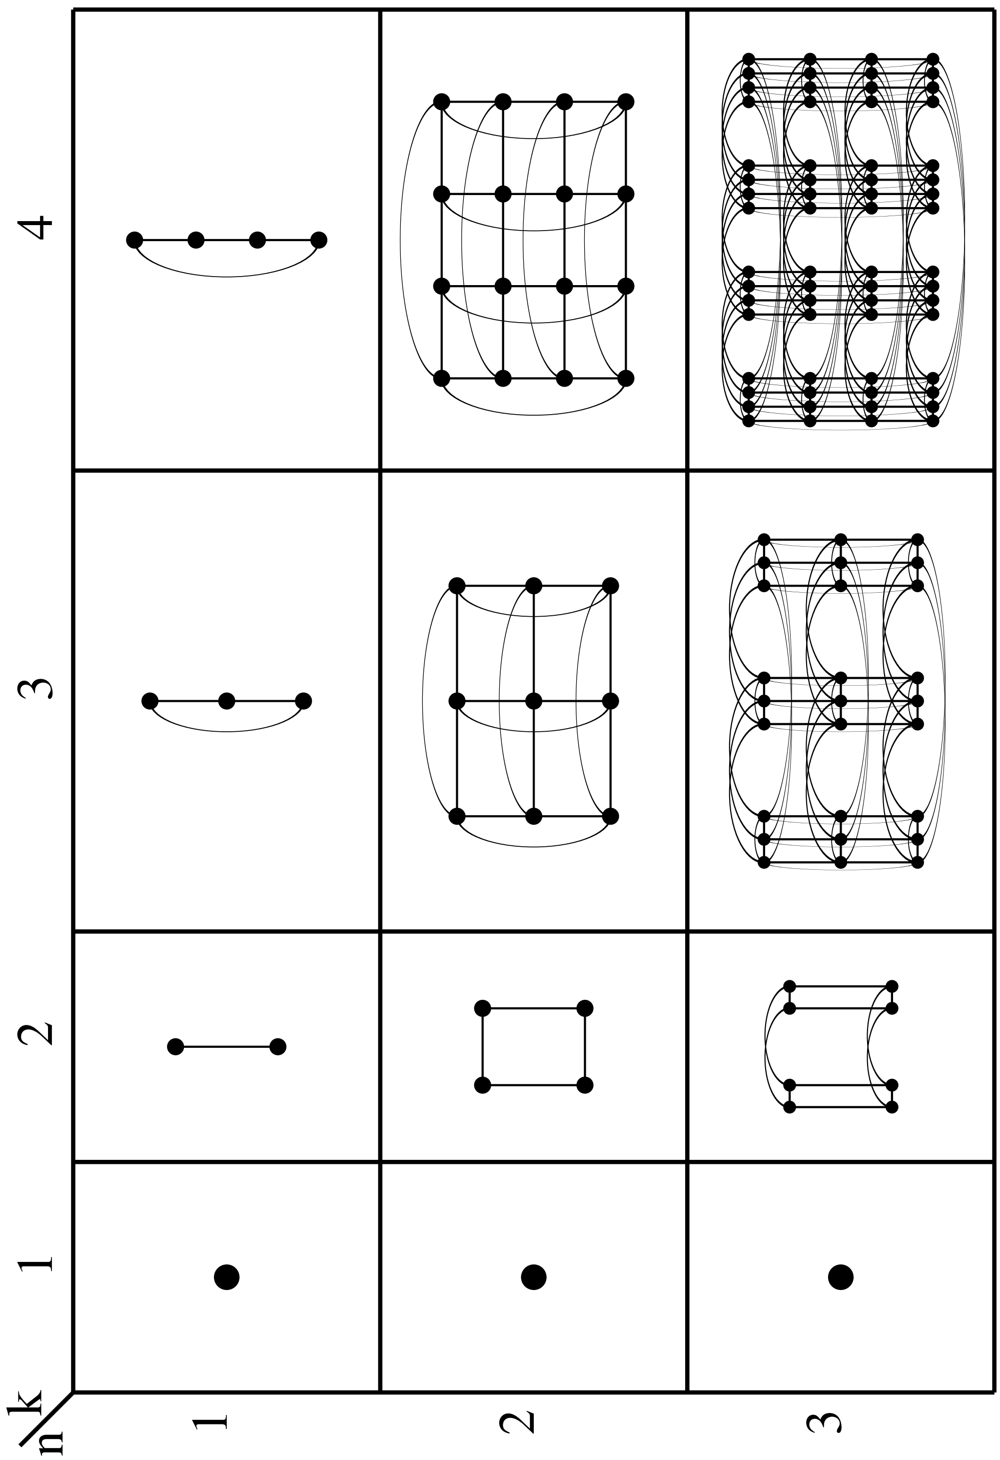

Figure 2.2: Examples of Q_n^k

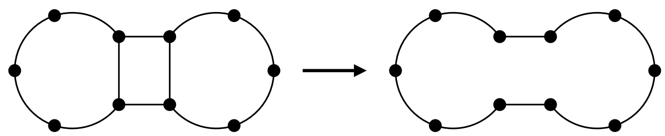

Figure 2.3: Joining Cycles

Theorem 1 *For $k \geq 4$ and $n \geq 2$ or $k = 3$ and $n \geq 3$, if Q_n^k has no more than $4n - 5$ faults and every node in Q_n^k is incident to at least two healthy links, then Q_n^k contains a Hamiltonian cycle.*

The hypothesis that every node is incident to at least two healthy links is a minimal condition for a graph to contain a Hamiltonian cycle. The proof follows by induction on n. This leaves as base cases the graph Q_3^3 and the family of graphs Q_2^k for $k \geq 4$ (see Figure 2.4). Chapter 3 proves the induction step. Chapter 4 covers the base case Q_2^k for $k \geq 4$, and chapter 6 addresses the base case Q_3^3.

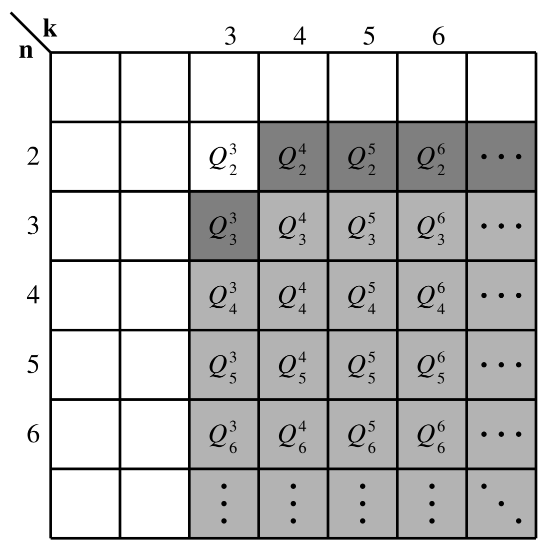

Figure 2.4: Inductive Structure

Chapter 3

The Induction Step

3.1 Argument Overview

To access the induction hypothesis, we partition the graph Q^k_{n+1} across a dimension to get k copies of Q^k_n.

Assume we have been given Q^k_{n+1} with $k \geq 4$ and $n \geq 2$ or $k = 3$ and $n \geq 3$ with $4(n+1) - 5 = 4n - 1$ faults such that every node is incident to two healthy links. With $4(n+1)-5$ faults, we cannot guarantee that one of the $n+1$ dimensions has four faults. However, with $n \geq 2$, we have $4(n+1)-5 > 2(n+1)$. Thus, by the pigeon hole principle, some dimension must have at least three faults. Without loss of generality (WLOG) we assume dimension 1 has at least three faults. We now partition Q^k_{n+1} across dimension 1. This leaves us with k copies of Q^k_n which we label $Q_1, Q_2, \cdots Q_k$. We now have the required graph to apply the induction hypothesis. With at least 3 faults committed to dimension 1, we have at most $4n - 4$ faults distributed among the Q_i's.

The induction argument follows three principal cases:

- case (i): Within each Q_i every node is incident to two healthy links in Q_i,
- case (ii): Some node x_i is incident to only one healthy link in Q_i,
- case (iii): Some node x_i is incident to zero healthy links in Q_i.

Note that these cases are distinct. The only overlap to consider is between cases (ii) and (iii). Suppose x_i and y_j are nodes in Q^k_{n+1} such that x_i is incident to only one healthy link in Q_i and y_j is incident to no healthy links in Q_j. Then x_i must be incident to $2n-1$ faults in Q_i and y_j must be incident to $2n$ faults in Q_j. It may be the case $i = j$, and one of the faults incident to x_i is also incident to y_j. Thus we have at least $(2n - 1) + (2n) - 1 = 4n - 2$ faults distributed among the Q_i's. Since we can have at most $4n - 4$ faults distributed among the Q_i's, cases (ii) and (iii) cannot occur simultaneously.

We further break cases (i) and (ii) into sub-cases.

- **case i:** Within each Q_i every node is incident with two healthy links in Q_i
    - **case i-a:** No Q_i contains all $4n-4$ faults distributed among the Q_i's,
    - **case i-b:** Some Q_i contains $4n-4$ faults.
- **case ii:** Some node x_i is incident to only one healthy link in Q_i,
    - **case ii-a:** There is an edge $\{w_i, x_i\}$ in Q_i that is faulty with the edge $\{w_i, w_{i+1}\}$ healthy.
    - **case ii-b:** There is no such edge $\{w_i, x_i\}$.
- **case iii:** Some node x_i is incident to zero healthy links in Q_i,

3.2 Case i: Within each Q_i every node is incident to two healthy links in Q_i

We have the hypotheses of the main theorem that Q^k_{n+1} has $4n-1$ faults such that every node is incident to at least two healthy links. We also have the assumption that dimension 1 contains at least three faults.

3.2.1 Case i-a: No Q_i has $4n-4$ faults

- **The graph Q_1 contains a Hamiltonian cycle C_1**

We can relabel the Q_i's such that Q_1 contains at least as many faults as any other Q_i. Thus each of $Q_2, \cdots, Q_k$ contains no more than $2n-2$ faults which is half the maximum number of possible faults distributed among the Q_i's. By our assumptions, Q_1 contains at most $4n-5$ faults with every node in Q_1 incident to at least two healthy links in Q_1. Thus, we can apply our induction hypothesis to Q_1 to get a Hamiltonian cycle C_1.

- **There exists a healthy 3-bridge between C_1 and Q_2**

We start by using a counting argument to establish a healthy 3-bridge between Q_1 and Q_k or a healthy 3-bridge between Q_1 and Q_2 (see Figure 3.1).

We partition C_1 into groups of three vertices to form desired configurations that only overlap in healthy links in C_1 (See Figure 3.2). Recall $Q_i = Q^k_n$. So, there are k^n nodes in C_1 which gives us $2\lfloor \frac{k^n}{3} \rfloor$ candidate 3-bridges. However the faults may be distributed among Q_2, Q_k, or dimension 1; so, we have at most $4n-1$ candidate 3-bridges eliminated due to containing a faulty edge.

The following edge case calculations demonstrate that for $k \geq 4$ and $n \geq 2$ or $k = 3$ and $n \geq 3$ we have $2\lfloor \frac{k^n}{3} \rfloor > 4n-1$.

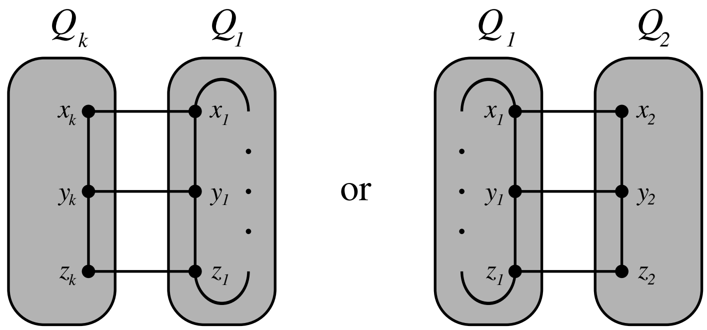

Figure 3.1: Putative healthy 3-bridge

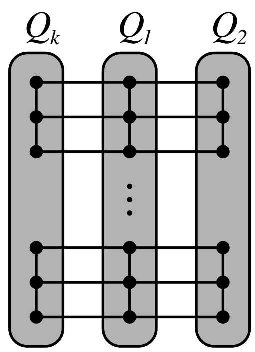

Figure 3.2: Partitioning C_1

$k = 4, n = 2:$

$$2\left\lfloor \frac{k^n}{3} \right\rfloor = 2\left\lfloor \frac{4^2}{3} \right\rfloor = 2 \cdot 5 = 10,$$

$$4n - 1 = 4 \cdot 2 - 1 = 7,$$

$k = 3, n = 3:$

$$2\left\lfloor \frac{k^n}{3} \right\rfloor = 2\left\lfloor \frac{3^3}{3} \right\rfloor = 2 \cdot 9 = 18,$$

$$4n - 1 = 4 \cdot 3 - 1 = 11.$$

So, there remains at least one desired healthy 3-bridge. By relabeling the Q_i's if necessary, assume we have the right hand image in Figure 3.1.

