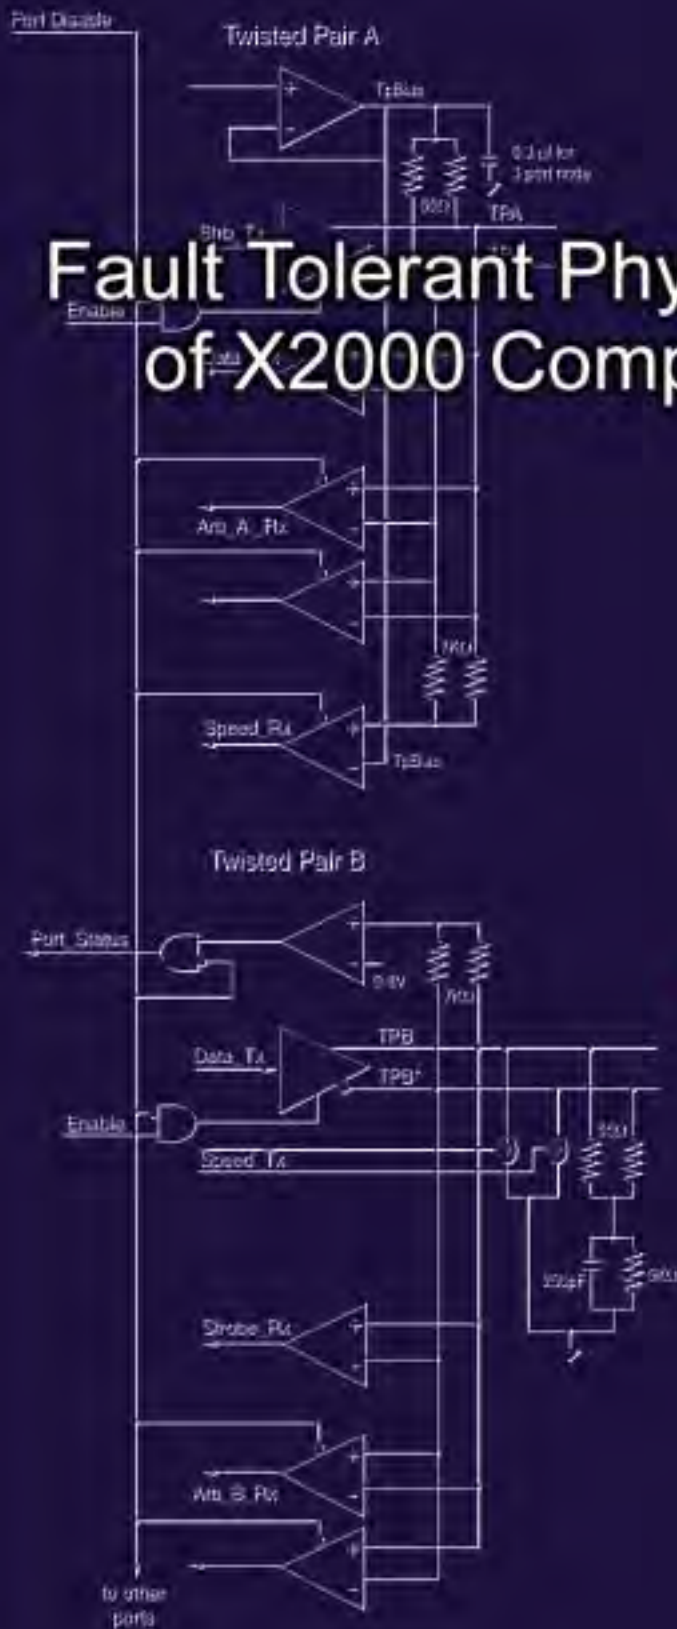


Fault Tolerant Physical Interconnection of X2000 Computational Avionics



Laurence E. LaForge

Jet Propulsion Laboratory
Document Number JPL D-16485




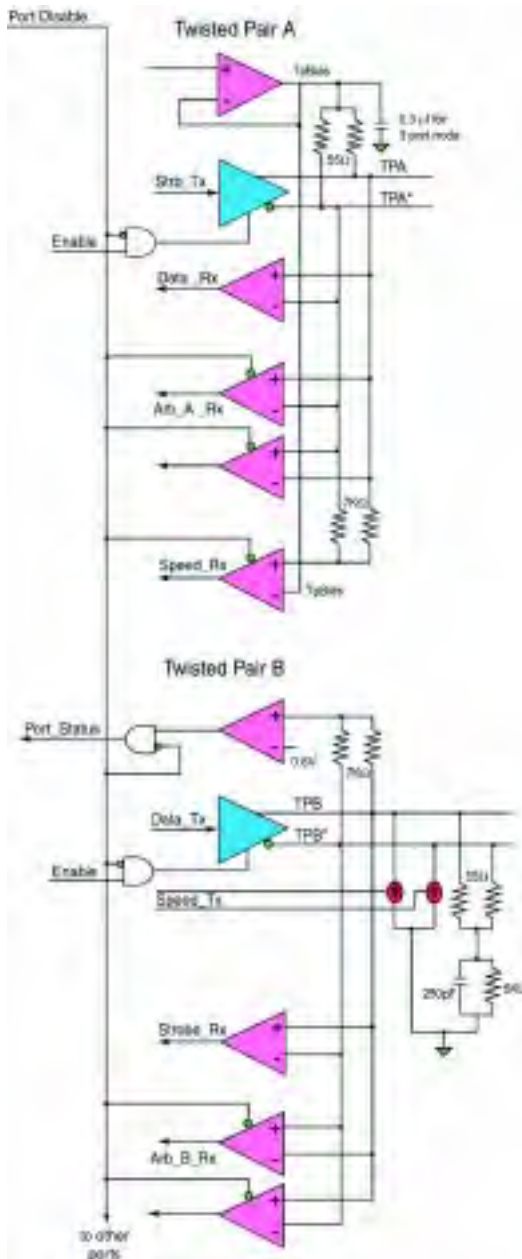
Fault Tolerant Physical Interconnection of X2000 Computational Avionics

Laurence E. LaForge

NASA/ASEE Summer Faculty Fellow
Embry-Riddle Aeronautical University

Jet Propulsion Laboratory
Document JPL D-16485

28-August-1998
revised 18-Oct-1999





Fault Tolerant Physical Interconnection of X2000 Computational Avionics

Copyright © 1999 Laurence E. LaForge

Fallon Naval Air Station
OEC / Code 100
Barracks 2 / Building 305
Fallon, NV 89496

NASA/ASEE Summer Faculty Fellow
Embry-Riddle Aeronautical University
28-Aug-1998, revised 13-Sep-1998,
30-Oct-1998, 16-Nov-1998, 8-Jan-1999,
26-Aug-1999, 18-Oct-1999

Tel: (775) 322-5186
Fax: (775) 423-0591

Larry@The-Right-Stuff.com

Table of Contents

1. Executive Summary.....	2
2. Findings and Recommendations	4
3. Fault Tolerance by Diagnosis and Configuration	8
3.1 Quorums from Trees, Cycles, and Cliques	9
3.2 General Lower Bound on Quorum Radius	14
3.3 Quorums from K-cubes	15
3.4 Quorums from K-cube-connected Cycles	24
3.5 Quorums from K-cube-connected Edges	40
3.6 Chordal Graphs and Cycles of K-cubes	42
3.7 Quorums from C-cubes	44
3.8 Choosing a Graph Architecture	52
3.9 Underware for Distributed Configuration	58
3.10 Application to X2000	65
A. References	71
A.1 Related NASA Documents	71
A.2 Related Books and Dissertations	71
A.3 Related Papers and Articles	72

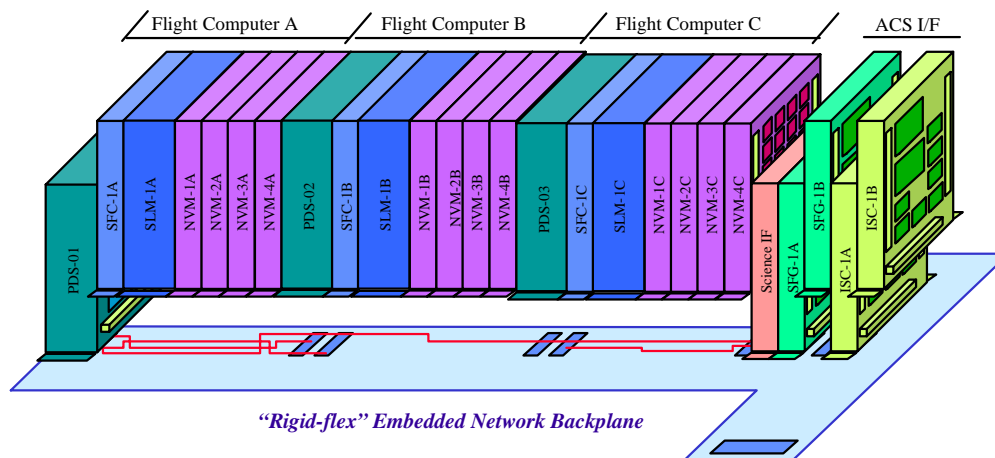


Figure 1: Proposed packaging for X2000 computational avionics ([Steiner 11-Mar-1997]).



Cover schematic reprinted with permission from [Anderson 1998], graph from Figure 20. Hardcover design by Derek Carlson.





1. Executive Summary

In early June of 1998 Glenn Reeves asked me to analyze the fault tolerance of the proposed bus structure for X2000 avionics. This report reflects my attendance at the 11-Jun-1998 X2000 avionics architecture review, as well as interviews with Savio Chau, Bob Barry, Bob Rasmussen, Don Hunter, and Carl Steiner. I have studied more than sixteen project documents, and have surveyed the related literature. Table 1 synthesizes my observations, conclusions, and recommendations, and is based on the following priorities:

- a. What do we want to build? That is, to what extent are *requirements* for bus fault tolerance
 - i) clear?
 - ii) complete?
 - iii) self-consistent?
- b. How will we build it? That is, to what extent is the *architecture* for bus fault tolerance
 - i) clear?
 - ii) complete?
 - iii) self-consistent?
- c. To what extent does the *what* match the *how*? Are requirements consistent with architecture?
- d. Does the architecture make best use of fault tolerant technology?

Somewhat surprisingly, criterion (d) seems to govern most people’s thinking, with relatively less emphasis on the operational utility of (a), (b), and (c). This is of concern since it is at best difficult to assess requirements or architectures that are neither well-specified nor complete. On a positive note, the requirements and architectural specifications that are in place can be augmented to provide operational utility.

Observations and conclusions	Overall recommendations	Details in
Tendency to overemphasize technology, under-emphasize requirements and architecture.	Broaden and deepen existing requirements and architectural specifications with clear, complete, consistent operational descriptions. (Estimate 180 hours to complete).	Tables 2, 3, 4, 5
Portions of requirements and architecture unclear, incomplete, or inconsistent.		
Viewgraphs are insufficient for capturing requirements and specifications.		
Diagnosis and configuration with software and back-door I ² C bus does not achieve physical level fault tolerance; cost and risk greater than with self-configuring 1394 bus alone.	Employ “underware” for mutual test, diagnosis, and distributed configuration of 1394 buses. Omit 50% of 1394 wires. Eliminate back door I ² C, but keep essentials of high level diagnosis.	Tables 4, 5 Section 3

Table 1: General observations, conclusions, and recommendations. X2000 will benefit more from clear, complete, and consistent requirements and specifications than from improvements in bus fault tolerance.

There is a wide variation in the clarity, completeness, and consistency of X2000 requirements and architectural specifications for bus fault tolerance. On the plus side, almost all of the requirements that are clear are also self-consistent. Five of the sixteen project documents that I examined contain explanatory narratives; three of these five are project documents, and two are viewgraph presentations. The remaining documents, all viewgraph presentations, lack explanatory narratives. Many of the gaps in requirements and architectural specifications reflect a tendency to use viewgraphs in place of explanatory narratives. I frequently found that different people interpreted the same viewgraphs quite differently. For this reason I conclude that viewgraphs are *not* adequate for capturing and communicating requirements and specifications. In part, “design-by-viewgraph” appears to be a reaction against the perceived overspecification of *Cassini*. Some people justified design-by-viewgraph by pointing to the success of *Mars Pathfinder*, wherein capture of requirements and architecture was minimized. Another justification, spoken by almost everyone I interviewed, is based on the perception that people have more work to do, in less time, than ever before.





As to criterion (c), the architecture for bus fault tolerance, such as it can be gleaned, is not entirely consistent with project requirements for bus fault tolerance. The good news is that both requirements and architectural specifications for fault tolerant bus interconnection can, in a straightforward manner, be rendered clear, complete, self-consistent, and consistent with each other. Assuming cooperation on the part of hardware and software groups, I estimate that it will take about 180 hours of work to accomplish this, with about one-third of the work devoted to requirements.