- **There is a cycle C_2 in Q_2 that contains the established 3-bridge**

In order to connect C_1 with Q_2, we need a Hamiltonian cycle in Q_2 that passes through $\{x_2, y_2\}$ or $\{y_2, z_2\}$. With our current assumptions, we may apply the induction hypothesis to get a Hamiltonian cycle in Q_2. If there are only two healthy links in Q_2 incident to y_2, then we are done since the given Hamiltonian cycle must pass through both $\{x_2, y_2\}$ and $\{y_2, z_2\}$. If y_2 is incident to more than two healthy links in Q_2, mark all but three links incident to y_2 in Q_2 as faulty making sure to leave $\{x_2, y_2\}$ and $\{y_2, z_2\}$ remaining as healthy links. Recall that Q_2 initially had at most $2n-2$ faults and we have introduced at most $2n-3$ faults for a current maximum of $4n-5$ faults in our modified Q_2. However, in order to apply the induction hypothesis, we must additionally ensure every node in the modified Q_2 is incident to two healthy links in Q_2.

Suppose some node in the modified Q_2 is not incident to two healthy links in Q_2. Every node other than y_2 is affected by at most by one of the temporary faults. Since each node was initially incident to at least two healthy links, and we explicitly leave three healthy links incident to y_2, every node in the modified Q_2 must be incident to at least one healthy link. Suppose after adding the temporary faults, w_2 is now incident to only one healthy link. Then w_2 must have initially been incident to $2n-2$ faults which is the maximum possible faults in the unmodified Q_2. So, w_2 is the only node in Q_2 incident to less than two healthy links. We now mark the third healthy link incident to y_2 as faulty and reinstate the link $\{y_2, w_2\}$. Every fault in Q_2 is incident to either y_2 or w_2. So, every other node in Q_2 loses at most two healthy links due to faulty connections with y_2 and w_2. By our construction, y_2 and w_2 are both incident to two healthy links.

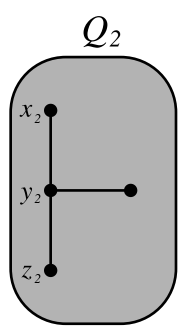

Figure 3.3: Healthy Links Incident To y_2

Since every node in our modified Q_2 is incident to at least two healthy links, we can apply our induction hypothesis. So, we have a Hamiltonian cycle C_2 in Q_2. Due to the faults we temporarily placed in incidence to y_2, C_2 must contain either the link $\{x_2, y_2\}$ or $\{y_2, z_2\}$ to incorporate the node y_2 (see Figure 3.3).

- **Form the cycle D_2 by joining C_1 and C_2**

By relabeling the nodes of Q_2 if necessary, we can assume WLOG that C_2 contains the link $\{x_2, y_2\}$. We can join the cycles C_1 and C_2 to form the cycle

D_2 as depicted in Figure 2.3.

- **There exists a healthy 3-bridge between D_2 and either Q_k or Q_3**

Note that $D_2 \cap Q_1$ and $D_2 \cap Q_2$ have k^n nodes each. As depicted in Figure 3.4, we can partition D_2 to get $2\lfloor \frac{k^n}{3} \rfloor$ 3-bridges that only overlap in healthy links in D_2.

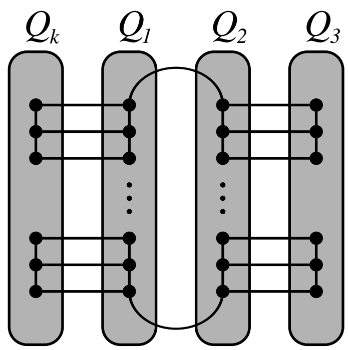

Figure 3.4: Partitioning D_2

We apply the same counting argument used to establish a 3-bridge between C_1 and Q_2 to conclude that we have at least one of our desired 3-bridges. Relabeling Q^k_{n+1} if necessary, we assume WLOG that we have a healthy 3-bridge connecting D_2 with Q_3.

- **Repeat the Last Two Steps**

The conditions that applied to Q_2 in the argument above, namely that Q_2 contains no more than $2n - 2$ faults, also apply to Q_3. Thus we have a Hamiltonian cycle in Q_3 that passes through one of the required edges of the 3-bridge connecting D_2 and Q_3 to create a cycle D_3 that passes through all of the nodes in Q_1, Q_2, and Q_3. We can apply this same argument repeatedly to create the cycles $D_4, \cdots D_{k-2}$.

- **Completing the Hamiltonian Cycle of Q^k_{n+1}**

Creating the final cycle D_{k-1} requires a slightly modified argument as the 3-bridges that are disjoint in Figure 3.4 may now share links in Q_k. Thus a fault in Q_k may now render two 3-bridges unavailable. Recall that there are at most $2n-2$ faults in Q_k. In the worst case, we have at least $2\lfloor \frac{k^n}{3} \rfloor - 2(2n-2)$ 3-bridges that survive the faults in Q_k, and we have at most $2n+1$ additional faults to consider. Recall that we have $4n-1$ faults in our Q^k_{n+1}. The following edge case calculations show for $k \geq 4$ and $n \geq 2$ or $k = 3$ and $n \geq 3$, $2\lfloor \frac{k^n}{3} \rfloor - 2(2n-2) > 2n + 1$

$k = 4, n = 2:$

$$2\left\lfloor\frac{k^n}{3}\right\rfloor - 2(2\cdot n-2) = 2\left(\left\lfloor\frac{4^2}{3}\right\rfloor - (2\cdot 2-2)\right) = 2\cdot(5-2) = 6$$

$$2n+1 = 2\cdot 2+1 = 5.$$

$k = 3, n = 3:$

$$2\left\lfloor\frac{k^n}{3}\right\rfloor - 2(2\cdot n-2) = 2\left(\left\lfloor\frac{3^3}{3}\right\rfloor - (2\cdot 3-2)\right) = 2\cdot(9-4) = 10$$

$$2n+1 = 2\cdot 3+1 = 7.$$

So, we have a desired healthy 3-bridge. The conditions that applied to Q_2 in constructing C_2 again apply to Q_k. So, we have a Hamiltonian cycle C_k in Q_k that passes through one of the required links in the 3-bridge connecting D_{k-2} with Q_k. Thus we can join C_k to D_{k-2} to get a Hamiltonian cycle of Q^k_{n+1}. This completes the subcase (i-a) in which no Q_i contains $4n-4$ faults.

3.2.2 Case i-b: Some Q_i contains $4n-4$ faults

In this case we assume every node in each Q_i is incident to two healthy links in Q_i. We also assume that the maximum number of faults not residing in dimension 1 all lie within a single Q_i. By relabeling the indices if necessary, we assume WLOG that Q_1 contains the maximum $4n-4$ faults. Let x_1 and y_1 be nodes in Q_1 connected by a faulty link. Suppose further that $\{x_1, y_1\}$ is a link in the 2-bridge $\{x_1, y_1, x_2, y_2\}$ where the links $\{x_1, x_2\}$, $\{x_2, y_2\}$, and $\{y_1, y_2\}$ are all healthy as depicted in Figure 3.5. We exhibit a Hamiltonian cycle in Q^k_{n+1} given this 2-bridge. We then show via contradiction that such a 2-bridge always exists.

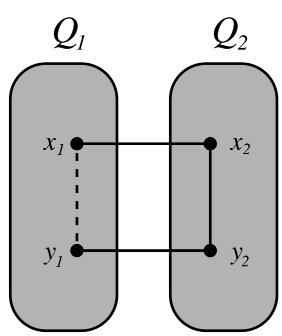

Figure 3.5: The Given 2-Bridge

- **Each Q_i contains an isomorphic cycle C_i**

Given the 2-bridge in Figure 3.5, temporarily mark the faulty link $\{x_1, y_1\}$ as healthy. We can then apply the induction hypothesis to Q_1 as Q_1 now only has $4n-5$ faults and each node in Q_1 is incident to two healthy links in Q_1. Thus, we have a Hamiltonian cycle C_1 in Q_1. Note that C_1 may or may not contain the link $\{x_1, y_1\}$. Either way, since all of the faults not in dimension 1 are in Q_1, each Q_i, $i = 2, \ldots, k$ contains a healthy cycle C_i that is an isomorphic copy of C_1.

- **There is a cycle D_1 containing all nodes in Q_1 and Q_2**

If $\{x_1, y_1\} \in C_1$, then we have a Hamiltonian Path $P_1 : x_1 \sim y_1\ Q_1$. We can join P_1 with C_2 via the 2-bridge $\{x_1, y_1, x_2, y_2\}$ to get the cycle D_1 that contains all of the nodes in Q_1 and Q_2.

If $\{x_1, y_1\} \notin C_1$, then we apply a counting argument similar to the one used in case (i-a) to argue for the existence of a healthy 2-bridge $\{u_1, v_1, u_2, v_2\}$ connecting C_1 and C_2. Since there are k^n nodes in Q_1, we have $\lfloor \frac{k^n}{2} \rfloor$ candidate 2-bridges. The only faults left to consider are the three faults in dimension 1. The following edge case calculations show that we have at least one such healthy 2-bridge. We use the 2-bridge to join C_1 and C_2 to form the cycle D_1 containing all of the nodes in Q_1 and Q_2

$$k = 4, n = 2: \qquad \left\lfloor \frac{k^n}{2} \right\rfloor = \left\lfloor \frac{4^2}{2} \right\rfloor = 8 > 3$$
$$k = 3, n = 2: \qquad \left\lfloor \frac{k^n}{2} \right\rfloor = \left\lfloor \frac{3^3}{2} \right\rfloor = 4 > 3$$

So, whether or not $\{x_1, y_1\} \in C_1$, we have a cycle D_1 containing all of the nodes in Q_1 and Q_2.

- **There is a healthy 2-bridge joining D_1 and C_3**

Note that $D_1 \cap Q_2$ contains k^n nodes. This gives us $\lfloor \frac{k^n}{2} \rfloor$ disjoint candidate 2-bridges connecting D_1 and C_3. The counting argument above that established the 2-bridge $\{u_1, v_1, u_2, v_2\}$ applies to guarantee a healthy 2-bridge between D_1 and C_3. We can then join D_1 and C_3 to form the cycle D_2 that contains all of the nodes in Q_1, Q_2, and Q_3.

- **Repeat the previous step**

We can repeatedly apply the argument of the previous step to get the cycles $D_3, \cdots, D_{k-1}$. By the nature of this construction, we avoid the extra considerations of case (i-a) in constructing D_{k-1}. So, given the 2-bridge in Figure 3.5, we have a Hamiltonian cycle in Q^k_{n+1}.

- **The 2-bridge in Figure 3.5 always exists**

Suppose no such 2-bridge exists. Then it must be the case that no such 2-bridge exists between Q_1 and Q_k as well as between Q_1 and Q_2 or else we could simply relabel the graph. Thus, each fault in Q_1 must be incident with at least two dimension 1 faults. That is, given a fault $\{x_1, y_1\}$ we must have one of the four scenarios depicted in Figure 3.6.

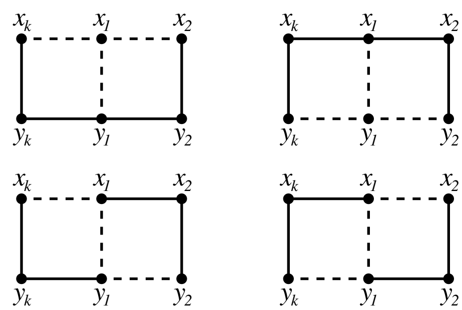

Figure 3.6: The Four Ways to Prevent the 2-bridge

So, if two faults in Q_1 are not incident to one another, at least four dimension 1 faults are required to prevent the existence of the 2-bridge in Figure 3.5. Thus, with only three dimension 1 faults, it must be the case that any two faults in Q_1 must incident to one another. This constraint on the faults in Q_1 permits only three possibilities: the faults form a 3-cycle, a star with x_1 at the center, or a star with y_1 at the center. Thus, in the worst case, we can have at most $2n-2$ faults (since x_1 and y_1 are both incident to at least 2 healthy links in Q_1) and still prevent the existence of the 2-bridge in Figure 3.5. Since $4n-4 > 2n-2$ for $n \geq 2$, we have enough faults in Q_1 to guarantee the existence of the 2-bridge in Figure 3.5.

This completes the proof of the induction step for sub-case (i-b) and case (i).

3.3 Case ii: Some node x_i is incident to only one healthy link in Q_i

In this case, we assume some Q_i has a node incident to only one healthy link within Q_i. Note that we still have at least three faults in dimension 1, and every node in Q_{n+1}^k is incident to two healthy links. By relabeling the Q_i's if necessary, we assume WLOG that $\{x_1, y_1\}$ is the only healthy link in Q_1 incident to x_1. We additionally label Q_{n+1}^k such that $\{x_1, x_2\}$ is a healthy link. It may or may not be the case that the link $\{x_1, x_k\}$ is healthy.

With only $4n-4$ faults distributed among the Q_i's, every node in Q_i $i \neq 1$ must be incident to at least three healthy links in Q_i. Otherwise, $(2n-1)+(2n-2) = 4n-3$ faults are required. It is possible that some node $y_1 \in Q_1$, $y_1 \neq x_1$ may be incident to only two healthy links in Q_1.

Case (ii) breaks into two subcases depending upon whether or not there exists some w_1 such that $\{x_1, w_1\}$ is faulty and $\{w_1, w_2\}$ is healthy.

3.3.1 Case ii-a: There exists w_1 such that $\{x_1, w_1\}$ is faulty and $\{w_1, w_2\}$ is healthy

- **There is a Hamiltonian Path $P_1 : x_1 \sim y_1$ in Q_1**

Every node in Q_1 other than x_1 is incident to at least two healthy links in Q_1. Temporarily mark the faulty link $\{x_1, w_1\}$ as healthy. Then every node in Q_1 is now incident to at least two healthy links in Q_1. Furthermore, Q_1 had at most $4n - 4$ faults initially. So, our modified Q_1 has at most $4n - 5$ faults. Thus, we may apply the induction hypothesis to our modified Q_1 to get a Hamiltonian cycle that must necessarily pass through the link $\{x_1, w_1\}$. Thus, we have a Hamiltonian Path $P_1 : x_1 \sim w_1$ in the unmodified Q_1.

- **There is a Hamiltonian Path $P_2 : x_2 \sim y_2$ in Q_2**

If necessary, temporarily mark the link $\{x_2, w_2\}$ as healthy. As argued above, w_2 is incident to at least 3 healthy links in Q_2. Select the link $\{x_2, w_2\}$ and one other healthy link incident to w_2 in Q_2 to leave as healthy. Temporarily mark the remaining links incident to w_2 in Q_2 as faulty.

We now argue that we can apply the induction hypothesis to Q_2. First note that Q_2 initially had $2n-3$ faults and we have introduced at most $2n-2$ faults for a maximum total of $4n-5$ faults. By the argument at the beginning of this section, each node in the unmodified Q_2 was incident to at least three healthy links in Q_2. Since each introduced fault is incident to w_2, every node in Q_2 other than w_2 is affected by at most one of the introduced faults. Thus, every node other than w_2 must be incident to at least two healthy links in Q_2. Since we explicitly left two healthy links incident to w_2, we may apply the induction hypothesis to our modified Q_2 to get a Hamiltonian cycle C_2 in Q_2. Whether or not the link $\{x_2, w_2\}$ is healthy in the original Q_2, we have a Hamiltonian Path $P_2 : x_2 \sim w_2$ in the unmodified Q_2.

- **There is a cycle D_1 containing all of the nodes in Q_1 and Q_2**

We explicitly chose x_1 and w_1 such that the links $\{x_1, x_2\}$ and $\{w_1, w_2\}$ are healthy. Thus we can join P_1 and P_2 to get the cycle D_1 that contains all of the nodes in Q_1 and Q_2.

- **There is a healthy 2-bridge joining D_1 with Q_3**

There are $\lfloor \frac{k^n}{2} \rfloor$ disjoint 2-bridges connecting D_1 and Q_3 (See Figure 3.7). With $2n - 1$ faults restricted to Q_1, we need only consider the remaining $2n$ faults. The following edge case calculations demonstrate that for $k \geq 2$ and $n \geq 3$ or $k \geq 4$ and $n = 2$, the inequality $\lfloor \frac{k^n}{2} \rfloor > 2n$ holds.

$$k = 4, n = 2: \qquad \left\lfloor \frac{k^n}{2} \right\rfloor = \left\lfloor \frac{4^2}{2} \right\rfloor = 8 \ > 2n$$

$$k = 3, n = 3: \qquad \left\lfloor \frac{k^n}{2} \right\rfloor = \left\lfloor \frac{3^3}{2} \right\rfloor = 13 > 2n$$

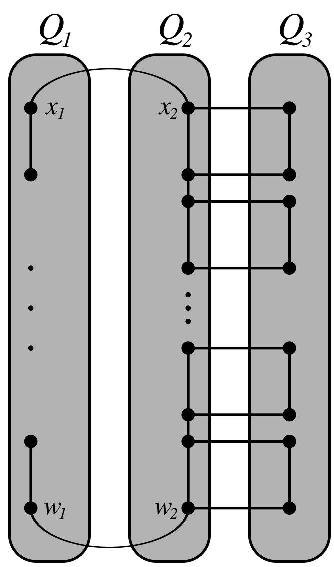

Figure 3.7: Partitioning D_1

Thus a 2-bridge $\{u_2, v_2, u_3, v_3\}$ between D_1 and Q_3 survives the $4n-1$ faults in Q^k_{n+1}.

- **There is a cycle D_2 containing all of the nodes in Q_1, Q_2, and Q_3**

The argument that established P_2 applies to establish a Hamiltonian path $P_3 : u_3 \sim v_3$ in Q_3. We then join D_1 and P_3 to get a cycle D_2 that contains all of the nodes in Q_1, Q_2, and Q_3.

- **Repeat the last two steps**

We can repeat the last two steps to produce the cycles $D_3, D_4, \cdots, D_{k-1}$ which gives us a Hamiltonian cycle of Q^k_{n+1}. This completes the induction step for the case (ii-a).

3.3.2 Case ii-b: There is no node w_1 such that $\{x_1, w_1\}$ is faulty and $\{w_1, w_2\}$ is healthy

In this case, each of the $2n-1$ neighbors of x_1 joined to x_1 via a faulty link in Q_1 is also incident to a dimension 1 fault between Q_1 and Q_2. This accounts for $4n-2$ faults which is all but one of the $4n-1$ faults in Q^k_{n+1}. Now suppose the link $\{x_1, x_k\}$ is healthy. With $n \geq 2$, x_1 is incident to at least three faulty links in Q_1. With only one fault in Q^k_{n+1} not accounted for, there is a node w_1 such that $\{x_1, w_1\}$ is faulty and $\{w_1, w_k\}$ is healthy. So, by symmetry, we are in case (ii-a), and we construct the cycles $D_1, D_2, \cdots, D_{k-1}$ by incorporating Q_i's in the opposite order. So, for the remainder of this section suppose $\{x_1, x_k\}$ is faulty which gives us a precise accounting of all of the faults in Q^k_{n+1}.

Ashir and Stewart offer two different such Hamiltonian cycles of Q^k_{n+1} depending upon the parity of k. We start by establishing a bridge between Q_k and Q_1 and a bridge between Q_1 and Q_2.

- **There is a Hamiltonian path P_1 in Q_1**

Let w_1 be a node in Q_1 such that $\{x_1, w_1\}$ is faulty. Temporarily mark the link $\{x_1, w_1\}$ as healthy. Since Q_1 contained only $2n-1$ faults, the modified Q_1 contains $2n < 4n-5$ faults. Recall from the discussion at the beginning of case (ii) that every node α_i other than x_1 is incident to at least two healthy links within Q_i. Temporarily marking $\{x_1, w_1\}$ as healthy places two healthy links in Q_1 incident with x_1 which allows us to apply the induction hypothesis to our modified Q_1. So, we have a Hamiltonian cycle C_1 in Q_1. By the nature of the faults incident to x_1, C_1 must contain the link $\{x_1, w_1\}$. Thus we have a Hamiltonian path P_1 in the original unmodified Q_1 with endpoints x_1 and w_1.

- **Q_2 and Q_k each contain a healthy path isomorphic whose vertices differ from those in P_1 only in the first coordinate**

In the current case, all of the faults in Q^k_{n+1} lie in Q_1 or in dimension 1 links with an endpoint in Q_1. Thus, Q_2 and Q_k each contain a healthy path isomorphic to P_1. We denote the two paths P_2 and P_k respectively.

- **A collection of parallel paths in dimension 1**

Consider a node α_2 in Q_2. As noted above, the faults in Q^k_{n+1} are restricted to Q_1 or dim-1 links with endpoints in Q_1. Thus the dim-1 links connecting $\alpha_2, \alpha_3, \ldots, \alpha_k$ are all healthy. Therefore, we have a collection of parallel paths $S_1, S_2, \ldots, S_{k^n}$ that collectively contain all of the nodes in $Q_2, Q_3, \ldots, Q_k$.

- **Knitting together $S_1, S_2, \ldots, S_{k^n}$ and P_1**

We will use links in P_2 and P_k to connect the endpoints of the paths $S_1, S_2, \cdots, S_{k^n}$. We then complete the Hamiltonian cycle of Q^k_{n+1} using the links $\{w_1, w_k\}$, $\{x_1, x_2\}$, and the path P_1.

The parity of k^n affects the manner in which we complete the knitting of the Hamiltonian cycle of Q^k_{n+1}. Since the parity of k^n is exactly that of k, Ashir and Stewart offer two alternative stitchings of the final Hamiltonian cycle depending on the parity of k. To facilitate describing the stitching process, we use a double indexing notation for the graph Q^k_{n+1}. Let $x_{1,1}, x_{1,2}, \cdots x_{1,k^n}$ denote the nodes in P_1 where $x_{1,1} = x_1, x_{1,2} = y_1, \cdots x_{1,k^n} = w_1$. As indicated in Figure 3.8, we similarly label the corresponding nodes in Q_i by $x_{i,1}, x_{i,2}, \cdots x_{i,k^n}$.

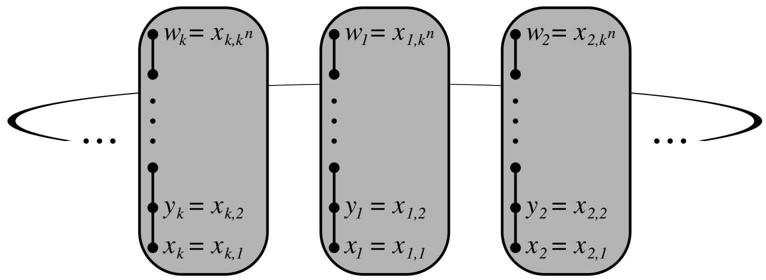

Figure 3.8: Double Index Labeling of Q^k_{n+1}

- **The Hamiltonian cycle when k is even**

To slightly compact the notation of the Hamiltonian cycle, let $S'_i = S_i - \{x_{k,i}\}$ and $P' = P - \{x_1, 2\}$. The following explicit Hamiltonian cycle of Q^k_{n+1} is depicted in Figure 3.9.

$$x_{1,k^n} \cdot S_{k^n} \cdot x_{2,k^n-1} \cdot S_{k^n-1} \cdot x_{k,k^n-2}, x_{k,k^n-2} \cdot S_{k^n-2} \cdots S'_3 \cdot x_{k-1,2} \cdot S'_2 \cdot x_{1,1} \cdot S_1 \cdot x_{k,1} \cdot x_{k,2} \cdot x_{k,3} \cdot x_{1,3} \cdot P'$$

This completes sub-case (ii-b) and case (ii).