Technical aspects of this report focus on criterion (d). To diagnose and configure in the presence of faults in computational nodes of the bus, the current approach ([Charlan et al 11-Jun-1998], Option D) combines a software lock and key with a “back-door” I²C bus. Such an add-on approach is at odds with X2000 written policies for built-in fault tolerance. My conclusion is that the current approach exposes the physical interconnection to a level of risk that is not commensurate with X2000 project goals. The good news is that the existing high-level approach can be modified to achieve robustness across the four layers of the 1394 protocol,¹ *without* a (redundant) back-door I²C bus ([Charlan et al 11-Jun-1998], Option C). Reallocating resources to the design and test of such “underware” would give more fault tolerance per dollar, and at the same time save at least two I²C wires per computational node. In keeping with goals of X2000 engineers, this underware requires no modifications to the 1394 hardware.

Refer to Figure 1. Since the switching functions of the 1394 are built into the avionics nodes themselves, a particular concern is the tolerance of the bus to nodes whose switching functions fail. Such failures tend to partition the bus. In this regard *the point-to-point connectivity as presented at the 11-Jun-1998 core avionics design review is single fault tolerant*, even if a back-door I²C bus scheme works perfectly. At a cost of six 1394 ports (36 wires) per node - we are substantially overpaying for single fault tolerance.² This report recommends halving the present 1394 bus wirecount (from 36 down to 18 wires per node), *at the same time doubling the tolerance to partitioning faults (up from one to two)*. Configuration in the presence of partitioning faults can be modeled as a bivariate optimization problem in extremal graph theory:

What $(f+1)$ -connected graphs with fewest edges minimize the maximum radius or diameter of trees spanning the quorums induced by deleting up to f of the n original vertices?

Here n is the number of nodes and f is the number of faults we want to tolerate, in the worst case. Minimizing the maximum number of hops between nodes in the tree configured is the same as minimizing the diameter, and essentially the same as minimizing the radius. As an absolute limit, the 1394 specification allows at most 16 hops between nodes [P1394 1995].

The remainder of this report is organized as follows. Section 2 synthesizes findings and recommendations with respect to project processes and procedures. Section 3 comprises the technical exposition. Sections 3.1 through 3.8 furnish lower and upper bounds on the radius and diameter of quorums for architectures based on stars, cycles, cliques, K-cubes, and C-cubes. I show how, with the exception of C-cubes, these structures are absolutely or asymptotically optimum. Section 3.9 illustrates how to formulate and analyze parallel algorithms for distributed diagnosis and configuration. In the presence of both partitioning faults and babbling nodes, these algorithms minimize the radius or diameter of a 1394 bus, *without the need for a back-door I²C bus*. Section 3.10 shows how to apply the theorems, formulae, and algorithms of Sections 3.1 through 3.9 to architectures capable of tolerating one, two, or three faults. To help the designer I have supplemented this report with an Excel workbook. For given number of nodes n and fault tolerance f , GRAFT (GRaph Architecture Fault Tolerance calculator) recommends an architecture with minimum number of point-to-point connections. GRAFT also reports the radius and diameter of quorums induced, as a function of the actual number of faults. Sections 3.8 and 3.10 describe how to use GRAFT.

Despite the technical emphasis of the bulk of this report, X2000 will benefit more from clear, complete, and consistent requirements and specifications than from improvements in bus fault tolerance.

1. The 1394 bus prescribes four layers of protocol: physical link, transaction, and bus management ([P1394 1995], Chapter 3). In effect, the MDS application will add at least one layer to this.
2. At six wires per port: twisted pairs A and B, plus power and ground [P1394 1995].





2. Findings and Recommendations

In the spirit of *open communications* and *independent review* ([Woerner, Spear, Parker 7-Aug-1998], 2.5.3; [Kemski 14-Jul-1998], 3.3.1), Tables 2 through 4 detail my findings with respect to evaluation criteria (a) through (d) as listed in Section 1. Table 5 lists my specific recommendations.

	I. Requirements	II. References	III. Points where requirements are unclear, incomplete, or inconsistent
A.	Significant risk list (SRL). “...Identified risks and associated decisions to either accept, mitigate, or eliminate those risks will be recorded in a Significant Risk List (SRL) ...”	Project Implementation Plan, 2.9 [Woerner, Spear, Parker 7-Aug-1998]	SRL does not appear to exist, despite risks (in part as described in this report). For example, it is apparently yet undocumented that a flight computer is <i>required</i> to execute the MDS software; this represents a significant risk to mission flyings with a single flight computer.
B.	Single point failure. “... The FDP design shall tolerate single faults. The definition of a fault shall include hardware failures (e.g., devices and sensors) and software failures (e.g., Single Event Upset, software bugs, sequence errors, and bad commands). No single fault shall result in the loss of a mission critical function ...”	Project Implementation Plan, 3.8.3.1 [Woerner, Spear, Parker 7-Aug-1998] Level 3 Requirements, 3-2141 [Guiar 23-Jul-1998]	Unclear what is meant by a “point”. Is this some location within a “fault containment region”? (See also row B of Table 3).
C.	Fault monitoring. “... The FDP design shall provide a method for detecting system failure modes. The preferred method for the detection shall be supplied within the subsystem component through the use of built in self-testing ...”	Project Implementation Plan, 3.8.3.4 [Woerner, Spear, Parker 7-Aug-1998]	Proposed software lock and key mechanism is not within the 1394 physical layer subsystem; proposed mechanism does not detect low level bus failure modes (e.g., switches stuck closed or faulty bus circuits). This report recommends <i>mutual</i> test and diagnosis [LaForge and Korver 1997]. See also Sections 3, 3.9.
D.	Reliability. “... The following reliability analyses shall utilize the methodology stated in JPL D-5703 or PEM/MAM approved methodologies...” Reliability/cost now highest priorities for X2000.	Mission Assurance Plan, 3.3 [Kemski 14-Jul-1998] Design Approach/Priorities [Guiar 11-Jun-1998]	References [JPL D-5703 23-Jul-1998] as containing methods for analysis; it does not. More relevant is [JPL 4-11 1-Apr-1984]. Unclear what the PEM/MAM approved methodologies are, or even if they exist. Lack of FMECA’s suggests reliability is not a highest priority
E.	Success-critical single failure point (SFP) “All system SFP’s shall be identified ...”	Mission Assurance Plan, 3.3.2 [Kemski 14-Jul-1998]	Apparently no list of SFP’s exist. However, this report serves to begin such a list
F.	FMECA’s. “... FMECA’s shall be performed and documented ...”	Mission Assurance Plan, 3.3.3.1 [Kemski 14-Jul-1998]	No FMECA published for avionics bus fault tolerance; this is at odds with 3.3.1: “... All analyses shall be maintained in a current state and reflect the currently approved design ...”

Table 2: Findings with respect to Project Implementation Plan and Mission Assurance Plan requirements.





With respect to Row A of Table 2, having a significant risk list (SRL) is an excellent idea. I recommend that an SRL be composed and, even if empty, placed in the online project library. The list should be readily locatable using the online library search engine and the keywords "SRL" or "Significant Risk List".

Row B of Table 2 and row B of Table 3 underscore a very important point. A clear analysis of fault tolerance requires a clearly defined fault model. Most people I interviewed claimed to understand "single fault" and "point of failure". When pressed, however, most could not articulate the meaning of these phrases in the context of X2000; those who did give definitions spoke in terms of "fault containment regions". I recommend that the respective definitions be unambiguously fleshed out, if necessary, by an exhaustive list of "faults, "points of failure", and "fault containment regions". First drafts of these lists are given by column III of row B of Table 3, and by row C of Table 3.

	I. Requirements	II. References	III. Points where requirements are unclear, incomplete, or inconsistent
A.	Fault protection built in, not on. "...X2000 has a major goal of building fault protection (a.k.a. redundancy management, fault tolerance, goal oriented commanding) INTO the spacecraft via its subsystem and subsystem components, rather than adding FP on top of the normal spacecraft functions. ..."	Level 3 Requirements 3-2643, 3-2139 [Guiar 23-Jul-1998]	Proposed software lock and key mechanism is directly at variance with this policy. Current design is on top of, not within, the 1394 physical layer subsystem; proposed mechanism does not detect low level bus failure modes (e.g., switches stuck closed or faulty bus circuits). This report recommends incorporating fault tolerance via <i>underware</i> . See also Sections 3, 3.9.
B.	Fault containment regions. "... All subsystems and block redundant units within subsystems shall be designed to be fault containment regions. Fault containment regions shall be designed such that any fault from the spacecraft fault set occurring in a fault containment region shall not propagate faults or undetectable errors into other fault containment regions. ..."	Level 3 Requirements 3-2147, 13-12, 13-13 [Guiar 23-Jul-1998]; Fault containment regions [Guiar Jun-1998]	Unclear what is meant by a "fault containment region" and "spacecraft fault set". Diagram is unaccompanied by narrative explanation. According to Savio Chau, each slice containing a microcontroller and its local memory is a fault containment region; each flight computer (3 slices) is a fault containment region; each pair consisting of a peripheral slice plus its respective controller slice is a fault containment region; within any PASM module inside a PCU slice, each switch is a fault containment region; the assemblage of alternative "global memory" (volatile and nonvolatile, all slices together) is a fault containment region.
C.	Faults tolerated. "... All permanent stuck-at faults ..." "... Any bridging fault (such as a short) occurring within, but not between an element of a redundant system ..."	Level 3 Requirements 3-2149, 3-2151 [Guiar 23-Jul-1998]	Proposed software lock and key mechanism is not tolerant to switches stuck closed or faulty bus circuits. Unclear what is meant by a bridging fault occurring "between an element of a redundant system". Is this related to fault containment regions?
D.	Distributed fault tolerance. "... Use low level, behavioral/reflexive fault detection and response where feasible... Use centralized (heuristic) fault detection and response where required ..."	Level 3 Requirements 3-2156, 3-2157, 13-16, 13-17 [Guiar 23-Jul-1998] [Barry 22-Jan-1998] [Barry 26-Feb-1998]	Proposed back-door mechanism for diagnosing and configuring faulty nodes on the bus is <i>centralized</i> . The approach recommended by this report is low-level behavioral/reflexive fault detection.

Table 3: Findings with respect to Level III and fault protection policy requirements.





As excerpted in row C of Table 2, the fault monitoring policy prescribed by the Project Implementation Plan reiterates a longstanding maxim for testing: systems are themselves best suited to detect faults in systems of like kind. “Like kind” includes details of interfaces, signal formats, and timing. The proposed lock and key mechanism is well-suited for detecting failures in low-level software. This mechanism is much less well suited to detecting failures at the transistor or gate level, or at the level of Goal Achieving Modules (GAMS). To be proper, we should design physical layer mutual test and diagnosis among bus nodes. However, in deference of the desire of project engineers not to modify bus controller hardware, I propose a version of mutual test and diagnosis that spans all layers of the 1394. An advantage of such high-level diagnosis is that the overall probability of fault detection is increased. A disadvantage is that the probability of fine-grained fault isolation is decreased. Mutual test and diagnosis is a superset of built in self-test, and dovetails well with point-to-point interconnections such as that used to support the redundancy of the 1394 bus [LaForge and Korver 1997].

	I. Architectural specifications	II. References	III. Points where specifications are unclear, incomplete, inconsistent, or do not reflect best use of fault tolerant technology.
A.	Fail silent. Software on each node sends out key on I ² C bus; opens connection to bus if key not retrieved. If key is sent and retrieved then lower layers on node are likely to work.	Symmetric Architecture [Rasmussen 11-Jun-1998]	Begs the question by relying on faulty nodes to dissociate themselves from I ² C bus. Nodes with SDA or SCL switch stuck closed (either due to hardware or software) will not dissociate themselves; fault propagates to all other nodes. Contravenes policies for significant risk, single fault tolerance, stuck-at faults, propagation of faults (cf. rows A and B of Table 2, rows B and C of Table 3, [Paret and Fenger 1997], Fig 3.1)
B.	Extra FET between node and I²C, 1394 bus. Prevent power shorts by adding pass transistor between ground and I ² C. Add pass transistor in 1394 physical layer circuitry.	Bus Tiger Team [Chau 11-Jun-1998] [Charlan et al 11-Jun-1998]	Unclear how pass gate is controlled. Extra pass gates increase the chances that nodes will be dissociated from the bus. Failure of FET's in series within any node (i.e., within some fault containment region) propagates fault across multidrop I ² C bus to all other nodes. This contravenes policies for significant risk, single fault tolerance, stuck-at faults, propagation of faults (cf. rows A and B of Table 2, rows B and C of Table 3, [Paret and Fenger 1997], Fig 3.1)
C.	Diagnosis of faulty nodes. 1394 root polls nodes through I ² C, relies on self-diagnosis of NUT (“node under test”)	Backup: Upstream Connection Failed [Chau 18-Aug-1998]	Unclear how and how well this works. Depends on 1394 root working properly. Contravenes policies for significant risk, single fault tolerance, stuck-at faults, propagation of faults, and decentralized detection (cf. rows A, B, and C of Table 2, rows B, C, and D of Table 3)
D.	Configuration of 1394. Duplicate buses, with pre-designated roots, leaves, and interior nodes	Bus Tiger Team [Chau 11-Jun-1998] [Charlan et al 11-Jun-1998] [Chau 17-Apr-1998]	This report recommends 1394 underware that uses dynamic configuration to achieve <u>twice</u> as much tolerance to partitioning faults, with <u>half</u> as many wires. This meets shuttle requirements for 2-fault tolerance (Europa orbiter), and at the same time configures from <i>n</i> nodes a tree having radius at most $1+n/4$. Cf. Table 18
E.	FMECA's.	Bus Failure Modes [Chau and Holmberg 17-Apr-1998]	FMECA's not yet performed. In at least 4 modes, (1a, 1b, 1e, 14), a single fault can cause failure of the entire avionics.

Table 4: Findings with respect to proposed architecture for bus fault tolerance.





Rows D and F of Table 2 concern FMECA's. *I recommend that FMECA's be carried out on all of the failure modes referenced in column II of row E of Table 4, and that the corresponding reports be posted in the online project library. I also recommend that Section 3.3 of the Mission Assurance Plan be updated to reference a document that, in fact, contains methodologies for performing FMECA's. I also recommend that each project element manager draft a policy delineating, perhaps by reference, the FMECA methodology for the respective portion of the project.*

Commensurate with row E of Table 2, *I recommend that a list of success-critical single failure points (SFPs) be composed and placed in the online project library.* Beginnings of such a list are given in column III of row A of Table 2. *I furthermore recommend that each of the X2000 documents listed in Section A be placed into the online project library.* It appears that eight of these documents are not yet online. Moreover, it appears that the bulk of X2000 meetings are not captured by written minutes. As a result, a great deal of intellectual effort is lost. *I therefore recommend that it be standard practice for each meeting's moderator to ensure that someone takes comprehensive minutes, and places these minutes in the online project library.* The benefit will be a clearer understanding of which issues are resolved, which are open, and who is supposed to do what, when.

Rows A and D of Table 3, and all the rows of Table 4, are related to issues raised by row C of Table 2 and discussed briefly on the preceding page. These issues may be summarized as 1) *in situ* decentralized diagnosis, at the appropriate level and 2) decentralized configuration, at the appropriate level. In lieu of a back-door I²C bus, *I recommend that (cf. Sections 3.9 and 3.10) diagnosis and configuration be implemented as underware.* The benefits of adopting such underware are substantial. Refer to Figure 1. According to Don Hunter and Carl Steiner, about 108 of the 112 pins in y-axis connectors of the avionics slices are allocated. This leaves an uncomfortable margin for enhancements. The approach I propose would free at least 20 pins per node, thus giving designers a bit more breathing room. The underware I propose is an application of the theory of diagnosis and configuration, and is described in Section 3.

	Specific recommendations	Findings, references
1.	Compose Significant Risk List (SRL); place in online project library.	Table 2, row A
2.	Spell out precise meaning of "fault", "point of failure", "fault containment region".	Table 2, row B; Table 3, rows B and C
3.	Implement mutual test and diagnosis among bus nodes.	Table 2, row C
4.	Update Section 3.3 of Mission Assurance Plan to reference a document that, in fact, contains methodologies for performing FMECA's.	Table 2, row D
5.	Each project element manager draft a policy delineating, perhaps by reference, the FMECA methodology for the respective project element.	
6.	Perform FMECA's on all failure modes referenced in Table 4, row E, column II; post reports in online project library.	Table 2, row F Table 4, row E
7.	Compose list of success-critical single failure points (SFPs); place in online project library.	Table 2, rows A and E
8.	Maintain <i>all</i> relevant X2000 documents in the online project library.	Appendix A
9.	Every X2000 meeting should be captured by minutes, which should then be placed in the online project library.	
10.	Substitute diagnosis and configuration, via 1394 <i>underware</i> , for back-door I ² C bus. Reduce number of ports per node.	Table 2, row C, Table 3, rows A and D, Table 4; Sections 3, 3.9, 3.10

Table 5: Summary of specific recommendations.





3. Fault Tolerance by Diagnosis and Configuration

Figure 2 synthesizes diagnosis and configuration [LaForge 1997]. Our recommendations incorporate several variations on this theme: distributed algorithms for diagnosis [Somani and Agarwal 1987], diagnosis interleaved with configuration [Preparata et al 1967], and degradable architectures [Koren and Pradhan 1986]. Physically, the architecture may be large or small; it may be a network of computers, or a systolic array on a chip. For the X2000 bus we adopt a model whereby computational nodes and point-to-point interconnections map to the vertices and edges of a graph [Hayes 1976]. Each node consists of at least one slice; a node that is a flight computer spans three or four slices. In the case of X2000 a faulty node that partitions the bus is naturally modeled by deleting a vertex from the corresponding graph. Commensurate with X2000 documentation, we focus on fault tolerance in the *worst case*.³




Figure 2: Diagnosis and configuration for fault tolerance: architectures versus algorithms.

A *diagnosis architecture* is an assignment of pairwise tests among n nodes, and may be modeled as a directed graph whose vertices map as elements and whose arcs correspond to test relations. The outcome of each test is either “pass” or “fail.” The ensemble of these outcomes is known as a *syndrome*. The syndrome serves as input to a *diagnosis algorithm*, purpose of which is to accurately identify the faulty elements. We assess a diagnosis algorithm in terms of its correctness and efficiency. As to diagnosis architectures, the most prevalent figure of merit is the *test redundancy*; that is, the average degree of a vertex in the underlying digraph.

As with diagnosis, configuration can be viewed in terms of architectures and algorithms. Often the target architecture is constrained by shape: a d -dimensional array or torus with extents prescribed [LaForge 1998], a d -dimensional hypercube [Armstrong and Gray 1981], or a j -ary balanced tree of height d [Chen and Upadhyaya 1993]. For example, [Hayes 1976] proposes and analyzes graph architectures whose target architectures include one-dimensional arrays, simple cycles, and balanced trees. For X2000, by contrast, we simply desire that all of the good nodes be connected by some spanning tree. In this case the target architecture is known as a *quorum*. We may place additional requirements on the quorum, such as graph diameter or graph radius.⁴ In our case the 1394 bus specification prescribes that the quorum must contain a tree whose diameter (*i.e.*, maximum number of network hops) is at most 16; in the interest of performance, moreover, we seek to minimize the maximum number of network hops in the tree configured [P1394 1995]. Furthermore, we desire that all of the faulty nodes be dissociated from the quorum.

The *connectivity* of a graph G is the minimum number of vertices whose removal from G results in a disconnected graph or a lone vertex.⁵ To tolerate f partitioning faults, therefore, we seek architectures whose corresponding graph is $(f+1)$ -connected. Since our primary cost function is the number of point-to-point

3. The worst-case graph model is the simplest that suits our purpose, and is to be contrasted with more general multi-hypergraph models, or those which probabilistically treat the distribution or behavior of faulty nodes, or the success of diagnosis or configuration [LaForge et al 1994], [LaForge 1994].

4. The *diameter* and *radius* of a graph are its maximum *resp.* minimum eccentricities. A vertex's *eccentricity* is the maximum distance to some other vertex. The (graph) *distance* between two vertices is the length of the shortest path connecting them. Depending on the graph, the diameter ranges between the radius and twice the radius ([Chartrand and Lesniak 1986], Thm 2.4).

5. Our definition and use of vertex connectivity is to be distinguished from the edge connectivity; the latter equals the minimum number of edges whose removal results in a disconnected graph or a lone vertex.



interconnections, we furthermore focus our attention on $(f+1)$ -connected graphs with minimum number of edges.⁶ A lower bound on this number is readily seen by noting that the connectivity of a graph is at most the minimum degree of a vertex in the graph – that is, the minimum number of edges impinging on any vertex.⁷ In consequence, the degree of every vertex in an $(f+1)$ -connected graph is at least $f+1$. If we sum the degrees of all the vertices then we have counted every edge twice. The number of edges in any $(f+1)$ -connected n -vertex graph is therefore at least $\lceil n(f+1)/2 \rceil$. For any integers $n > f > 0$, moreover, [Hayes 1976] achieves this bound with constructions from which we can configure a one-dimensional array. These constructions are chordal graphs of order n and size $\lceil n(f+1)/2 \rceil$ from which we can remove i vertices, $0 \leq i \leq f$, and still have an $n-i$ vertices connected together as a path P_{n-i} .⁸ Unfortunately, the diameter of P_{n-i} equals $n-i-1$ and is maximum over all quorums. Thus, chordal constructions that achieve a P_{n-i} depart from our objective.⁹ Moreover, although [Hayes 1976] and [Kwan and Toida 1981] consider graph architectures from which we can configure trees, the analyses do not apply in the case of X2000 avionics. Table 6 indexes our notation for a development that *does* model X2000 avionics.

3.1 Quorums from Trees, Cycles, and Cliques

Suppose that an n -vertex graph G is $(f+1)$ -connected and, for $0 \leq i \leq f$, denote by H an arbitrary quorum induced by deleting i vertices of G . A graph T of order n is a tree if and only if T is connected and cycle-free; equivalently, T is connected and has minimum size $n-1$ ([Chartrand and Lesniak 1986], Chapter 3). T is said to *span* H if T and H have the same vertices and every edge of T is an edge of H . For our purposes it will be more convenient to formulate the problem in terms of graph *radius* than in terms of diameter.⁴ This is largely a consequence of Theorem 1, Corollary 1.1, and Theorems 2 and 36, which free us from having to distinguish the radius of the induced quorum H from the radius of a tree spanning H . In the case of partitioning faults, our candidates for configuration architectures are members of the set $\mathcal{G}_{n,f,k}^+$ of minimum size $(f+1)$ -connected graphs of order n whose quorums, induced by deletion of up to f vertices, have radii at most k . For given n and f , we naturally wish to assure that k is the exact minimum, in which case we write $\mathcal{G}_{n,f}$, perhaps with an extra subscript k . We denote the corresponding radius by $\rho(n, f)$. Although the general solution to this problem appears to be unknown,¹⁰ we can enumerate $\mathcal{G}_{n,0,k=2}$, $\mathcal{G}_{n,1,k=\lfloor n/2 \rfloor}$, and $\mathcal{G}_{n,n-2,k=2}$; that is, $\rho(n, 0) = 2$, $\rho(n, 1) = \lfloor n/2 \rfloor$, and $\rho(n, n-2) = 1$. For other values of f , we provide upper and lower bounds on $\rho(n, f)$, and give sets $\mathcal{G}_{n,f,k}^+$ whose induced quorums have radii logarithmic in n .




Figure 3: While a spanning tree that minimizes diameter also minimizes radius, the converse is not true.

6. Minimizing this cost is equivalent to minimizing the average degree of a vertex, and is therefore analogous to the objective of minimum test redundancy in the case of diagnosis.

7. [Chartrand and Lesniak 1986], Theorem 5.1: vertex connectivity \leq edge connectivity \leq minimum degree.

8. The *size* e and *order* n of a graph are the number of edges *resp.* number of vertices it contains.

9. Section 3.6 establishes that the radius of chordal graph quorums exceeds that of secant graph quorums.

10. The closest body of work seems to be related to the function $\varphi(n, d_0, d, f)$, introduced by [Murty and Vijayan 1964]. Here j counts the minimum number of edges in an n -vertex graph with diameter at most d_0 , such that deletion of any f of the vertices induces a graph of diameter at most d . Even for this relatively well-studied problem, results are confined primarily to the cases $d \leq 4$, $f = 1$ or $d_0 = 2$ ([Bollabás, 1978], Chapter IV, Sections 2 and 3). Moreover, our formulation differs in that we *fix* the number of edges at $\lceil (f+1)n/2 \rceil$, and then ask for the minimum diameter or radius achievable in a *tree* that spans the induced quorum.