- **The Hamiltonian cycle when k is odd**

The following explicit Hamiltonian cycle is depicted in Figure 3.10

$$x_{1,k^n} \cdot S_{k^n} \cdot x_{2,k^n-1} \cdot S_{k^n-1} \cdots x_{1,k^n} \cdot S_2 x_{k,1} \cdot S_1 \cdot P_1$$

3.4 Case iii: Some node x_i is incident to zero healthy links in Q_i

We still have our initial assumptions that our given Q^k_{n+1} contains $4n-1$ faults with every node incident to at least two healthy neighbors. We still have at least 3 faults confined to dimension 1. For case (iii) we further assume that some node $x_i \in Q_i$ is incident to zero healthy links in Q_i. To see that every

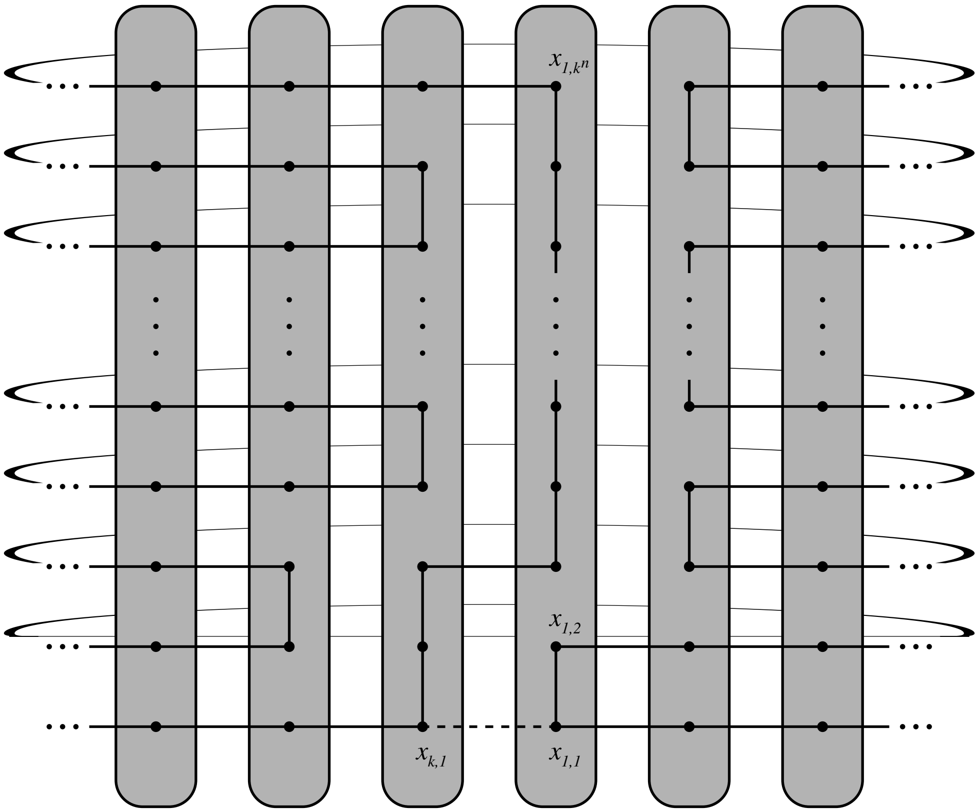

Figure 3.9: (Case ii-b) The Hamiltonian cycle when k is even.

node α_i other than x_1 must be incident to at least three healthy links in Q_i consider the remaining $2n-4$ faults that may be in the Q_i's $i \neq 1$. In the worst case all of the remaining $2n-4$ faults may be incident to a node w_1 already incident to x_1 via a faulty link. Thus, in this worst case, w_1 is still incident to three healthy links in Q_1.

Relabeling the Q^k_{n+1} if necessary we assume WLOG that $x_1 \in Q_1$ is incident to zero healthy links in Q_1. Thus the links $\{x_1, x_2\}$ and $\{x_1, x_k\}$ must be healthy.

- **The Subgraph (Figure 3.12) Exists**

We begin by establishing the existence of one of the configurations in Figure 3.11. The following counting argument shows that a desired configuration survives all of the faults in Q^k_{n+1}.

Note that x_1 is the center of a star consisting of $2n$ faults in Q_1. Partition those faults into n pairs. In order to not have one of the configurations in Figure 3.11, two dim-1 faults are required to prevent both possible configurations. Thus $4n$ faults in total are required. However, we have only $4n-1$ faults in our Q^k_{n+1}.

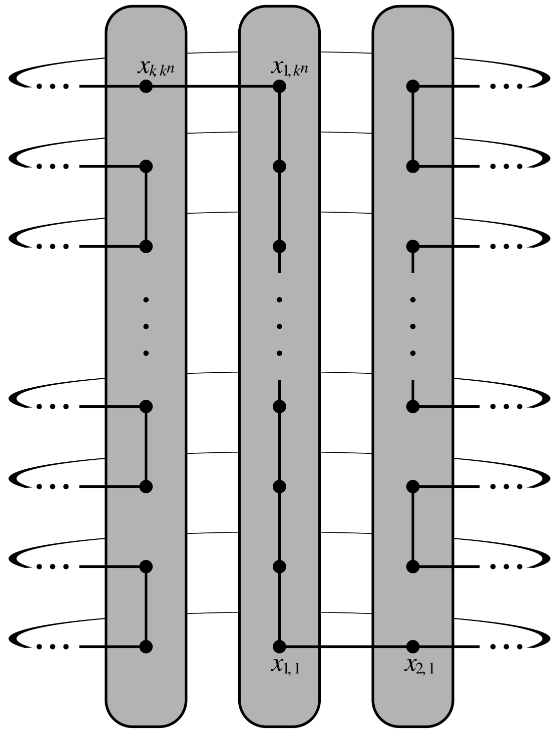

Figure 3.10: (Case ii-b) The Hamiltonian cycle when k is odd.

So, at least one such configuration exists. By relabeling Q^k_{n+1} if necessary,we assume WLOG that we have the configuration in Figure 3.12.

- The path $P_1 : y_1 \sim z_1$ in Q_1 with $|P_1| = k^n - 1$ is healthy.

Temporarily mark the links $\{x_1, y_1\}$ and $\{x_1, z_1\}$ as healthy. Originally, Q_1 contained at most $4n - 4$ faults. So, with the two newly instated healthy links, the modified Q_1 contains at most $4n - 6$ faults which is less than the $4n - 5$ faults of the induction hypothesis. By the argument at the beginning of case (iii), every node in the original Q_1 other than x_1 is incident to two healthy links in Q_1. The newly marked healthy links place two healthy links incident with x_1. Thus every node in the modified Q_1 is incident to two healthy links. Thus, we can apply the induction hypothesis to our modified Q_1. Thus, we have a Hamiltonian cycle C_1 in Q_1. Note that C_1 must pass through the links $\{x_1, y_1\}$ and $\{x_1, z_1\}$. Thus we have a path $P_1 : y_1 \sim z_1$ that contains all of the nodes of Q_1 except x_1.

- The Hamiltonian path $P_k : x_k \sim z_k$ in Q_k is healthy.

Recall that Q_k contains no more than $2n-4$ faults. If necessary, temporarily mark the link $\{x_k, z_k\}$ as healthy. Recall that x_k was initially incident to at

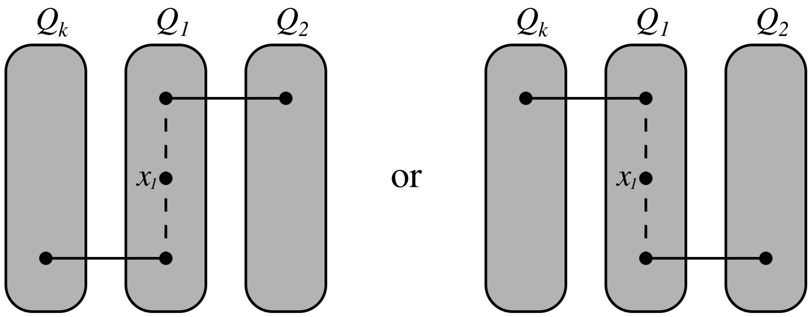

Figure 3.11: Case iii: Needed Configuration

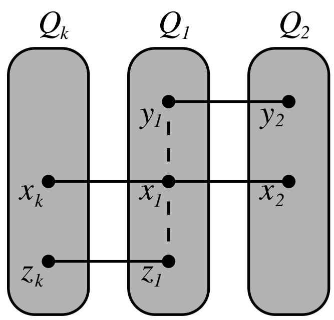

Figure 3.12: Case iii: Configuration In Hand

least 3 healthy links in Q_k. Choose the link $\{x_k, z_k\}$ and a healthy link incident to x_k to preserve. Mark the remaining $2n-2$ links incident to x_k in Q_k as faulty. Note that our modified Q_k has at most $(2n-4)+(2n-2) = 4n-6 < 4n-5$ faults. We now verify the second condition of the induction hypothesis holds. Each fault we introduced was incident to x_k. So, each node in Q_k other than x_k is affected by at most 1 of the introduced faults. Since every node in the original Q_k was incident to at least 3 healthy links, every node other than x_k in Q_k is healthy to at least 2 healthy links. We explicitly left 2 healthy links incident to x_k. So, we may apply the induction hypothesis to Q_k to get a Hamiltonian cycle that must necessarily contain the link $\{x_k, z_k\}$. Therefore, we have the Hamiltonian path $P_k : x_k \sim z_k$ in the unmodified Q_k.

- There exists a healthy Hamiltonian path $P_2 : x_2 \sim y_2$ in Q_2.

All of the conditions that existed in the previous case exist for Q_1. Thus we can apply the previous argument mutatis mutandis to get a Hamiltonian path $P_2 : x_2 \sim y_2$ in Q_2.

- There is a cycle D_2 containing all of the nodes in Q_k, Q_1, and Q_2.

$$D_2 = P_k \cdot x_1 \cdot x_2 \cdot P_2 \cdot y_1 \cdot P_1 \cdot z_k$$

- There is a healthy 2-bridge connecting D_2 and Q_3.

With $|D_2 \cap Q_2| = k^n$, we have $\lfloor \frac{k^n}{2} \rfloor$ disjoint 2-bridges between D_2 and Q_3. We have at most $2n-1$ faults outside of Q_1 to consider. The following counting argument shows that we do not have enough faults to prevent the existence of a required 2-bridge. Denote the 2-bridge by $\{u_3, v_3, u_2, v_2\}$.

$k = 4, n = 2:$

$$2\left\lfloor \frac{k^n}{2} \right\rfloor = 2\left(\left\lfloor \frac{4^2}{2} \right\rfloor\right) = 16$$
$$2n - 1 = 2 \cdot 2 - 1 = 3.$$

$k = 3, n = 3:$

$$2\left\lfloor \frac{k^n}{2} \right\rfloor = 2\left(\left\lfloor \frac{3^3}{2} \right\rfloor\right) = 2 \cdot (4) = 8$$
$$2n - 1 = 2 \cdot 3 - 1 = 5.$$

- There is a Hamiltonian path $P_3 : u_3 \sim v_3$ in Q_3.

The argument establishing P_k and P_1 applied directly to give us the Hamiltonian path P_3 in Q_3.

- There is a cycle D_3 containing all of the nodes in Q_k, Q_1, Q_2, and Q_3

Connect D_2 and P_3 with our usual construction to get the cycle D_3.

- Repeat to construct the Hamiltonian cycle D_{k-1} in Q^k_{n+1}

Repeat the last four steps to successively construct the cycles $D_4, \ldots, D_{k-1}$. This concludes the proof of the induction step.

Chapter 4

The Base Case Q_2^k, $k \geq 4$

4.1 Argument Overview

We assume we have 4(2) - 5 = 3 faulty links. With 3 faults distributed among 2 dimensions, some dimension must contain at least 2 faults. Assume WLOG that dimension 1 contains at least 2 faults. We now partition Q_2^k across dimension 1 to get $Q_1, Q_2, \ldots, Q_k$. Note that each Q_i is a k-cycle. We now consider if the third fault is in dimension 1 or in one of the Q_i's.