Symbol	Significance	Page(s)
$\lceil x \rceil; \lfloor x \rfloor$	Ceiling (least integer no less than x); floor (greatest integer no greater than x)	9, 9
$\langle u, v \rangle; \langle P \rangle$	Graph distance between vertices u and v ; length of path P	15, 51
$O(g(n)); \Omega(g(n))$	Set of functions no greater <i>resp.</i> no less than $c \cdot g(n)$, for $n > k$, constants c, k	58
$o(g(n)); \omega(g(n))$	Set of functions $h(n)$ such that $\lim_{n \rightarrow \infty} h/g = 0$ <i>resp.</i> $\lim_{n \rightarrow \infty} g/h = 0$	55
$\Theta(g(n))$	Intersection of $O(g(n))$ and $\Omega(g(n))$	58
$B_j(d,i); B_j(d,i,m); B_j^C(d,i)$	Number of vertices at graph distance i from any vertex in K_j^d ; in $K_{m,j}^d$; in C_j^d	17, 32, 45
$C_n; C_{n,f+1}; C(m,f); C_j^d$	n -vertex cycle; n -vertex $(f+1)$ -regular chordal graph; (m,f) -vertex secant graph; d -dimensional j -ary C-cube	12, 42, 42, 44
$\Delta(n, f)$	Maximum diameter among quorums induced by f or fewer faults	57
$e, e_K(d,j); e_{K(d,j,n)}; e_C(d,j)$	Size (number of edges) of a graph; ⁸ of a K_j^d ; of a $K_j^d(n)$; of a C_j^d	15, 24, 45
f, f_{frac}	Number, fraction f/n of faulty elements (deleted vertices) that can be tolerated	3, 57
G	Graph, often one that represents the configuration architecture	8
$\mathfrak{G}_{n,f,k}^+$	Set of minimum size $(f+1)$ -connected graphs of order n whose quorums, induced by deletion of up to f vertices, have radii at most k	9
$\mathfrak{G}_{n,f}, \mathfrak{G}_{n,f,k}$	Set $\mathfrak{G}_{n,f,k}^+$ that minimizes the maximum radius k	9
$H; T$	Quorum induced by deleting vertices from G ; tree, often one that spans H	9
$K_n = K_j^1; K_j^d$	n -vertex clique; d -dimensional j -ary K-cube	12, 15
$K_j^d(n); K_{m,j}^d$	d -dimensional j -ary K-cube-connected cycle on n <i>resp.</i> $m \cdot j^d$ vertices	24, 28
$n; n_K(d,j); n_C(d,j)$	Order (number of vertices) of a graph; of a K_j^d ; of a C_j^d	3, 15, 45
$\rho(n, f)$	Maximum radius among quorums induced by f or fewer faults	9
$P_n; S_n$	n -vertex path; n -vertex star	10, 11
$V_j^C(d,i)$	Number of vertices graph distance at most i from any vertex in C_j^d	45

Table 6: Notation.

Refer to Figure 3. A vertex is *central* if its eccentricity equals the graph radius. If k is odd then a path P_k of order k has one central vertex (at the midpoint of P_k , graph distance $(k-1)/2$ from either endpoint). If k is even then P_k has *two* central vertices (at midpoints whose graph distance is $\lceil (k-1)/2 \rceil$ from one end and $\lfloor (k-1)/2 \rfloor$ from the other). More generally:

Theorem 1. In any tree T having longest path P_k of length $k-1$, there is a unique central vertex u at the midpoint of P_k , distance $\lceil k/2 \rceil$ from one end of P_k , if and only if P_k is unique and k is odd, or if there is a second maximum length path Q_k . In the latter case, u lies at the intersection of all maximum length paths of T . If P_k is unique and k is even then T has two central vertices, each of which lies distance $\lceil k/2 \rceil$ from an endpoint of P_k .

Proof. Suppose that a central vertex u does not lie on arbitrary path P_k of maximum length $k-1$. In any tree



there is a unique path between any two vertices ([Chartrand and Lesniak 1986], Thm 3.4). The path P_j from u to the its intersection with P_k has nonzero length j . The distance from u to the farthest endpoint of P_k is therefore strictly greater than $\lceil k/2 \rceil$. The sum of the distances from a midpoint v of P_k to two farthest leaves in the tree is at most $k-1$ (if not then P_k is not a maximum length path). Since the eccentricity of u is strictly greater than the eccentricity of v , vertex u cannot be central. Thus any central vertex u of the tree lies on P_k . If u lies on P_k , but is not graph distance $\lceil k/2 \rceil$ from one end and $\lfloor k/2 \rfloor$ from the other end of P_k , moreover, then there is at least one vertex v (and another w , if $k-1$ is odd) whose eccentricity is less than that of u . Therefore, the only candidates for central vertices of the tree are midpoints of P_k . If P_k is unique and $k-1$ is even then the central vertex is the unique midpoint of P_k . If P_k is unique and $k-1$ is odd then there are two midpoints of P_k and these are the central vertices of the tree. If there is more than one maximum length path then any two of these paths P_k, Q_k intersect at v , midpoint of both P_k and Q_k (if not then P_k and Q_k are not maximum length). Since any central vertex lies on both P_k and Q_k , u is the unique central vertex and lies at the intersection of all maximum length paths. \square

Corollary 1.1. The diameter of a tree is either twice its radius, or one less than twice its radius.

By Theorem 1, choosing a graph G whose every induced quorum has a spanning tree T with diameter at most k is equivalent to choosing a graph whose every induced quorum has radius at most $\lceil k/2 \rceil$. In particular, if G minimizes the diameter of T then G also minimizes the radius of T .¹¹ As Figure 3 illustrates, the converse is *not* true. However, Corollary 1.1 reveals how the converse is “almost” true: choosing a structure G whose every induced quorum H has a spanning tree with minimum radius either minimizes the diameter, or comes within one of a minimum diameter spanning tree. In essence, that is, we do as well to minimize radius as to minimize the diameter of a spanning tree. In terms of 1394 specifications, it suffices to ensure that the radius of the spanning tree does not exceed 8. But why is it more *convenient* to formulate the problem in terms of radius? A principal reason is the following.

Theorem 2. ([Chartrand and Lesniak 1986], Thm. 3.5; [Ore 1962], p. 102) For every vertex u of a connected graph H , there exists a spanning tree T of H that is distance-preserving from u .

By controlling the structure of G , we should be able to influence and profit from the structure of H (hence trees that span H). We devote the remainder of this section, as well as Sections 3.2 through 3.8, to characterizing this structure. For the sake of completeness, we begin with G which are 1-connected; that is, graph architectures that cannot tolerate any partitioning faults. Refer to Figure 4. An n -vertex *star* S_n is a tree having $n-1$ leaves, all connected to a single vertex. The following theorem exemplifies perhaps the simplest use of *expectation* in graph theory: in any set of real numbers whose arithmetic average is x , at least one element of the set is no less than x and at least one element of the set is no greater than x .

Theorem 3. For integers $n > f = 0$, the star S_n is the unique $(f+1)$ -connected unlabeled graph of order n and minimum size $n-1$ having minimum radius 1 and minimum diameter 2.

Proof. For $n \leq 2$ the theorem holds by inspection. For $n \geq 3$, consider the set of 1-connected graphs of minimum size $n-1$: that is, the trees of order n ([Chartrand and Lesniak 1986], Chapter 3). The average degree of any vertex in any such tree is $2 - 2/n > 1$. Therefore, at least one vertex u has two neighbors v and w . Since the path between v and w is unique ([Chartrand and Lesniak 1986], Theorem 3.4), the distance between u and w equals two. Thus, any n -vertex tree has radius and diameter at least 1 *resp.* 2. By inspection, S_n matches this bound. Further, any n -vertex tree other than S_n has more than one interior vertex, and therefore has radius and diameter strictly greater than 1 *resp.* 2. Hence S_n is the unique 1-connected unlabeled graph of order n and minimum size $n-1$ having minimum radius 1 and minimum diameter 2. \square

11. If not then there is a spanning tree Q whose radius is at most $\lceil t/2 \rceil - 1$. But by Corollary 1.1, $\text{diam } Q \leq 2\lceil t/2 \rceil - 2 < \lceil t/2 \rceil + \lfloor t/2 \rfloor = \text{diam } T$, contradicting the hypothesis that T has minimum diameter.




Figure 4: The n -vertex star S_n is the unique element of $\mathcal{G}_{n,0}$. $e(n) = n-1$. $\rho(n,0) = 2$.

Note that there is a gap between the number $n-1$ of edges in a tree (in particular, in S_n) and the lower bound $\lceil n(f+1)/2 \rceil = \lceil n/2 \rceil$ derived on page 9. It is curious that, by simply allowing one more edge, we can obtain a 2-connected graph. Refer to Figure 5. As with 1-connected graphs having minimum radius and diameter, the analogous unlabeled 2-connected graph, an n -vertex cycle C_n , is unique. As is the case for all $f > 0$, moreover, the edge count $\lceil n(f+1)/2 \rceil = n/2$ of C_n matches exactly the lower bound derived on page 9. By contrast to S_n , however, the uniqueness of C_n is due to constraints on connectivity, and not on radius or diameter. These observations are formalized by the following.

Theorem 4. For integers $n-1 > f = 1$, the cycle C_n is the unique $(f+1)$ -connected unlabeled graph of order n and minimum size n . C_n has (minimum) radius and diameter $\lfloor n/2 \rfloor$. If u is any vertex of C_n then the quorum $C_n \setminus u$, induced by deleting u , has radius $\lfloor n/2 \rfloor - 1$ and diameter $n-2$.

Proof. C_n is 2-connected by inspection. Therefore, any minimum size 2-connected n -vertex graph G has exactly n edges. The degree of every vertex of G must be exactly two, else some vertex has degree less than two and G cannot be 2-connected.⁷ That is, any vertex u in G has two neighbors, say v and w . By the results of Menger and Whitney, there are at least (exactly, in this case) two paths between u and v , and these paths are disjoint except for their endpoints ([Chartrand and Lesniak 1986], Theorems 5.10 and 5.11). One of these paths $P_{u,v}$ is just the edge (u,v) . The other path $P_{u,w,v}$ traverses some number $i > 2$ of vertices of G , including u , w , and v . Therefore, $P_{u,v} \cup P_{u,w,v}$ is a cycle, and each vertex in $P_{u,v} \cup P_{u,w,v}$ has degree exactly two. Suppose that some vertex z in G is not a member of the cycle $P_{u,v} \cup P_{u,w,v}$. Since G is connected, there must be a path from z to some vertex q of $P_{u,v} \cup P_{u,w,v}$. Let r be the last vertex along this path from z to q that is not a member of the cycle $P_{u,v} \cup P_{u,w,v}$. The edge (q,r) is therefore not a part of the cycle $P_{u,v} \cup P_{u,w,v}$. But this means that the degree of q is at least three, a contradiction. Therefore the cycle $P_{u,v} \cup P_{u,w,v}$ includes all the vertices of G . That is, G is identically C_n , the unique 2-connected unlabeled graph of order n and minimum size n . The results for radius and diameter follow by inspection. □




Figure 5: The n -vertex cycle C_n is the unique element of $\mathcal{G}_{n,1}$. $e(n) = n$. $\rho(n,1) = \lfloor n/2 \rfloor$.

For either minimum radius or diameter, Theorems 3 and 4 record the exact membership of the set $\mathcal{G}_{n,f}$ when $f = 0$ or $f = 1$; that is, when the number f of faults tolerated is as far from n as possible. It is noteworthy to remark on the membership of $\mathcal{G}_{n,f}$ when f is close to n . A *clique*, or *complete graph*, is a graph K_n with every possible edge in place. A clique of order n has size $n(n-1)/2$; that is, each vertex has degree $n-1$.

Theorem 5. For integers $f+1 = n-1$, K_n is the unique $(f+1)$ -connected unlabeled graph of order n and (minimum) size $n(n-1)/2$. For $0 \leq i \leq f$, if U_i is any set of i vertices of K_n then the quorum $K_n \setminus U_i$, induced by deleting the vertices of U_i , has radius and diameter 1.






Figure 6: The n -vertex clique K_n is the unique element of $\mathcal{G}_{n,n-2} \cdot e(n) = n$. $\rho(n, n-2) = 1$.

Proof. The degree of each vertex in an $(n-1)$ -connected graph of order n is at least $n-1$.⁷ Further, $n-1$ is the maximum degree of any vertex in a (simple) graph of order n . K_n is the unique such unlabeled graph. The sum $n(n-1)$ of the degrees of the vertices counts each edge twice; hence K_n has size $n(n-1)/2$. It remains to establish the connectivity, radius, and diameter. We proceed by induction on n . For a basis the theorem holds at $n = 2$ by inspection. Assume the theorem holds on $2, \dots, (n-1)$ vertices and consider arbitrary set U_i of i vertices to be deleted from K_n , $n > 2$, $0 \leq i \leq f$. Since the distance between any two vertices equals one, the theorem holds whenever U_i is empty. Otherwise let $u_1 \in U_i$ be the first vertex deleted from K_n . Since the induced quorum is a K_{n-1} , we recursively apply the theorem with $n-1, f-1$, and the set $U_i \setminus u_1$. \square

Fault tolerance f	$\mathcal{G}_{n,f,k}^+$: $(f+1)$ -connected graphs of minimum size $\lceil n(f+1)/2 \rceil$, induced quorums have radii at most k	Radius of quorum and of tree spanning quorum, as a function of the number $i \leq f$ of vertices deleted	Diameter of tree spanning quorum, as a function of the number $i \leq f$ of vertices deleted	References
0	Best possible $\mathcal{G}_{n,0,1}$ uniquely the set of n -vertex stars S_n	1	2	Figure 4 Theorem 3
1	Best possible $\mathcal{G}_{n,1,\lfloor n/2 \rfloor}$ uniquely the set of n -vertex cycles C_n	$\lfloor n/2 \rfloor$ if $i = 0$ $\lfloor n/2 \rfloor - 1$ otherwise	$n - 1$ if $i = 0$ $n - 2$ otherwise	Figure 5 Theorem 4
$n-2$	Best possible $\mathcal{G}_{n,n-2,1}$ uniquely the set of n -vertex cliques K_n	1	1 if $i = f$ 2 otherwise	Figure 6 Theorem 5
$n-1$	Best possible $\mathcal{G}_{n,n-1,1}$ uniquely the set of n -vertex cliques K_n	0 if $i = f$ 1 otherwise	0 if $i = f$ 1 if $i = f-1$ 2 otherwise	Discussion following Theorem 5

Table 7: Characteristics of quorums at either end of the range of the fault tolerance $f < n$, $n \geq 3$.

Theorems 3, 4, and 5 establish the exact memberships, as well as the respective exact values of $\rho(n, f)$, at the endpoints of the range of f . Further, if we want to tolerate up to $n-1$ faults then we must tolerate $n-2$ faults. Therefore, the set $\mathcal{G}_{n,n-1}$ is identically $\mathcal{G}_{n,n-2}$. Table 7 summarizes these results. Both the radius and diameter of an induced quorum H tend to change as we delete more vertices from G . While the radius of H is the same as that of a minimum-radius spanning tree T of H , the diameter of H is in general less than the diameter of T . By Corollary 1.1 and Theorem 2, the minimum diameter of a tree that spans a quorum is at least one less than twice a lower bound on the radius of the quorum, and at most twice an upper bound on the quorum radius. Combining this observation with Theorem 5, for example, we conclude that the diameter of a tree spanning a quorum induced by deleting a single vertex from a cycle is between $2\lfloor n/2 \rfloor - 3$ and $2\lfloor n/2 \rfloor - 2$. In some cases (such as this one) we can tighten these bounds even further. The diameters of the fourth column in Table 7, for example, are exact. In other cases we will establish bounds on the radius and diameter of quorums, and of trees that span quorums.





3.2 General Lower Bound on Quorum Radius

Sections 3.3 through 3.7 characterize a taxonomy of graphs $\mathcal{G}_{n,f,k}^+$ covering K-cubes, K-cube-connected cycles, K-cube-connected edges, chordal graphs, and C-cubes. Analysis of each set in this taxonomy gives *constructive* upper bounds on $\rho(n, f)$. In each case we furnish lower bounds on the radius of induced quorum. Short of expanding our taxonomy to include every possible graph, however, there remains the question whether other constructions have quorums with smaller radii. For this reason we seek a lower bound on the problem; *i.e.*, one that is *independent* of our choice of the embedding graph.

In what follows we take the *principal value* of $i \bmod j$ as the least nonnegative integer h such that, for some integer q , $i = qj + h$. We use the equal sign to denote evaluation of the congruence to its principal value. For example, $-6 \equiv 16 \bmod 11$, while $5 = 16 \bmod 11$. In any *rooted tree*, we define the *level* of vertex as its distance from the root; the *height* of the tree is its maximum level. We will also make use of the formula for summing terms j through m of a geometric series with common ratio $x \neq 1$ and constant coefficient a ([Thomas 1969], p. 623):

$$\frac{a(x^{m+1} - x^j)}{x - 1} = \sum_{i=j}^m ax^i \quad (1)$$

Theorem 6. (General lower bound on radius). For $1 < f < n-2$: $\rho(n, f) \geq \left\lceil \log_f \left[\frac{n(f-1)+3}{f+2} \right] \right\rceil$

Proof. Let H be any quorum induced by deleting i vertices of G , where G is a minimum size $(f+1)$ -connected graph of order n , and $0 \leq i \leq f$. Let T be a spanning tree of H with radius the same as that of H . Since G is of minimum size, $n-1$ of the vertices of G have degree $f+1$; one vertex of G has degree $(f+1) + [n(f+1) \bmod 2]$. These values bound as well the degree of any vertex in T . Let u be a central vertex of T (by Theorem 1, there are at most two). The radius of T may be viewed as the *height* of T when rooted at u . The height of any such tree is minimized when the number of children of every interior vertex is maximized. Therefore, the height of T is at least the height h of such a tree T' on $n-i$ vertices. All but at most two of the vertices of T' have at most f children. Since the root of T' has no parent, it may have as many as $f+1$ children. Further, if n and $f+1$ are both odd then G contains an “extra” edge, which may add one to the number of children spawned by some interior vertex v . Denote by j the level of v in T' . The total number of vertices in T' is maximized when T' is complete; that is, the number of vertices in T' is at most

$$1 + (f+1) \sum_{i=0}^{h-1} f^i + [n(f+1) \bmod 2] \sum_{i=0}^{h-j-1} f^i \quad (2)$$

For given value of f , expression (2) is maximized when the vertex v is at level $j = 0$; that is, when v is the root of T' . Applying formula (1), we see that the number of vertices in T' is at most

$$1 + (f+1 + [n(f+1) \bmod 2]) \frac{f^h - 1}{f - 1} \quad (3)$$

Since T' contains $n-i$ vertices, it follows that $n-i$ is no greater than (3). Recalling that the height (and thus the radius) of T is at least the height h of T' , we conclude

$$\rho(n, f) \geq h \geq \left\lceil \log_f \left[\frac{(n-i-1)(f-1) + f + 1 + [n(f+1) \bmod 2]}{f + 1 + [n(f+1) \bmod 2]} \right] \right\rceil \quad (4)$$

Verify that (4) is minimized when $1 = n(f+1) \bmod 2$. Since (4) must hold for *every* integer i between 0 and f inclusive, it must in particular hold when $i = 0$; that is, when the righthand side is maximized with respect to i . This gives the result of the theorem. \square





Though independently derived for this work, the lower bound of Theorem 6 is reminiscent of the development of Bollobás and Harary for graphs of order n , diameter $2m$, connectivity $f+1$, and size at most $n(f+2)/2$. While the latter is quite close to the minimize size $\lceil n(f+1)/2 \rceil$ of graphs that are $(f+1)$ -connected, the corresponding constructions admit induced quorums having arbitrary diameter. At first blush, therefore, the development of Bollobás and Harary is not directly applicable to X2000 ([Bollobás, 1978], p 182). Theorem 6 also generalizes the bound attributed to Moore (circa 1958), that relates diameter, maximum degree, and number of vertices, when no vertices are deleted [Sampels 1997]. This said, let us proceed to expand our taxonomy of graphs and their induced quorums.

3.3 Quorums from K-cubes

A d -dimensional Gray-coded j -ary K-cube K_j^d is recursively constructed as follows.¹² K_j^0 is a lone vertex labeled with the null string. For K_j^d we i) make j copies of K_j^{d-1} ; ii) join with an edge vertices u and v (from different copies of K_j^{d-1}) if and only if u and v have with identical labels; iii) prepend i to the label of each vertex of the i th copy of K_j^{d-1} . Note that K_j^1 is just the clique K_j whose vertices have been labeled from 0 to $j-1$. Figure 7 illustrates binary and ternary K-cubes in 3 *resp.* 2 dimensions.




Figure 7: Gray-code labeling of a three-dimensional K_2 -cube and a two-dimensional K_3 -cube.

Since our constructions for members of $\mathcal{G}_{n,f,k}^+$ are based on d -dimensional j -ary K-cubes, it pays to know their salient properties. By step (i) above, K_j^d contains j copies of K_j^{d-1} ; therefore the order $n_K(d, j)$ of K_j^d equals $j \cdot n_K(d-1, j)$. Subject to the initial condition $n_K(0, j) = 1$, verify that the unique solution of this recurrence relation is

$$n_K(d, j) = j^d \tag{5}$$

By step (ii) above, the degree of a vertex in K_j^d equals its degree in K_j^{d-1} plus $j-1$, the number of edges that connect it to vertices with the same labels in the other copies of K_j^{d-1} . Subject to the initial condition of zero edges in K_j^0 , the degree of each vertex in K_j^d is therefore

$$d(j-1) \tag{6}$$

Summing (6) over all j^d vertices counts every edge twice. Hence the total number $e_K(d, j)$ of edges in K_j^d is

$$e_K(d, j) = \frac{1}{2} \cdot d(j-1) \cdot j^d \tag{7}$$

Since graph *distance* is a theme underlying both radius and diameter, we seek to characterize the distances among vertices in K_j^d . We abbreviate the distance between vertices u and v as $\langle u, v \rangle$.¹³ The Gray-code labeling prescribed by step (iii) above does not change the distances among vertices of K_j^d , but it *does* help to elucidate the distances.

12. Unlike C-cubes (Section 3.7), K-cubes are *cliqued based*; with notation based on that K_j for a j -vertex clique.

13. Over all vertices in any connected graph, the graph distance, defined in footnote 4, forms a *metric space*: i) $\langle u, v \rangle \geq 0$ (nonnegative definite); ii) $\langle u, v \rangle = \langle v, u \rangle$ (symmetry); iii) $\langle u, v \rangle = 0$ if and only if $u = v$ (identity); iv) $\langle u, w \rangle \leq \langle u, v \rangle + \langle v, w \rangle$ (triangle inequality). The latter can be used to prove that the diameter is at most twice the radius ([Chartrand and Lesniak 1986], Chapter 2).





Theorem 7. If u and v are vertices of K_j^d , Gray-code labeled according to steps (i) – (iii) on page 15, then $\langle u, v \rangle$ equals the number of digits where the respective labels for u and v are different.

Proof. By induction on d . For a basis, the theorem holds at $d = 0$ by inspection. Assume that the theorem holds in $0, \dots (d-1)$ dimensions, and regard arbitrary vertices u and v in K_j^d . By steps (i) and (ii) on page 15, u and v are contained in copies K' and K'' of K_j^{d-1} . If $K' \neq K''$ then, by step (iii), the labels of $u = u'$ and $v = v''$ differ in the high order digit. Let u'' be the vertex in K'' having the same label as u' , except for the high order digit. By induction, there is a shortest path P'' from u'' to v'' , strictly contained in K'' , whose length equals the number of digits where the respective labels for u'' and v'' are different. By steps (ii) and (iii) on page 15, u' adjoins u'' . Therefore, there is a path $P'' + (u', u'')$ from u' to v'' whose length equals the number of digits where the respective labels for u' and v'' are different. Suppose that some other path Q between u' and v'' has length strictly less than that of $P'' + (u', u'')$. Traversing Q from u' to v'' , there is a vertex w' where Q first leaves K' and a vertex z'' where Q last enters K'' . By induction, the length of the portion of Q from u' to w' is at least the number of digits where the labels of u' to w' differ. Similarly, the length of the portion of Q from z'' to v'' is at least the number of digits where the labels of z'' to v'' differ. Moreover, the portion of Q from w' to z'' is at least one edge long. If the labels on u' and v'' differ in the h^{th} digit then, as we traverse from u' to v'' along Q , the value of digit h changes at least once. If the value of the h^{th} digit changes more than once then Q is strictly longer than $P'' + (u', u'')$. Thus, as we traverse from u' to v'' along Q , the digits where the labels on u' and v'' are different change only once. But this means that Q is at least as long as $P'' + (u', u'')$, contradicting the assumption that Q is shorter than $P'' + (u', u'')$. When $u = u'$ and $v = v''$ are contained in different copies K' resp. K'' of K_j^{d-1} , therefore, $\langle u', v'' \rangle$ equals the number of digits where the respective labels for u' and v'' are different, and $P'' + (u', u'')$ is one such shortest path.