4.2 Case i: All 3 faults lie in dimension 1

With all of the faults lying in dimension 1, each Q_i is a healthy k-cycle which is necessarily Hamiltonian. For consistency with the rest of this paper, we let $C_1, \ldots, C_k$ stand for the Hamiltonian cycles in $Q_1, \ldots, Q_k$ respectively. We now need a 2-bridge between C_1 and C_2 or between C_1 and C_k. Since each C_i contains k nodes, we have $2\lfloor \frac{k}{2} \rfloor$ disjoint 2-bridges to connect C_1 with either C_2 or C_k. With $k \geq 4$, we have $2\lfloor \frac{k}{2} \rfloor > 3$. So, a desired 2-bridge exists. Assume WLOG that we have a 2-bridge between C_1 and C_2 which gives a cycle D_1 encompassing all of the nodes in Q_1 and Q_2. We now need a 2-bridge connecting D_1 with either C_k or C_3. Again we have $2\lfloor \frac{k}{2} \rfloor$ disjoint 2-bridges with $k \geq 4$ and only 3 faults in dimension 1. So, some 2-bridge is left unaffected by the faults in Q_2^k. Assume WLOG that we have a 2-bridge between D_1 and C_3 which yields a cycle D_2 encompassing all of the nodes in Q_1, Q_2, and Q_3. This construction continues to get the cycle D_{k-2} which encompasses all of the nodes in $Q_1, Q_2, \ldots, Q_{k-1}$. Up until now, at each stage of the construction, we had $2\lfloor \frac{k}{2} \rfloor$ disjoint 2-bridges to consider. When considering 2- bridges between D_{k-1} and Q_k, we have 2-bridges between Q_1 and Q_k as well as between Q_{k-1} and Q_k. It may be the case that these 2-bridges are not disjoint within Q_k. However, this poses no problem since all of the faults are in dimension 1 by assumption. Thus, the counting argument above applies again to conclude there is a healthy

2-bridge between D_{k-1} and C_k. Thus, we can construct the cycle D_{k-1} which is our desired Hamiltonian cycle of Q_2^k.

4.3 Case ii: Only 2 faults lie in dimension 1

Assume WLOG that the third fault lies in Q_1. We label the fault in Q_1 as $\{x_1, y_1\}$. If the 2-bridge $\{x_1, x_2, y_1, y_2\}$ or the 2-bridge $\{x_1, x_k, y_1, y_k\}$ consists of healthy dim-1 links, then we may apply case (i) above by considering the Hamiltonian path $P_1 : x_1 \sim y_1$ in Q_1 as the starting point. The 2-bridge with healthy dim-1 links connects P_1 with C_2 to create D_1. The argument then proceeds identically to case (i).

Now assume that neither 2-bridge has two healthy dim-1 links. With only two dim-1 faults, it must the case that one of the dim-1 links in each of the 2-bridges $\{x_1, x_2, y_1, y_2\}$ and $\{x_1, x_k, y_1, y_k\}$ is faulty. Since each node is incident to at least two healthy links by hypothesis, it cannot be the case that the edges $\{x_1, x_2\}$ and $\{x_1, x_k\}$ are both faulty. Similarly, only one of $\{y_1, y_2\}$ and $\{y_1, y_k\}$ can be faulty. Thus, we have two possibilities for the two dim-1 faults. Either $\{x_1, x_k\}$ and $\{y_1, y_2\}$ are both faulty or $\{x_1, x_2\}$ and $\{y_1, y_k\}$ are both faulty. We can assume WLOG that the links $\{x_1, x_2\}$ and $\{y_1, y_k\}$ are faulty. In the case k is even, the Hamiltonian cycle in Figure 3.9 applies with $x_{1,3}, x_{1,2}, x_{2,3}$, and $x_{k,2}$ replaced with x_1, y_1, x_2, and y_k respectively. Figure 4.1 depicts this situation when $k = 6$. When k is odd, then figure 3.10 applies with $x_{1,m}, x_{1,1}, x_{2,m}$, and $x_{k,1}$ replaced with x_1, y_1, x_2, and y_k respectively. Figure 4.2 depicts this scenario when $k = 5$.

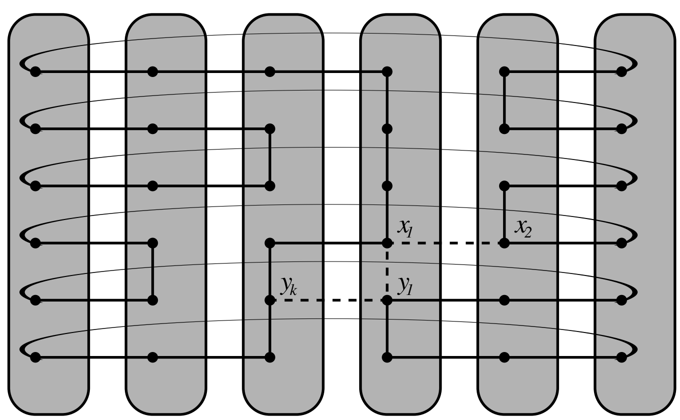

Figure 4.1: Hamiltonian Cycle for Q_2^k when $k = 6$

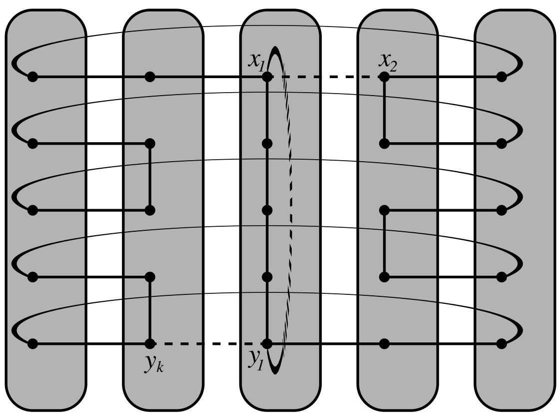

Figure 4.2: Hamiltonian Cycle for Q_2^k when $k = 5$

Chapter 5

A Lemma on Q_2^3

Lemma 1 *If Q_2^3 has three faulty links such that every node is incident to two healthy links, then Q_2^3 has a Hamiltonian cycle unless the three faults form a cycle of length 3.*

With three faults spread across two dimensions, some dimension must contain at least two faults. We assume WLOG that dimension 1 contains at least two faults. We partition Q_2^3 over dimension 1 and label the nodes of Q_2^3 as x_i, y_i, z_i, $1 \leq i \leq 3$. We have two cases to consider depending upon whether or not the third fault is also in dimension 1 or if it is in dimension 2.

5.1 Case i: Both dimensions contain faults

.

In this case the third fault must lie in one of the three Q_i's. We may assume WLOG that Q_1 contains the third fault with the link $\{x_1, y_1\}$ as faulty. With all faults accounted for, we have the healthy Hamiltonian paths $P_i : x_i \sim y_i$, $1 \leq i \leq 3$. Furthermore, the healthy links $\{x_2, y_2\}$ and $\{x_3, y_3\}$ permit us to extend P_2 and P_3 to form the cycles C_2 and C_3 in Q_2 and Q_3 respectively. The current case of having only two faults in dimension one now breaks down into two more sub-cases. We consider whether or not one of the following two 2-bridges has healthy dim-1 links:

$$\{x_1, x_2, y_1, y_2\} \quad \text{or} \quad \{x_1, x_3, y_1, y_3\}.$$

5.1.1 Case i-a: A 2-bridge contains healthy dim-1 links

Assume WLOG that the 2-bridge spans Q_1 and Q_2. Figure 5.1 shows the one known fault and a cycle D_1 joining Q_1 and Q_2 along with an isomorphic redrawing of the graph.

Figure 5.2 shows the four possible 2-bridges we can use to join D_1 and Q_3 to construct a Hamiltonian cycle of Q_2^3. If one of these four 2-bridges is healthy,

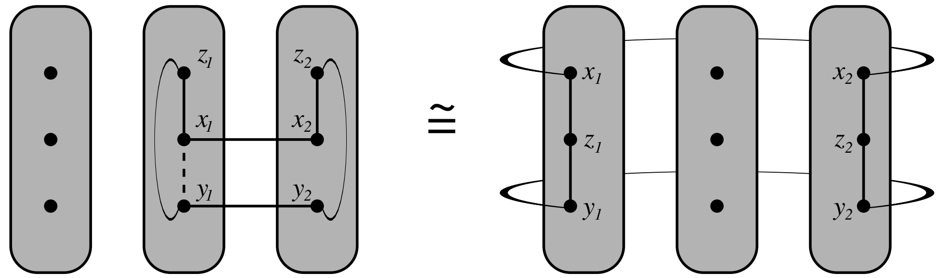

Figure 5.1: Q_2^3 lemma - case 1a

then we get our Hamiltonian cycle with the usual construction. The only way the two dimension 1 faults can destroy all four 2-bridges is if the links $\{z_2, z_3\}$ and $\{z_1, z_3\}$ are both faulty. In this case we have the Hamiltonian cycle depicted in Figure 5.3.

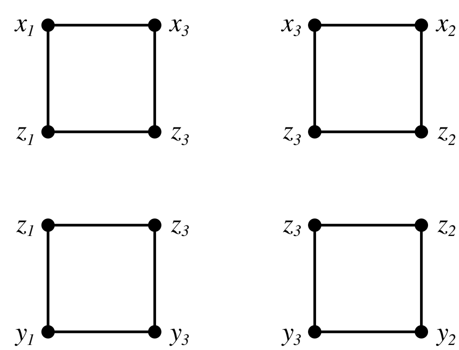

Figure 5.2: Q_2^3 lemma - case 1a: the four possible 2-bridges

5.1.2 Case i-b: Neither 2-bridge has two healthy dim-1 links

With only two dim-1 faults, it must the case that one of the dim-1 links in each of the 2-bridges $\{x_1, x_2, y_1, y_2\}$ and $\{x_1, x_k, y_1, y_k\}$ is faulty. Since each node is incident to at least two healthy links by hypothesis, it cannot be the case that the edges $\{x_1, x_2\}$ and $\{x_1, x_3\}$ are both faulty. Similarly, only one of $\{y_1, y_2\}$ and $\{y_1, y_3\}$ can be faulty. Thus, we have two possibilities for the two dim-1 faults. Either $\{x_1, x_3\}$ and $\{y_1, y_2\}$ are both faulty or $\{x_1, x_2\}$ and $\{y_1, y_3\}$ are both faulty. We can assume WLOG that the links $\{x_1, x_2\}$ and $\{y_1, y_3\}$ are faulty. In this case, we have the Hamiltonian cycle of Q_2^3 depicted in Figure 5.4

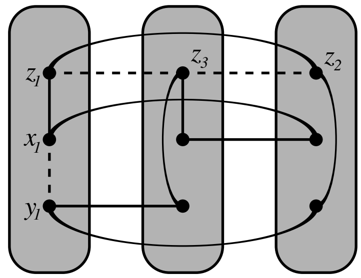

Figure 5.3: Q_2^3 lemma - case 1a : the cycle with faulty $\{z_1, z_3\}$ and $\{z_3, z_2\}$

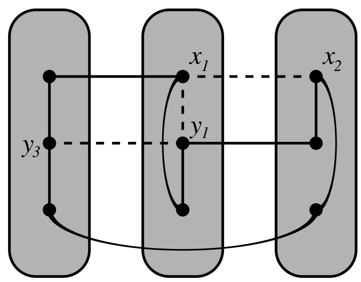

Figure 5.4: Q_2^3 lemma - case 1b

5.2 Case ii: All faults lie In one dimension

In this case, we argue that up to isomorphism there are only six ways to place the three faults in dimension one in Q_2^3. We start by encoding the dim-1 edges in a Table T as depicted in Figure 5.5. Let C_{ij} denote swapping columns i and j in T. Similarly, let R_{ij} denote swapping rows i and j in the Table T. By direct inspection of the adjacency relation in Figures 5.6 and 5.7 we see that C_{ij} and R_{ij} respectively encode isomorphisms of Q_2^3.

With these isomorphisms in hand, we consider the distinct ways we may place 3 faults within the nine entries in T. In subsequent drawings of T, the edge numbers are omitted and the faulty edges are indicated with disks.

The argument successively considers how to distribute the three faults among a single column, among two columns, and then among all three columns of T. Starting with restricting the faults to lie in a single column, we get one way to place the faults as depicted in Figure 5.8.

Now, we consider how to distribute the three dim-1 faults among two columns in T. Note that one column will contain 2 faults and the other will contain 1 fault. We start by considering how to place 2 faults in a single column. Figure 5.9 shows that by combining C_{12}, C_{23}, R_{12}, and R_{23}, there is only one way up to isomorphism to place two faults in a single column.

With the two faults placed as in the top middle table depicted in Figure

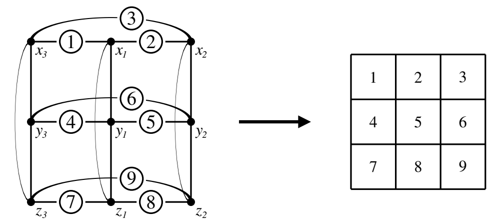

Figure 5.5: Q_2^3 lemma - case 2: encoding the dim-1 links in Table T

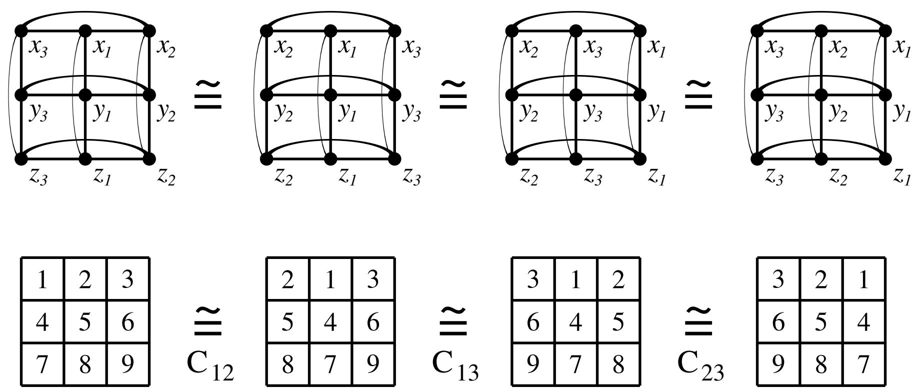

Figure 5.6: C_{ij} encodes an isomorphism of Q_2^3

5.9, we consider how to place the third fault. The two rows of isomorphisms in Figure 5.10 show that there are two ways to place the third fault. To see that that we do indeed have two isomorphism classes, note that the upper row indicates the existence of a node incident to two faults and that there is no such node in the lower row.

All we have left to consider, is how to distribute the faults among the three columns. To analyze this case, we successively consider distributing the three faults among one row, two rows, and then all three rows. Starting with placing all three faults among a single row, Figure 5.11 shows that up to isomorphism there is only way to do this.

Now we consider how to place all three faults in 3 columns and 2 rows. In this case, we must have 2 faults in one row and the third fault in a second row. Figure 5.12 shows that by combining C_{12}, C_{23}, R_{12}, and R_{23}, there is up to isomorphism only one way to place two faults in a single row.

Starting with the table in the middle left of Figure 5.12, we now consider the ways to add the third fault. Figure 5.13 shows that up to isomorphism, there are only two ways to spread the three faults across all three columns and two rows. To see that that the two rows represent distinct isomorphism classes,

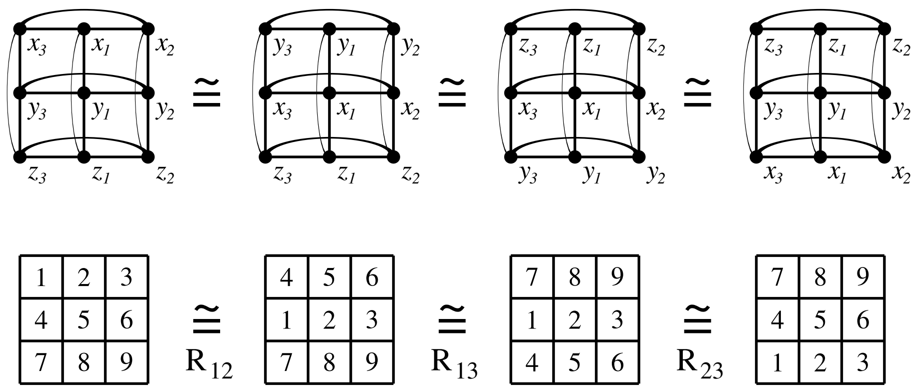

Figure 5.7: R_{ij} encodes an isomorphism of Q_2^3

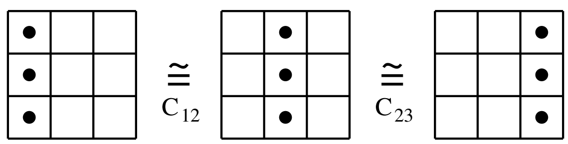

Figure 5.8: The one way to have 3 faults in a single column

consider that the upper row contains a 4-cycle containing three faults and that no such 4-cycle exists in the lower row. Furthermore, Note that the upper row of Figure 5.13 contains a configuration that is also present in the upper row of Figure 5.10.

Finally, we consider how many ways three faults can be distributed across three columns and three rows. Note that we have $3 \cdot 2 \cdot 1 = 6$ ways to place the three faults in the table. Figure 5.14 shows that all six such arrangements are isomorphic.

From Figure 5.8, we get 1 configuration. From Figure 5.10, we get 2 configurations. From Figure 5.11 we get 1 configuration. From Figure 5.13, we get only 1 additional configuration, and we get the sixth configuration from Figure 5.14. Note that the only configuration in which the three dim-1 faults form a cycle is in Figure 5.11.

Figure 5.15 displays an example from each of the six congruence classes, and Figure 5.16 shows the Hamiltonian cycle for all of the graphs except for (a). The faults are indicated as usual with dotted lines. In addition to the faults, only those healthy links required to exhibit the Hamiltonian cycle are drawn. Figure 5.17 show the graph of 5.16-(a) with the faulty 3-cycle. The healthy links displayed are required for a Hamiltonian cycle to pass through the nodes lying on the faulty 3-cycle. It is clear form this drawing that there is no Hamiltonian cycle for Q_2^3 in this case.

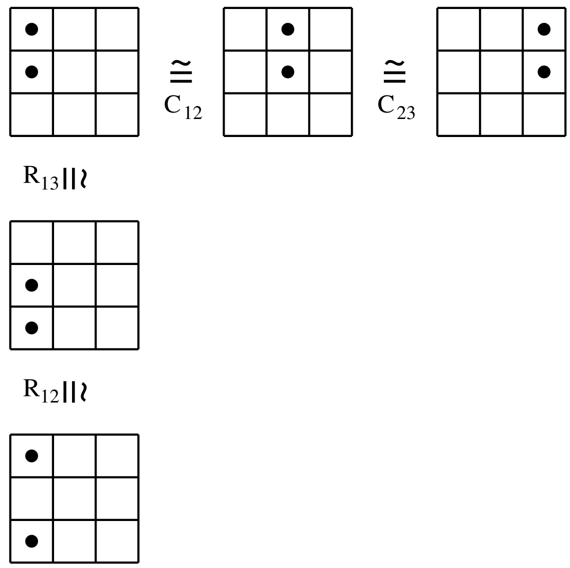

Figure 5.9: The one way to have 2 faults in a single column

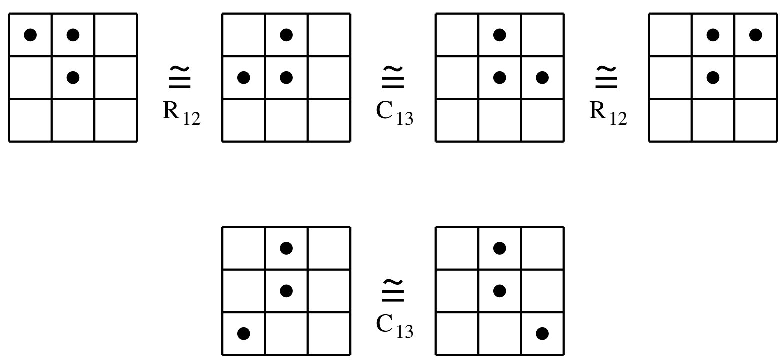

Figure 5.10: The two ways to have 3 faults in two columns

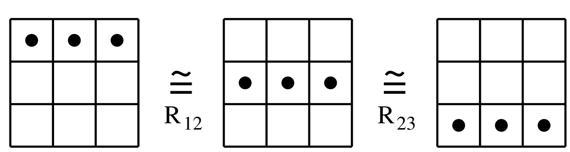

Figure 5.11: The one way to have 3 faults in three columns and one row

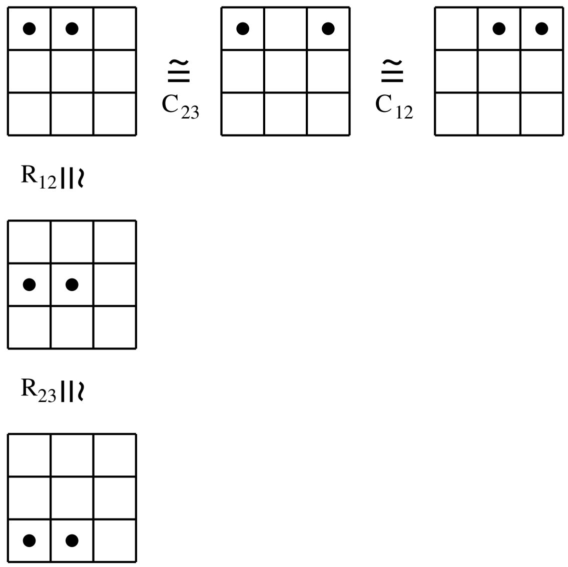

Figure 5.12: The one essential way to have two faults in a single row

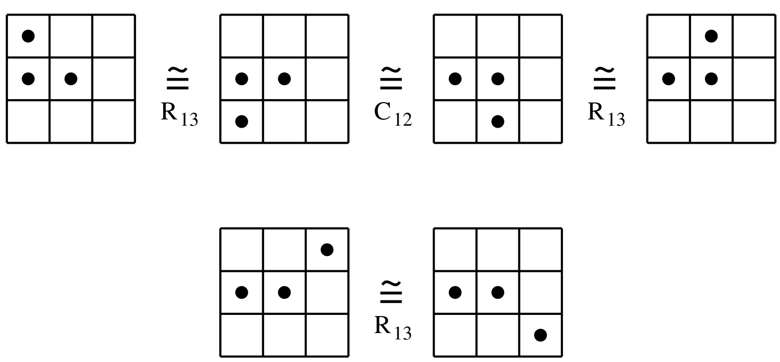

Figure 5.13: The two ways to have 3 faults in two rows

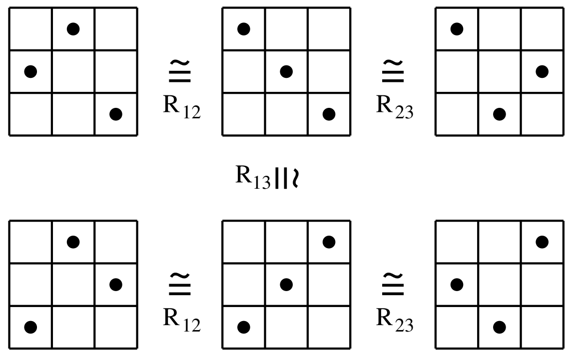

Figure 5.14: The one essential way to have 3 faults in distinct columns and rows

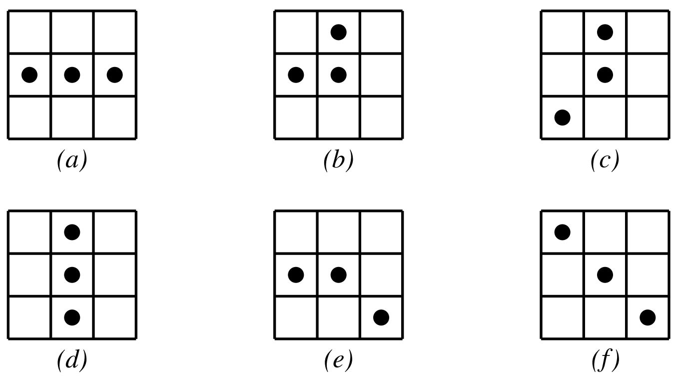

Figure 5.15: The Six Ways to have 3 faults in dim 1

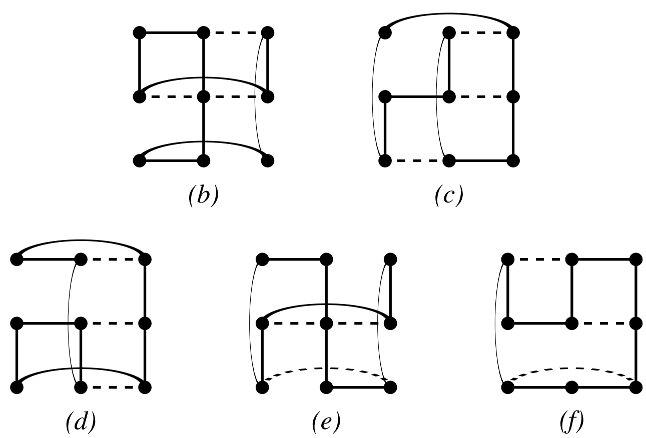

Figure 5.16: The Five Hamiltonian Cycles

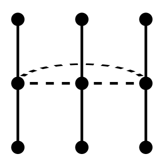

(a)

Figure 5.17: No Hamiltonian Cycle

Chapter 6

The Base Case Q_3^3

We assume we have Q_3^3 with $4{\cdot}(3)-5 = 7$ faults such that each node is incident to at least two healthy links. Our general strategy in this case is to recapitulate the argument of the induction step replacing the induction hypothesis with Lemma 1. In order to use Lemma 1 we must insure that our Q_2^3 subgraphs do not contain a 3-cycle of faults. Thus, we handle this scenario separately. First, note that a 3-cycle of faults must lie within a single dimension. Thus, our troublesome case occurs when our Q_3^3 contains two 3-cycles in distinct dimensions.