If $K' = K''$ then, by induction, there exists a path P' , strictly contained in K' , whose length is equal to the number of digits where the respective labels are different; furthermore, P' is a shortest path from $u = u'$ to $v = v'$ that is strictly contained in K' . If there is path Q between u' and v' whose length strictly less than that of P' , then Q must necessarily exit and re-enter K' . Traversing Q from u' to v' , there is a vertex w' where Q first leaves K' and a vertex z' where Q last enters K' . By induction, the length of the portion of Q from u' to w' is at least the number of digits where the labels of u' to w' differ. Similarly, the length of the portion of Q from z' to v' is at least the number of digits where the labels of z' to v' differ. Moreover, the portion of Q from w' to z' is at least two edges long. If the labels on u' and v' differ in the h^{th} digit then, as we traverse from u' to v' along Q , the value of digit h changes at least once. If the value of the h^{th} digit changes more than once then Q is strictly longer than P' . Thus, as we traverse from u' to v' along Q , the digits where the labels u' and v' are different change only once. Since the portion of Q outside of K' is at least two edges long, Q is strictly longer than P' , contrary to assumption. When $u = u'$ and $v = v'$ are contained in the same copy K' of K_j^{d-1} , therefore, $\langle u', v' \rangle$ equals the number of digits where the respective labels for u' and v' are different, and all such paths are strictly contained in K' . □

K_j^d is *vertex symmetric*; that is, the perspective of K_j^d is the same from every vertex.¹⁴ A valid argument about properties of K_j^d , with respect to some fixed vertex u , remains valid when u is replaced by any other

14. More precisely, a graph G is *vertex-symmetric* of the group $A(G)$ of graph automorphisms of G acts transitively on V ; i.e., for any $v, w \in V$, there is a graph automorphism $\alpha \in A(G)$ such that $\alpha(v) = w$ [Sampels 1997].



vertex. For example, how many vertices lie at distance i from any given vertex u in K_j^d ? The answer does not depend on our choice of u , and so, without loss of generality, we may assume that u lies at the origin, *i.e.*, all d digits of its label equal 0. Geometrically, this is equivalent to asking for the number $B_j(d,i)$ of vertices on the surface of a ball in K_j^d centered at u and having integer radius i .¹⁵ To answer this question, we invoke the structure revealed by the proof of Theorem 7. Suppose that $u = u'$ is in K' , one of j copies of K_j^{d-1} comprising K_j^d . The number of vertices distance i from u' equals the number of vertices in K' that are distance i from u' plus the number of vertices in the $j-1$ other copies of K_j^{d-1} that are distance i from u' . Focus on any one of these other copies $K'' \neq K'$. By Theorem 7, any vertex v'' in K'' that is distance i from u' can be reached by a shortest path $(u', u'') + P''$, where u'' is connected to u' and, except for the high order digit, is labeled the same as u' . Further P'' is a shortest path of length $i-1$, strictly in K'' , from u'' to v'' . Thus, in addition to the vertices in K' that are distance i from u' , the surface of the ball includes the vertices in K'' that are distance $i-1$ from u' . Since this is true for each $K_j^{d-1} \neq K'$, we have the recurrence

$$B_j(d,i) = B_j(d-1,i) + (j-1) \cdot B_j(d-1,i-1) \quad \text{with boundary conditions} \quad B_j(d,0) = 1, B_j(d,i>0) = 0 \quad (8)$$

↓ d → i	j = 2							j = 3							j = 4							
	0	1	2	3	4	5	6	0	1	2	3	4	5	6	0	1	2	3	4	5	6	
0	1							1							1							
1	1	1						1	2						1	3						
2	1	2	1					1	4	4					1	6	9					
3	1	3	3	1				1	6	12	8				1	9	27	27				
4	1	4	6	4	1			1	8	24	32	16			1	12	54	108	81			
5	1	5	10	10	5	1		1	10	40	80	80	32		1	15	90	270	405	243		
6	1	6	15	20	15	6	1	1	12	60	160	240	192	64	1	18	135	540	1215	1458	729	

Table 8: Number $B_j(d,i)$ of vertices at graph distance i from any other in a d -dimensional j -ary K-cube . The table may be verified or extended using any of the formulae (8) or (9), and for $j=2$ reduces to Pascal's triangle.

Refer to Table 8. When $j = 2$ the recurrence of (8) reduces to that for *Pascal's triangle*, and we have $B_2(d,i) = \binom{d}{i} = \frac{d!}{(d-i)!i!}$, the binomial coefficient “ d choose i ”. More generally, Theorem 7 allows us to solve (8) by combinatorial argument [Comtet 1974]. A *j-nomial coefficient* whose numerator equals $d!$ has j factors $q_0!, \dots, q_{j-1}!$ in the denominator. Absent the factorial, the values of q_h sum to $d = \sum_{0 \leq h \leq j-1} q_h$; in consequence, only $j-1$ of the factors need be explicated. Such is the custom for example, with $j = 2$, *i.e.*, the binomial coefficient $\binom{d}{i} = \frac{d!}{(d-i)!i!} = \frac{d!}{i!(d-i)!} = \binom{d}{d-i} = \binom{d}{d-i, i}$. For the general case, we let q_0 be the number of *digits* having value 0 in the label on any vertex v distance i from u . That is, q_0 is the number of digits where the label of v is the same as the label of u . For $0 < h < j$, q_h equals the number of digits of v whose value equals h . Therefore, the sum of all such q_h is just the number of digits where the label of v dif-

15. If $j = 2$ then the labels of K_2^d are bit strings, and we have a special case of the L_1 metric: the *Hamming distance*. In 1950 Richard Hamming introduced detection and correction codes that bear his name. These codes are based on K_2^d packings of balls of given Hamming radius ([Wakerly 1990], Section 2.14).





fers from the label of u . By Theorem 7, that is, $i = \sum_{1 \leq h \leq j-1} q_h$ and, for $0 < h < j$, the ordered sequence q_1, \dots, q_{j-1} , abbreviated as $\mathbf{q}^+_{i,j}$, determines the value $q_0 = d - i$. Note that $\mathbf{q}^+_{i,j}$ may be viewed as a vector whose components are nonnegative integers no greater than i . Denote by $\mathcal{Q}^+_{i,j}$ the set of all vectors $\mathbf{q}^+_{i,j}$. For any vector $\mathbf{q}^+_{i,j}$ the multinomial coefficient $\binom{d}{q_0, \dots, q_{j-1}} = \binom{d}{(d-i), \dots, q_{j-1}} = \frac{d!}{(d-i)!q_1! \dots q_{j-1}!} \stackrel{\text{def}}{=} \binom{d}{\mathbf{q}^+_{i,j}}$ counts the number of vertices at distance i from u . Over all such $\mathbf{q}^+_{i,j} \in \mathcal{Q}^+_{i,j}$, summing these multinomial coefficients gives

$$B_j(d, i) = \sum_{\mathbf{q}^+_{i,j} \in \mathcal{Q}^+_{i,j}} \binom{d}{\mathbf{q}^+_{i,j}} \quad (9)$$

For example, we use (9) to verify two entries of Table 8. $\mathcal{Q}^+_{2,3} = \{(0,2), (2,0), (1,1)\}$. $B_3(3,2)$ is therefore $\binom{3}{1,0,2} + \binom{3}{1,2,0} + \binom{3}{1,1,1} = 3 + 3 + 6 = 12$. $\mathcal{Q}^+_{2,4} = \{(0,0,2), (0,1,1), (1,0,1), (0,2,0), (1,1,0), (2,0,0)\}$. Thus, $B_4(3,2)$ is $\binom{3}{1,0,0,2} + \binom{3}{1,0,1,1} + \binom{3}{1,1,0,1} + \binom{3}{1,0,2,0} + \binom{3}{1,1,1,0} + \binom{3}{1,2,0,0} = 27$. Although (9) illuminates the underlying combinatorics, its convenience is limited by the difficulty of enumerating the membership of $\mathcal{Q}^+_{i,j}$. This subproblem appears to be even more difficult than computing the unordered partitions of an integer,¹⁶ and leads to seek an alternative formula for $B_j(d, i)$.

We can express $B_j(d, i)$ in closed form by noting that there are $\binom{d}{i}$ ways for vertices in a Gray-coded K_j^d to have i digits that are different from those of the origin $u_0 = u$. Using all the nonzero digit values radix j , we can form $(j-1)^i$ labels from any given set of i such digits. By the fundamental counting rule, that is, we obtain an expression that, by contrast to (9), is readily computed:

$$B_j(d, i) = (j-1)^i \binom{d}{i} \quad (10)$$

Verify by substitution that (10) satisfies the recurrence and boundary conditions of (8).

Equations (8), (9), and (10) are both relevant and illustrative of the power of Theorem 7. In addition to providing three ways of computing the number of vertices at distance i from any vertex in K_j^d , the theorem springs forth other analytic results. For example, (9) and (10) imply the double identity:

$$\sum_{i=0}^d \sum_{\mathbf{q}^+_{i,j} \in \mathcal{Q}^+_{i,j}} \binom{d}{\mathbf{q}^+_{i,j}} = \sum_{i=0}^d (j-1)^i \binom{d}{i} = j^d \quad (11)$$

Two paths are *interior-disjoint* if, apart from their endpoints, they do not intersect. In what follows we will once again invoke two important facts: a) the connectivity of a graph is the minimum number of interior-disjoint paths joining any two vertices;¹⁷ b) the connectivity of a graph is at most its minimum degree.¹⁸ With respect to (a), suppose that graph G has order at least $f+3$ and that between every pair of vertices in G there are $f+1$ interior-disjoint paths of length at most q . Any quorum induced by deleting $0 \leq i \leq f$ vertices of G , has diameter at most q . This observation motivates our formulations and proofs of Theorems 8 and 9.

16. An *inexact m -maximum unordered partition of the integer i* is an m -vector of nonnegative integers whose inner product with a vector containing the first m positive integers equals i . ([Anderson 1998], pp. 65-67).

17. Attributed to Menger and Whitney, this result is used in the proof of Theorem 4. It is recounted by Theorems 5.10 and 5.11 of [Chartrand and Lesniak 1986].

18. This is used on page 9 to bound from below the size of an $(f+1)$ -connected graph. Also see footnote 7.



Theorem 8. (Connectivity, upper bound on diameter.) If $j \geq 3$ then between vertices u and v in K_j^d there are $d(j-1)$ interior-disjoint paths, each of whose length is at most $d+1$. At least d of these paths have length at most d .

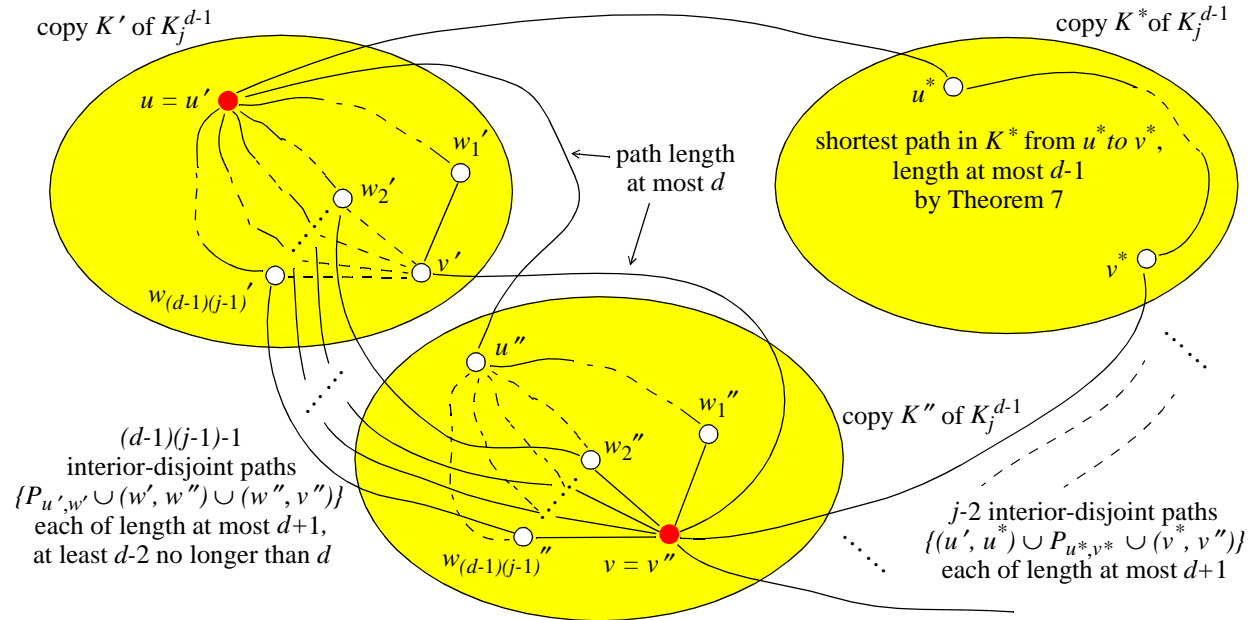


Figure 8: Illustration of the first part of the proof of Theorem 8.

Proof. By induction on d . For a basis, the theorem holds by inspection at $d = 0$ and $d = 1$. Assume that the theorem holds in $0, \dots, (d-1)$ dimensions, and regard arbitrary vertices u and v in K_j^d , $d > 1$, $j > 2$. By steps (i) and (ii) of the construction on page 15, u and v are contained in copies K' and K'' of K_j^{d-1} . Refer to Figure 8. If $K' \neq K''$ then in K' inductively form $(d-1)(j-1)$ interior-disjoint paths, of length at most d , from $u = u'$ to v' . Here v' is the vertex in K' whose label, except in the high order digit, matches that of $v = v''$. Pick the first $d-1$ of these paths to have length at most $d-1$. Since these paths are interior-disjoint, each such path passes through exactly one of the neighbors $w_1', \dots, w_{(d-1)(j-1)}'$ of v' . The path $P_{u',w_1'} \cup (w_1', v') \cup (v', v'')$ has length at most d . Exclusive of w_1' , there remain $(d-1)(j-1)-1$ neighbors $w_2', \dots, w_{(d-1)(j-1)}'$ of v' in K' . By steps (ii) and (iii) of the construction on page 15, each endpoint $w' \in \{w_2', \dots, w_{(d-1)(j-1)}'\}$ of the paths from u' to a neighbor w' of v' is adjoined to a counterpart $w'' \in \{w_2'', \dots, w_{(d-1)(j-1)}''\}$ in K'' . Except for the high order digit, the labels on w' and w'' are the same. Furthermore, the construction assures that each such w'' in K'' is a neighbor of v'' . Augmenting each of the $(d-1)(j-1)-1$ paths $P_{u',w'}$ with edges (w', w'') and (w'', v'') yields a set of $(d-1)(j-1)-1$ paths $\{P_{u',w'} \cup (w', w'') \cup (w'', v'')\}$, each of length at most $d+1$. The first $d-2$ of these paths have length at most d . These paths are interior-disjoint with each other and with $P_{u',w_1'} \cup (w_1', v') \cup (v', v'')$. For the remaining $j-1$ paths choose $(u', u'') \cup P_{u'',w_1''} \cup (w_1'', v'')$, a path of length at most d , along with the set $\{(u', u^*) \cup P_{u^*,v^*} \cup (v^*, v'')\}$. Here u^* and v^* have labels whose low order $d-1$ digits are identical to those on u' resp. v'' , but whose corresponding high order digit differs from that of either u' or v'' . P_{u^*,v^*} is a shortest path between u^* and v^* , strictly contained in $K^* \neq K' \neq K''$. By Theorem 7, P_{u^*,v^*} has length at most $d-1$; thus $(u', u^*) \cup P_{u^*,v^*} \cup (v^*, v'')$ has length at most $d+1$. The $j-1$ paths



$(u', u'') \cup P_{u'', w_1''} \cup (w_1'', v'')$, $\{(u', u^*) \cup P_{u^*, v^*} \cup (v^*, v'')\}$ are interior-disjoint with each other as well as with $P_{u', w_1'} \cup (w_1', v') \cup (v', v'')$ and $\{P_{u', w'} \cup (w_1', v') \cup (v', v'')\}$. The ensemble constitutes $d(j-1)$ interior-disjoint paths. Each of has length at most $d+1$, and at least d of them have length at most d .

Refer to Figure 9. If $K' = K''$ then in K' inductively form $(d-1)(j-1)$ interior-disjoint paths, of length at most d , from $u = u'$ to $v = v'$. At least $d-1$ of these paths have length at most $d-1$, and each of the others has length at most d . For the remaining $j-1$ paths choose the set $\{(u', u^+) \cup P_{u^+, v^+} \cup (v^+, v')\}$. Here u^+ and v^+ have labels whose low order $d-1$ digits are identical to those on u' and v' , but whose corresponding high order digit is different. P_{u^+, v^+} is a shortest path between u^+ and v^+ , strictly contained in $K^+ \neq K'$. By Theorem 7, P_{u^+, v^+} has length at most $d-1$; thus $(u', u^+) \cup P_{u^+, v^+} \cup (v^+, v')$ has length at most $d+1$. The $j-1$ paths $\{(u', u^+) \cup P_{u^+, v^+} \cup (v^+, v')\}$ are interior-disjoint with each other as well as with the $(d-1)(j-1)$ interior-disjoint paths in K' from $u = u'$ to $v = v'$. The ensemble constitutes $d(j-1)$ interior-disjoint paths, each of which has length at most $d+1$. Since $j > 2$, at least $(d-1)(j-1) \geq d$ of these paths have length at most d . \square

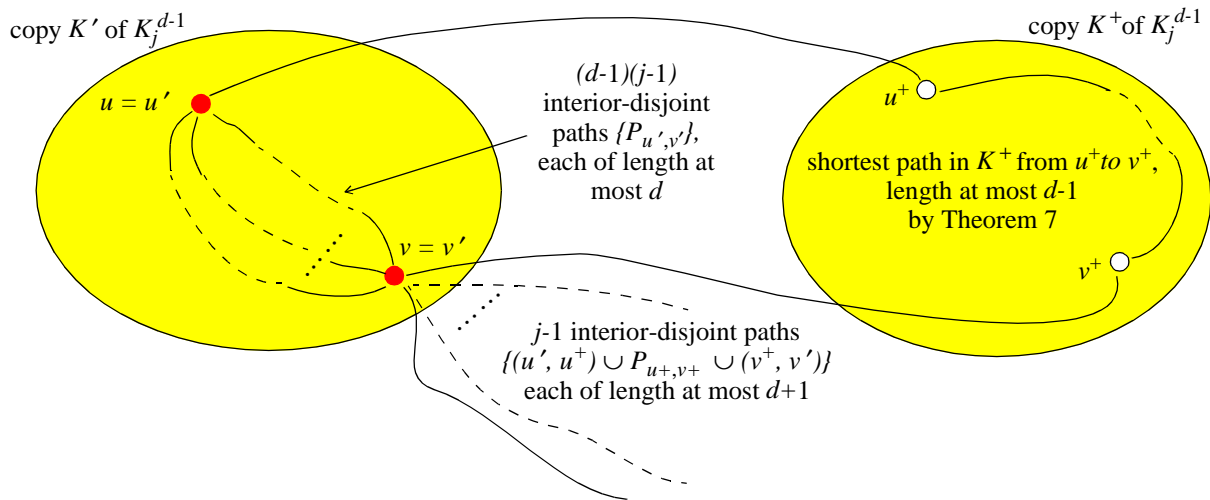


Figure 9: Illustration of the second part of the proof of Theorem 8.

For $j = 2$ the assertion of Theorem 8 is false; e.g., for $C_4 = K_2^2$ there are two paths between adjacent vertices, and one of these paths has length $3 > d = 2$. Regrettably, throughout this work we will often be relegated to treating the case $j = 2$ separately. This admittedly awkward situation is, apparently, a consequence of the natural order of things. For example, the radix 2 analog of Theorem 8 is:

Theorem 9. (Connectivity, upper bound on diameter.) Between two vertices in K_2^d , $d > 1$, there are d interior-disjoint paths of length at most $d+1$. At least $d-1$ of these paths have length at most d .

Proof. By induction on d . For a basis, the theorem holds by inspection at $d = 2$. Assume that the theorem holds in $2, \dots (d-1)$ dimensions, and regard arbitrary vertices u and v in K_2^d , $d > 2$. By steps (i) and (ii) of the construction on page 15, u and v are contained in copies K' and K'' of K_2^{d-1} . If $K' \neq K''$ then, in K' and analogous to Figure 8 (but without K^*), inductively form $d-1$ interior-disjoint paths, of length at most d , from $u = u'$ to v' . Here v' is the vertex in K' whose label, except in the high order bit, matches that of $v = v''$. Pick the first $d-2$ of these paths to have length at most $d-1$. Since these paths are interior-disjoint, each such path passes through exactly one of the neighbors $w_1', \dots, w_{(d-1)}'$ of v' . The path $P_{u', w_1'} \cup (w_1', v') \cup (v', v'')$ has length at most d . Exclusive of w_1' , there remain $d-2$ neighbors





$w_2', \dots, w_{(d-1)'}$ of v' in K' . By steps (ii) and (iii) of the construction on page 15, each endpoint $w' \in \{w_2', \dots, w_{(d-1)'}\}$ of the paths from u' to a neighbor w' of v' is adjoined to a counterpart $w'' \in \{w_2'', \dots, w_{(d-1)}''\}$ in K'' . Except for the high order bit, the labels on w' and w'' are the same. Furthermore, the construction assures that each such w'' in K'' is a neighbor of v'' . Augmenting each of the $d-2$ paths $P_{u',w'}$ with edges (w', w'') and (w'', v'') yields a set of $d-2$ paths $\{P_{u',w'} \cup (w', w'') \cup (w'', v'')\}$, each of length at most $d+1$. The first $d-2$ of these paths have length at most d . These paths are interior-disjoint with each other and with $P_{u',w_1'} \cup (w_1', v') \cup (v', v'')$. For the remaining path choose $(u', u'') \cup P_{u'',w_1''} \cup (w_1'', v'')$, a path of length at most d , interior-disjoint with $P_{u',w_1'} \cup (w_1', v') \cup (v', v'')$ and with $\{P_{u',w'} \cup (w_1', v') \cup (v', v'')\}$. The ensemble constitutes d interior-disjoint paths. Each of these paths has length at most $d+1$, and at least $d-1$ of them have length at most d .

If $K' = K''$ then, in K' and analogous to Figure 9, inductively form $d-1$ interior-disjoint paths, of length at most d , from $u = u'$ to $v = v'$. At least $d-2$ of these paths have length at most $d-1$, and the other has length at most d . For the remaining path choose $(u', u^+) \cup P_{u^+,v^+} \cup (v^+, v')$. Here u^+ and v^+ have labels whose low order $d-1$ bits are identical to those on u' and v' , but whose corresponding high order bit is different. P_{u^+,v^+} is a shortest path between u^+ and v^+ , strictly contained in $K^+ \neq K'$. By Theorem 7, P_{u^+,v^+} has length at most $d-1$; thus $(u', u^+) \cup P_{u^+,v^+} \cup (v^+, v')$ has length at most $d+1$. The path $(u', u^+) \cup P_{u^+,v^+} \cup (v^+, v')$ is interior-disjoint with the $d-1$ interior-disjoint paths in K' from $u = u'$ to $v = v'$. The ensemble constitutes d interior-disjoint paths, each of which has length at most $d+1$, with at least $d-1$ having length no greater than d . \square

Combining equations (5) and (7), remarks (a) and (b) on page 18, and Theorems 8 and 9:

Corollary 9.1. K_j^d has order j^d , connectivity $d(j-1)$ and minimum size $\frac{1}{2}d(j-1)j^d$.

Corollary 9.1 naturally leads us back to our central question: "What is the radius or diameter of quorums induced by deleting up to $d(j-1)-1$ vertices of K_j ?" We can now begin to address this issue. Theorems 8 and 9 give upper bounds on the diameter of any quorum H induced by deleting i vertices from K_j^d , $0 \leq i \leq f = d(j-1)-1$. If $j > 2$ and $0 \leq i \leq d-1$ then from any vertex u in H we can reach all other vertices of H by a path of length at most d . If $j > 2$ and $d \leq i \leq d(j-1)-1 = f$ then from any vertex u in H we can reach all the vertices of H by a path of length at most $d+1$. If $j = 2$ and $0 \leq i \leq d-2$ then from any vertex u we can reach all the vertices of H by a path of length at most d . If $j = 2$ and $i = d-1 = f$ then from any vertex u , we can reach all the vertices of H by a path of length at most $d+1$. In an extremal (*i.e.*, worst-case) sense, these bounds are best possible. To see this we construct a class of counterexamples.

Consider arbitrary vertex u' in copy K' of $K_j^{d-1} \subset K_j^d$. For $d > j = 2$, form H by deleting $d-1$ neighbors of u' , leaving undeleted one neighbor $u'' \in K''$ of u' . The label of u'' is same as that on u' , except for the high order bit. Let v' be the vertex in K' having label whose low order $d-1$ bits are all different from those of u' . In H any path of shortest distance $\langle u', v' \rangle$ necessarily enters K'' via the edge (u', u'') , follows a path to some vertex $z'' \in K''$, re-enters K' via edge (z'', z') , and follows a path from z' to v' . The length of this path is at least $\langle u', u'' \rangle + \langle u'', z'' \rangle + \langle z'', z' \rangle + \langle z', v' \rangle \geq 2 + \langle u', v' \rangle = d+1$. By equation (5), this construction holds for each of the 2^d values that u can take on. Of the $\binom{2^d}{d-1}$ quorums formed by deletion of $i = d-1$ vertices from a Gray-coded d -dimensional binary K-cube, that is, at least 2^d have diameter $d+1$. This class of counterexamples refutes an "almost correct" claim of [Armstrong and Gray 1981]:

Between any two vertices in K_2^d , $d > 1$, there are d interior-disjoint paths of length at most d .





The preceding claim is corrected by Theorem 9, and generalized by Theorem 8.