6.1 Case i: Q_3^3 contains two 3-cycles in distinct dimensions

Denote the two 3-cycles by C and D. Relabel our Q_3^3 if necessary to have C in dimension one and partition along dimension one. Label the Q_i's such that Q_1 contains as many or more faults than Q_2 which contains as many or more faults than Q_3. This labeling forces D to exist in Q_1. Furthermore, Q_2 contains at most one fault, and Q_3 does not contain any faults.

- Q_1 **contains a Hamiltonian path** $P_1 : y_1 \sim z_1$

We now suppose that the seventh fault lies within Q_1 and consider whether or not the seventh fault is incident to D. We label the Q_3^3 depending upon how D sits within Q_1. We distinguish two nodes in D with the labels y_1 and z_1 such that y_1 is incident to as many faults in Q_1 as either of the other nodes in D and z_1 is arbitrarily assigned to an unlabeled node in D. We now consider whether or not the seventh fault is incident to D. In the case that the seventh fault is incident to D, Figure 6.1 shows the possibilities along with our imposed labeling. Temporarily mark the link $\{y_1, z_1\}$ as healthy. Then we may apply Lemma 1 to our modified Q_1 to get a Hamiltonian cycle that must necessarily contain the link $\{y_1, z_1\}$. Thus we have a Hamiltonian path $P_1 : y_1 \sim z_1$ in the original Q_1.

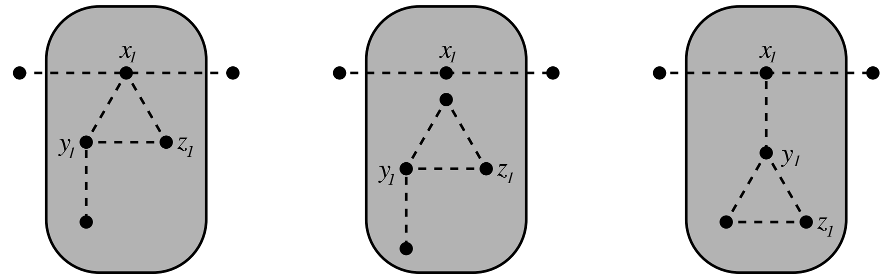

Figure 6.1: Q_3^3 case 1-a: Seventh Fault Is Incident to D

Now suppose the seventh fault is not incident to D. Since Q_1 is the graph Q_2^3, Figure 5.2 shows the two cases for the seventh fault along with P_1.

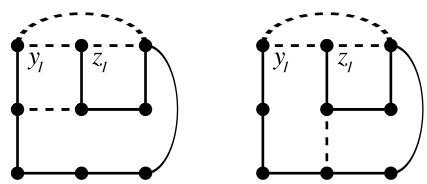

Figure 6.2: Q_3^3 case 1-a: Seventh Fault Not Incident to D

If the seventh fault is not in Q_1, then we temporarily mark a link in Q_1 incident to D as faulty and apply the preceding argument to get the Hamiltonian path $P_1 : y_1 \sim z_1$.

- **The 2-bridge $\{y_1, y_2, z_1, z_2\}$ has healthy dim-1 links**

We have determined six of the seven faults in our Q_3^3. None of the six faults affects the dim-1 links in either of the following 2-bridges: $\{y_1, y_3, z_1, z_3\}$ and $\{y_1, y_2, z_1, z_2\}$. With only one fault not accounted for, one of these 2-bridges must contain healthy dim-1 links. We may relabel the graph if necessary and assume WLOG that the 2-bridge $\{y_1, y_2, z_1, z_2\}$ contains healthy dim-1 links.

- **Q_2 contains a Hamiltonian path $P_2 : y_2 \sim z_2$**

Note that it may be the case that Q_2 contains the seventh fault in our Q_3^3. If the link $\{y_2, z_2\}$ is faulty, then temporarily mark two of the three healthy links in Q_2 incident with y_2 as faulty and the link $\{y_2, z_2\}$ as healthy. Then we may apply Lemma 1 to our modified Q_2 to get a Hamiltonian cycle in Q_2 that must include the link $\{y_2, z_2\}$. Thus we have a Hamiltonian path in the original Q_2 with endpoints y_2 and z_2.

If the link $\{y_2, z_2\}$ is healthy and there is a fault in Q_2 incident to y_2, then mark one of the three healthy links incident to y_2 as faulty, making sure to leave the link $\{y_2, z_2\}$ healthy. Then we may apply Lemma 1 to get a Hamiltonian cycle in our modified Q_2 that must contain the link $\{y_2, z_2\}$. Thus we have a Hamiltonian path in the original Q_2 with endpoints y_2 and z_2. The argument in this paragraph applies mutatis mutandis if there is a fault in Q_2 incident to z_2.

If there is a fault in Q_2 that is not incident to either y_2 or z_2 then y_2 is incident to three healthy links in addition to the link $\{y_2, z_2\}$. Leave the link $\{y_2, z_2\}$ as healthy, and temporarily mark two of the other three links as faulty. It may be the case that the two faults we introduced form a cycle with the fault already present in Q_2. In this case, reinstate one of the temporary faults and mark the third untouched link as faulty. We may now apply Lemma 1 to get a Hamiltonian cycle in our modified Q_2 that must contain the link $\{y_2, z_2\}$. Thus we have Hamiltonian path in the original Q_2 with endpoints y_2 and z_2.

If there is no fault in Q_2, then mark two of the healthy links in Q_2 incident with y_2 as faulty, making sure to leave the link $\{y_2, z_2\}$ remaining as healthy. We may then apply Lemma 1 to get a Hamiltonian cycle in Q_2 that must contain the link $\{y_2, z_2\}$. Thus we have a Hamiltonian path in the original Q_2 with endpoints y_2 and z_2.

Thus, whether or not the seventh fault is in Q_2, we have a Hamiltonian path $P_2 : y_2 \sim z_2$ in Q_2.

We use the healthy links $\{y_1, y_2\}$ and $\{z_1, z_2\}$ to join the paths P_1 and P_2 in our usual construction to get a cycle E that encompasses all of the nodes in Q_1 and Q_2.

- **There is a healthy 2-bridge bewteen E and Q_3**

Recall that Q_3 does not contain any faults. This gives us four 2-bridges between $E \cap Q_2$ and Q_3 that are disjoint in dim-1 links. We need only consider three faults as threats to our 2-bridges since three faults are in Q_1 and one of the faults in C has endpoints in Q_1 and Q_3. Thus one of the four 2-bridges has healthy dim-1 links. We label the 2-bridge $\{u_2, u_3, v_2, v_3\}$. By our labeling of the Q_i's, Q_3 does not contain any faults. We may apply the argument for P_2 above to conclude there is a Hamiltonian path $P_3 : u_3 \sim v_3$ in Q_3. We use our usual construction to get a Hamiltonian cycle in Q_3^3.

6.2 Case ii: Q_3^3 does not contain two 3-cycles in distinct dimensions

If our Q_3^3 contains one or more cycles, then those cycles are restricted to a single dimension that we assume WLOG is dimension one. If our Q_3^3 does not contain a 3-cycle, then some dimension must contain at least three faults, since seven faults are distributed among three dimensions. In this case, we label our Q_3^3 such that dimension one contains at least three faults. We now partition our Q_3^3 along dimension one. We wish to reuse the argument of the induction hypothesis. However, we can only reuse the argument if dimension one contains no more than five faults. So, we handle the case that dimension one contains more than five faults separately.

6.2.1 Case ii-a: Dimension 1 contains more than five faults

In this case, Q_1 contains at most 1 fault with Q_2 and Q_3 not containing any faults. Thus, we may apply Lemma 1 to get a Hamiltonian cycle C_1 in Q_1. We partition C_1 into four disjoint pairs of adjacent nodes to get eight 2-bridges connecting C_1 with Q_2 or Q_3 that only overlap in healthy links in C_1. With at most seven faults in dimension 1 and no faults in Q_2, one of these eight 2-bridges is healthy. Assume WLOG that the 2-bridge $\{u_1, u_2, v_1, v_2\}$ is healthy.

We now establish a Hamiltonian cycle in Q_2 containing the link $\{u_2, v_2\}$. Temporarily mark two of the links in Q_2 incident to v_2 as faulty, making sure to leave the link $\{u_2, v_2\}$ healthy. Then we may apply Lemma 1 to our modified Q_2 to get the Hamiltonian cycle C_2 that must contain the link $\{u_2, v_2\}$. Thus we have our desired Hamiltonian cycle in our original Q_2. We can join C_1 and C_2 using the 2-bridge $\{u_1, u_2, v_1, v_2\}$ to get a cycle D containing all of the nodes in Q_1 and Q_2. We partition D as shown in Figure 3.4 so that $D \cap Q_1$ and $D \cap Q_3$ each consist of four disjoint pairs of adjacent nodes. Then we have eight 2-bridges between D and Q_3 that are disjoint in D. Since there are no faults in Q_3, we do not need to consider if these 2-bridges are disjoint within Q_3. With at most seven faults in dimension 1, one of the 2-bridges is healthy.

The argument that established C_2 applies to get a Hamiltonian cycle C_3 in Q_3 that shares a healthy link with the 2-bridge connecting D and Q_3. We join C_3 and D with our usual construction to get a Hamiltonian cycle in the original Q_3^3.

6.2.2 Case ii-b: Dimension 1 contains no more than five faults

We are now in a position to reuse the argument of the induction hypothesis. We replace the induction hypothesis with Lemma 1. For reference, we label the pieces of the argument as they relate to the cases of the induction step argument.

- IS Case (i): Within each Q_i every node is incident to two healthy links
- IS Case (i-a): None of Q_1, Q_2, or Q_3 contains four faults.

We label the Q_i's such that Q_1 contains at least as many faults as either Q_2 or Q_3. Thus Q_2 and Q_3 each contains at most two faults. With our hypotheses on Q_1, we can apply Lemma 1 to get the Hamiltonian cycle C_1 in Q_1. By our labeling of Q_1, Table 6.2.2 shows the possible combinations of faults that may lie in dimension 1 and Q_2 or Q_3.

Partition C_1 into groups of three nodes as depicted in Figure 3.2 ($k = 3$). We then have six disjoint 3-bridges between C_1 and Q_2 or Q_3. In the worst case

	dim-1 faults	max Q_2 or Q_3 faults
(a)	3	2
(b)	4	1
(c)	5	1

Table 6.1: Possible Combinations of Faults

(Table 6.2.2 (c)) , one of these 3-bridges consists of healthy dim-1 links with one fault in the Q_i with $i \neq 1$. We may relabel our Q_3^3 such that the 3-bridge $\{x_1, x_2, y_1, y_2, z_1, z_2\}$ as depicted in the right-hand image of Figure 3.1 consists of healthy dim-1 links. Furthermore, we relabel Q_3^3 such that if our 3-bridge contains a fault, then $\{x_2, y_2\}$ is faulty.

In either case, there is a Hamiltonian cycle in Q_2 containing the link $\{y_2, z_2\}$. If the link $\{x_2, y_2\}$ is healthy, temporarily mark it as faulty. Mark an additional link incident to y_2 other than $\{y_2, z_2\}$ as faulty. It may be the case that $\{x_2, y_2\}$ was initially healthy and our two introduced faults have formed a 3-cycle with the possible fault already present in Q_2. In this case, reinstate the link that is not $\{x_2, y_2\}$ and mark the fourth link incident to y_2 in Q_2 that is neither $\{x_2, y_2\}$ or $\{y_2, z_2\}$ as faulty. We now have at most three faults in Q_2, and these three faults do not form a 3-cycle. Furthermore, every node in Q_2 is incident to at least two healthy links in Q_2. Thus we may apply Lemma 1 to get a Hamiltonian cycle C_2 in our unmodified Q_2 that must necessarily contain the link $\{y_2, z_2\}$. We may then join C_1 and C_2 via the 2-bridge $\{y_1, y_2, z_1, z_2\}$ to get the cycle D_2 containing all of the nodes in Q_1 and Q_2.

We now partition D_2 into disjoint groups of 3 nodes as depicted in Figure 3.4 with $k = 3$ to get six disjoint 3-bridges connecting D_2 with Q_3. As argued for the 2-bridge between C_1 and either Q_2 or Q_3, one of the 3-bridges between D_2 and Q_3 must consist of healthy dim-1 links with at most one fault in Q_3. Label the 3-bridge as $\{u_1, u_3, v_1, v_3, w_1, w_3\}$ such that $\{u_3, v_3\}$ and $\{v_3, w_3\}$ are links in Q_3. Furthermore, label the 3-bridge such that if the 3-bridge contains a fault Q_3, the link $\{u_3, v_3\}$ is faulty. Then we may apply the argument for the existence of C_2 to get a Hamiltonian cycle C_3 in Q_3 containing the link $\{v_3, w_3\}$. By joining D_2 and C_3 with the 2-bridge $\{v_1, v_3, w_1, w_3\}$, we have a Hamiltonian cycle in our Q_3^3.

This completes IS case (i-a).

- IS case (i-b): Some Q_i contains four faults.

By relabeling our graph if necessary, we may assume WLOG that Q_1 contains four faults.

We begin by establishing the existence of the 2-bridge depicted in Figure 3.5. If two faults in Q_1 are incident to one another, then their respective 2-bridges between Q_1 and Q_2 will share a dim-1 link. If we consider 2-bridges between Q_1 and Q_3, in addition to 2-bridges between Q_1 and Q_2, then regardless how the faults lie in Q_1, we are guaranteed four disjoint 2-bridges connecting Q_1 with either Q_2 or Q_3. With only three dim-1 faults to consider, at least one of the candidate 2-bridges has healthy dim-1 links. We may relabel our Q_3^3 such that our 2-bridge is given by $\{x_1, y_1, x_2, y_2\}$.

We next establish the isomorphic Hamiltonian cycles C_1, C_2, and C_3 in Q_1, Q_2, and Q_3 respectively. Temporarily mark the faulty link $\{x_1, y_1\}$ as healthy. We may then apply Lemma 1 to get the Hamiltonian cycle C_1 that may or may not contain the link $\{x_1, y_1\}$. Under the current case of the induction step, there are no faults in Q_2 or Q_3. Thus, we have the isomorphic cycles C_2 and C_3.

We now exhibit the cycle D_1 containing all of the nodes in Q_1 and Q_2. If $\{x_1, y_1\} \in C_1$ we may then we may construct D_1 by joining C_1 and C_2 with our usual construction. If $\{x_1, y_2\} \notin C_1$, then we consider the four disjoint 2-bridges that exist between C_1 and C_2. With only the three dim-1 faults to consider, one of these four 2-bridges must be healthy. We then construct D_1 using this healthy 2-bridge. To complete the Hamiltonian cycle for our Q_3^3, we establish a 2-bridge between D_1 and C_3. We can partition $D_1 \cap Q_2$ into four pairs of disjoint nodes which gives us four disjoint 2-bridges between D_1 and Q_3. With only the three dim-1 faults to consider, one of these four 2-bridges is healthy. Thus, we may complete the Hamiltonian cycle of Q_3^3 with our usual construction.

This completes the proof for IS case (i-b) and case (i).

- IS case(ii): Some node x_i is incident to only one healthy link in some Q_i

We relabel our Q_3^3 if necessary to assume WLOG that x_1 is incident to only one healthy link in Q_1. We furthermore assume WLOG that $\{x_1, y_1\}$ is a healthy link.

- IS case (ii-a): There exists a node w_1 such that $\{x_1, w_1\}$ is faulty and $\{w_1, w_2\}$ is healthy

We begin by establishing a Hamiltonian path in $P_1 : x_1 \sim w_1$ in Q_1. Temporarily mark the link $\{x_1, w_1\}$ as healthy. We may now apply Lemma 1 to our modified Q_1 since every node is incident to at least two healthy links, there are no more than three faults, and there is no 3-cycle of faults. Thus, we have a Hamiltonian cycle in our modified Q_1. By the nature of the links incident to x_1, the Hamiltonian cycle must contain the link $\{x_1, w_1\}$. Thus we have the Hamiltonian path P_1 in our unmodified Q_1.

We next establish the Hamiltonian path $P_2 : x_2 \sim w_2$ in Q_2. At most Q_2 contains one fault. If necessary, mark the link $\{x_2, w_2\}$ as healthy. Choose a healthy link $\{x_2, u_2\}$ and mark the remaining two links in Q_2 incident to x_2 as faulty. It is possible that $\{x_2, w_2\}$ was healthy initially and that the two faults we introduced formed a 3-cycle with the possible fault in Q_2. In this case, we reinstate one of the links we marked as faulty, and mark the link $\{x_2, u_2\}$ as faulty. We have introduced at most two faults to Q_2 which limits us to a maximum of three faults in our modified Q_2. We have arranged the faults in Q_2 so that there is no 3-cycle of faults, and every node in Q_2 is incident to at least two healthy links in Q_2. Thus, we may apply Lemma 1 to get a Hamiltonian cycle in our modified Q_2. Whether or not $\{x_2, w_2\}$ is healthy in our original Q_2, we have our desired Hamiltonian path P_2.

We now use the 2-bridge $\{x_1, w_1, x_2, w_2\}$ with our usual construction to form the cycle D_1 containing all of the nodes in Q_1 and Q_2.

To complete the construction of a Hamiltonian cycle in our Q_3^3, we establish a 2-bridge between D_1 and Q_3. We have eight potential 2-bridges: four between $D_1 \cap Q_1$ and Q_3 and four between $D_1 \cap Q_2$ and Q_3. With three faults committed to Q_1, we only have four faults to consider. So, at least one of these 2-bridges is healthy. Label the link of our 2-bridge that is in Q_3 by $\{v_3, z_3\}$. We may apply the argument above that established P_2 to get the Hamiltonian path $P_3 : v_3 \sim z_3\}$. We complete our desired Hamiltonian cycle by joining D_1 and P_3 with our usual construction.

This completes the proof for IS case(ii-a).

- IS case(ii-b): There is no w_1 such that $\{x_1, w_1\}$ is faulty and $\{w_1, w_2\}$ is healthy.

Since x_1 is incident to three faults in Q_1, our current case determines all but one of the faults in our Q_3^3. With only one fault left undetermined, at least one of the three faults incident to x_1 in Q_1 must be incident to a healthy link between Q_1 and Q_3. Thus, if the link $\{x_1, x_3\}$ is healthy, we are back in IS case (ii-a) by symmetry.

So, for the remainder of this case, we assume $\{x_1, x_3\}$ is faulty which gives us a complete accounting of the seven faults in our Q_3^3. Furthermore, in this case, Figure 3.10 depicts the Hamiltonian cycle for Q_3^3.

This completes the proof for IS case (ii-b) and case (ii).

- IS case(iii): Some node x_i is incident to zero healthy links in Q_i

We assume WLOG that x_1 in not incident to any healthy links in Q_1. Thus the links $\{x_1, x_2\}$ and $\{x_1, x_3\}$ must be healthy.

We begin by arguing for the existence of one of the configurations depicted in Figure 3.11. We have two disjoint pairs of faults incident to x_1 in Q_1. Each such pair of faults requires two dim-1 faults to prevent both configurations shown in Figure 3.11. With only three faults in dimension 1, one of the desired configurations survives the dim-1 faults. By relabeling the graph if necessary, assume we have the configuration depicted in Figure 3.12.

We begin by establishing the path $P_1 : y_1 \sim z_1$ in Q_1 that contains all of the nodes in Q_1 except x_1. If we temporarily mark the links $\{x_1, y_1\}$ and $\{x_1, z_1\}$ as healthy then we may apply Lemma 1, since every node in our modified Q_1 is incident to two healthy links and there are only two faults. Thus we get a Hamiltonian cycle in our modified Q_1 that must contain the links we marked as healthy. Thus we have our desired path P_1.

We now show there is a Hamiltonian path $P_2 : x_2 \sim y_2$ in Q_2. Mark two of the links incident to x_2 in Q_2 as faulty. Make sure to leave the link $\{x_2, y_2\}$ as healthy. Then we may apply Lemma 1 to our modified Q_2 to get a Hamiltonian cycle that contains the link $\{x_2, y_2\}$. Thus we have our desired Hamiltonian path P_2 in the original Q_2. Similarly, we have a Hamiltonian path $P_3 : x_3 \sim z_3$ in Q_3.

The path $x_1 \cdot x_2 \cdot P_2 \cdot y_1 \cdot P_1 \cdot z_3 \cdot P_3 \cdot x_1$ explicitly describes the Hamiltonian cycle in our Q_3^3.

This completes the proof for IS case (iii), and case (ii) of Q_3^3.

Chapter 7

Conclusion

Now that we have completed the proof of the main theorem, we consider whether or not we can, under the assumptions of the theorem, strengthen the bound $4n-5$ to include additional faults. Figure 7.1 depicts a configuration in which $4n-4$ faults preclude the existence of a Hamiltonian cycle. Note that the square subgraph depicted always occurs as a subgraph of Q_n^k for our values of k and n. (For example, consider the 4-cycle $\{(0,0,\star),(1,0,\star),(1,1,\star),(0,1,\star)\}$ where $\star$ represents the remaining entries in the n-tuples).

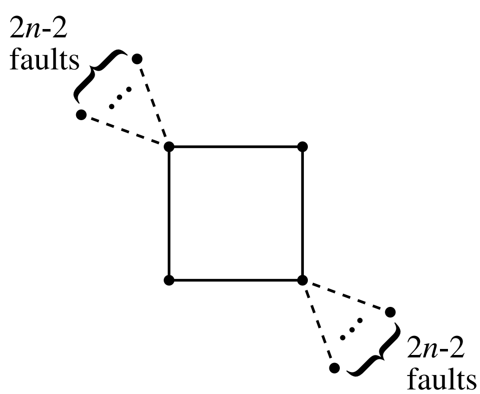

Figure 7.1: No Hamiltonian Cycle

To see that there does not exist a Hamiltonian cycle in the presence of this configuration, observe that the depicted faults force a Hamiltonian cycle to contain all four of the edges in the square. Since any Hamiltonian cycle must account for exactly two edges at each vertex, none of the four vertices above can be incident to an edge joining it with the rest of the graph. It follows that the bound $4n-5$ is sharp.

Possible directions for further research include considering a probabilistic analysis in which we consider a random collection of more than $4n - 5$ faults and determine the likelihood of the existence of a Hamiltonian cycle.

Chapter 8

Acknowledgements

My thanks to John Caughman for his invaluable input on this writing and encouragement in all things mathematical. To the many professors in the math department who added to my training, I am truly grateful for the education I have received from the Portland State University mathematics and statistics department.

The figures were crafted with the Postscript programming language. The knowledge to do so came from seeds planted in the Mathematical Science Research Institute (http://www.msri.org) 2003 workshop on mathematical graphics. Notes on the two-week workshop are on the websites of the conductors: Bill Casselman (http://www.math.ubc.ca/∼cass) and David Austin (http://mersanger.math.gvsu.edu/david).

Bibliography

[1] Y.A. Ashir and I.A. Stewart. Fault-tolerant embeddings of Hamiltonian circuits in k-ary n-cubes. *SIAM J. Discrete Math*, 15:317 – 328, 2002.

[2] B. Bose, B. Broeg, Y. Kwon, and Y. Ashir. Lee distance and topological properties of k-ary n-cubes. *IEEE Transactions on Computers*, 44:1021 – 1030, 1995.

[3] M.Y. Chan and S.-J. Lee. On the existence of Hamiltonian circuits in faulty hypercubes. *SIAM J. Discrete Math*, 4:511 – 527, 1991.

[4] S. Lakshmivarahan and S. K. Dhall. *Analysis and Design of Parallel Algorithms: Arithmetic and Matrix Problems.* McGraw-Hill, New York, 1990.

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

[6] Miro Kraetzl Miltos D. Grammatikakis, D. Frank Hsu. *Parallel Systems Interconnections and Communications.* CRC, Boca Raton, 2000.

[7] Robert Suaya and Eds. Graham Birtwistle. *VLSI and Parallel Computation Frontiers.* Morgan Kaufmann Publishers, Inc., San Mateo, 1990.

[8] D. West. *Introduction To Graph Theory, Second Edition.* Prentice Hall, New Jersey, 2001.