Continuing the counterexamples, for $j > 2$ let u be the vertex whose label is all zeros. H_j^1 is just K_j^1 with u deleted. For $d = 2$ delete from K_j^2 the vertex whose label is $[0, j-1]$, along with the vertex whose label is $[j-1, 0]$. Any path from $[0,0]$ to $[j-1, j-1]$ cannot go through the deleted vertices. The digits of the labels in any such path change at least three times. By Theorem 7, any path from $[0,0]$ to $[j-1, j-1]$ in H_j^2 must therefore have length at least $3 = d+1$. For $d > 2$ recursively form H_j^d by replacing K_j^1 , the zeroth copy of $K_j^{d-1} \subset K_j^d$, with H_j^{d-1} ; connect vertex h of H_j^{d-1} to vertex h of every other copy of K_j^{d-1} ; prepend the label on each vertex of H_j^{d-1} with a zero. Complete the construction of H_j^d by deleting u'' , vertex 0 of the $(j-1)^{\text{st}}$ copy of $K_j^{d-1} = K''$. Consider any shortest path P from u' , vertex 0 in H_j^{d-1} , to v'' , vertex $j-1$ in K'' . Let v' be the vertex of H_j^{d-1} whose label is the same as that on v'' , except that the high order digit of v' equals zero instead of $j-1$. If P runs through H_j^{d-1} and K'' but through no other copy of K_j^{d-1} then P cannot be of length d ; if it were then there would be a path in H_j^{d-1} of length $d-1$ from u' to v' , a contradiction. If P runs through some copy of K_j^{d-1} other than K'' then, by Theorem 7, P has length at least $d+1$. Therefore P has length at least $d+1$, and this is the diameter of H_j^d . By changing the label for u' , there are j ways to construct H_j^1 . In higher dimensions d , there are $\binom{j}{1, 1, j-2} = j(j-1)$ ways to choose the placement of H_j^{d-1} and K'' . Of the $\binom{j^d}{d}$ quorums formed by deletion of $i = d$ vertices from a Gray-coded d -dimensional j -ary K-cube, that is, at least $j^{d(j-1)^{d-1}}$ have diameter $d+1$.

Theorem 10. Let H be any quorum induced by deleting i vertices from K_j^d , $0 \leq i \leq f = d(j-1)-1$. The diameter of H is at least d .

Proof. Vertices u and v are *opposite* if they are distance d apart; *i.e.*, their labels differ in every position. By equation (10), any given vertex u has $(j-1)^d$ opposites; that is, there are $(j-1)^d$ opposite pairs that include u . Summing over all j^d vertices counts every pair of opposites twice, and the total number of opposite pairs equals $\frac{1}{2}j^d(j-1)^d$. Each vertex we delete from K_j^d removes at most $(j-1)^d$ opposite pairs. Therefore, there remains at least one opposite pair as long as

$$[d(j-1)-1](j-1)^d < \frac{1}{2}j^d(j-1)^d \quad (12)$$

(12) is satisfied if $d \leq 2^{d-2}$. For $d \geq 3$, that is, (12) is valid for all integers $j \geq 2$. For $d = 2$ verify by substitution that (12) is satisfied for all $j \geq 2$. At $d = 1$ the theorem follows by inspection. \square

Theorem 11. Let H be any quorum induced by deleting i vertices from K_j^d , $0 \leq i \leq f = d(j-1)-1$. If $d = 1$, $i = 0$, or $j \geq 3$ then the radius of H is at least d . If $j = 2$ and $i \geq 1$ then the radius of H is at least $d-1$.

Proof. By (10), any undeleted vertex u has at least one opposite as long as $d(j-1)-1 < (j-1)^d$ (13)

(13) is satisfied if $d \leq 2^{d-1}$. For $d \geq 4$, that is, (13) is valid for all integers $j \geq 3$. For $d = 3$ and $d = 2$ verify by substitution that (13) is satisfied for all $j \geq 3$. The theorem follows by inspection for $i = 0$, as well as for $d = 1, j \geq 2$. It remains to consider the case $i \geq 1, d \geq 2, j = 2$. By equation (10), every vertex u in a binary K-cube has exactly one opposite. Thus, deleting any vertex from K_2^d leaves its opposite with eccentricity at least $d-1$. By equation (10), there are $\binom{d}{d-1} + 1 = d + 1$ vertices in K_2^d at distance $d-1$ or d from u . Since at most $d-1$ vertices are deleted, the eccentricity of every vertex is at least $d-1$. \square



Radix j of K-cube	Number i of vertices deleted, $0 \leq i \leq f$ $f = [(j-1) \cdot \log_j n] - 1$	Radius		Diameter		Number i of vertices deleted, $0 \leq i \leq f$ $f = [(j-1) \cdot \log_j n] - 1$
		At least	At most	At least	At most	
2	0	$\log_2 n$ Theorem 7		$\log_2 n$ Theorems 9, 10		from 0 to $[\log_2 n] - 2$
	from 1 to $[\log_2 n] - 2$	$[\log_2 n] - 1$ Theorem 11	$[\log_2 n]$ Theorem 9	$\log_2 n$ Theorem 10	$[\log_2 n] + 1$ Theorem 9	$[\log_2 n] - 1$
	$[\log_2 n] - 1$		$[\log_2 n] + 1$ Theorem 9			
≥ 3	from 0 to $[\log_j n] - 1$	$\log_j n$ Theorems 8, 11		$\log_j n$ Theorems 8, 10		from 0 to $[\log_j n] - 1$
	from $[\log_j n]$ to $[(j-1) \cdot \log_j n] - 1$	$\log_j n$ Theorem 11	$[\log_j n] + 1$ Theorem 8	$\log_j n$ Theorem 10	$[\log_j n] + 1$ Theorem 8	from $[\log_j n]$ to $[(j-1) \cdot \log_j n] - 1$

Table 9: Characteristics of quorums induced by deleting vertices of d -dimensional j -ary K-cubes K_j^d . K_j^d is constructible if and only if the maximum number of faults f equals $[(j-1) \cdot \log_j n] - 1$ and $d = \log_j n$.

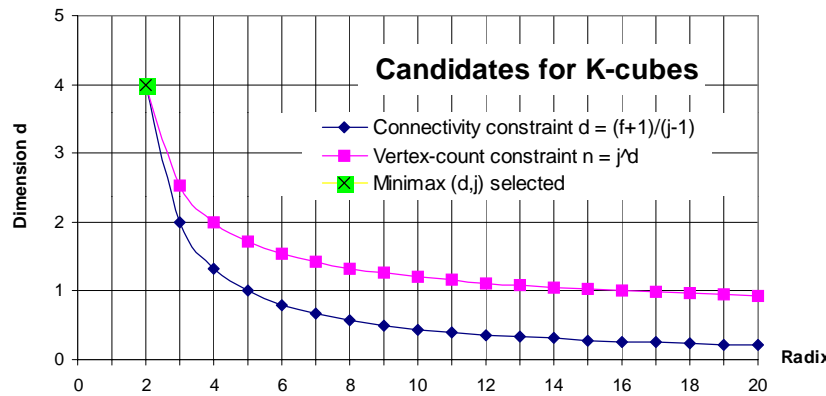


Figure 10: Exact conditions on the number $n = 16$ of vertices and maximum number $f = 3$ of faults tolerated for constructibility of a K-cube, $j = 2$, $d = 4$. A K-cube is constructible if and only if the constraints intersect at integer values of $d \geq 1$ and $j \geq 2$. Plot by GRAFT calculator described in Section 3.8.

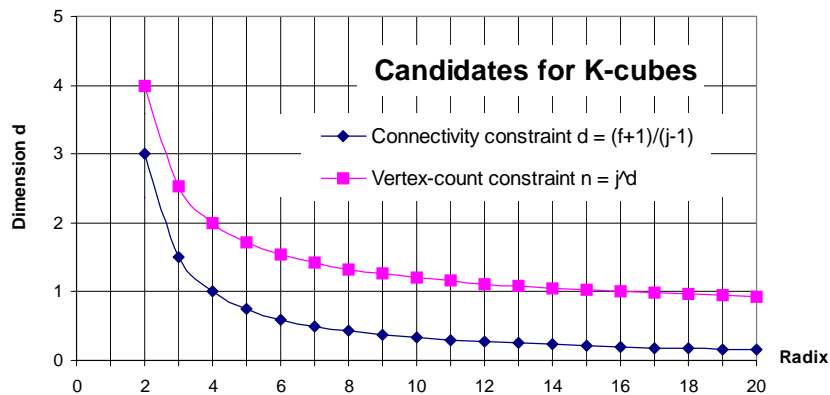


Figure 11: K-cube not feasible. At $(n, f) = (16, 2)$, constraints on connectivity and vertex count fail to intersect at integer values. However, as Figure 13 shows, it is possible to construct a K-cube-connected cycle. Plotted by GRAFT calculator described in Section 3.8





Refer to Table 9. In between the extremes $f = 0, 1$ and $f = n-2, n-1$, this section covers values of n and f satisfying $n = j^d$ and $f+1 = d(j-1)$; that is, K-cubes $\{K_j^d\} = \mathcal{G}_{n, \log n, \log n}^+$ whose dimension, radix, and vertex degree are related by formulae (5) and (6). Figure 10 illustrates these relations. However, and as illustrated by Figure 11, our results do not include the cases where degree of a vertex is not equal to that of some K-cube. As shown by Figure 13, we may be able to substitute a *K-cube-connected cycle* if the degree is less that of some K-cube. That is, when

$$f+1 < d(j-1) \tag{14}$$

3.4 Quorums from K-cube-connected Cycles

A d -dimensional j -ary K-cube-connected cycle of order n , denoted $K_j^d(n)$, is the result of replacing each of the j^d vertices of K_j^d with $n \bmod j^d$ cycles, each of which contains $\lceil n/j^d \rceil$ vertices, along with $j^d - n \bmod j^d$ cycles, each of which contains $\lfloor n/j^d \rfloor$ vertices. Refer to Figure 12. For a basis, a zero-dimensional K-cube-connected cycle $K_j^0(n)$ is a cycle with vertices labeled from 0 to $\lfloor n/j^0 \rfloor - 1$ (i.e., from 0 to $n-1$). The high order d digits of the label on a vertex u in cycle h of $K_j^d(n)$ are identical to the d digits on the label of vertex h of the corresponding K_j^d . The low order digit on u is its label in the corresponding $K_j^0(n)$. Vertex u shares an edge with vertex v if and only if i) u and v are neighbors in a basic cycle $K_j^0(n)$; or ii) the low order digits of u and v are identical, and the high order digits differ in exactly one position, or iii) there are $\lfloor n/j^d \rfloor$ vertices in the basic cycle of which u is a member, $\lceil n/j^d \rceil$ vertices in the basic cycle of which v is a member, and u and v have the highest labels in their respective basic cycles.




Figure 12: A K-cube-connected cycle $K_j^d(n)$ has minimum size if and only if (19a) or (19b) holds.

Since each basic cycle must contain at least three vertices, it follows that $n/3 \geq j^d$ (15)

is a constraint on the number of vertices in any $K_j^d(n)$. If $0 = n \bmod j^d$ then $n = m \cdot j^d$ for some positive integer m . Since it contains exactly m vertices per basic cycle, we denote such a $K_j^d(m \cdot j^d)$ by $K_{m \cdot j^d}^d$.

Each vertex of $K_{m \cdot j^d}^d$ has degree $d(j-1)+2 = f+1$ (16)

Summing the degree of every vertex counts each edge twice, hence

$$e_K(d, j, n) = \frac{1}{2} \cdot m \cdot j^d \cdot (d[j-1]+2) = \frac{1}{2} \cdot m \cdot d \cdot j^d [j-1] + m \cdot j^d \quad (\text{for } n = m \cdot j^d) \tag{17}$$

Since either j or $j-1$ is even, the first term on the righthand side of (17) is an integer; (17) is therefore an integer. Substituting $d(j-1)+2 = f+1$, we see that (17) equals $\lceil n(f+1)/2 \rceil$. Thus $0 = n \bmod j^d$ implies that the number of edges in $K_{m \cdot j^d}^d$ is exactly that of any minimum size $(f+1)$ -connected graph on $m \cdot j^d$ vertices. Suppose on the other hand that $0 \neq n \bmod j^d$. By step (iii) above, we connect the vertex with the highest label in each of the $n \bmod j^d$ long cycles to the vertex with the highest label in each of the $j^d - n \bmod j^d$ short cycles; moreover, we count these $(n \bmod j^d)(j^d - n \bmod j^d)$ “extra” edges only at $d=1$.





Summing the degree of each vertex counts each edge twice. The number of edges in $K_j^d(n)$ is therefore

$$e_K(d, j, n) = \lceil \frac{1}{2} \cdot [n \cdot (d(j-1)+2) + (n \bmod j^d)(j^d - n \bmod j^d)] \rceil \tag{18}$$

Substituting $d(j-1)+2 = f+1$, we see that (18) equals $\lceil \frac{1}{2} \cdot [n(f+1) + (n \bmod j^d)(j^d - n \bmod j^d)] \rceil$. That is, $K_j^d(n)$ has minimum size $\lceil n(f+1)/2 \rceil$ if and only if

$$\text{either} \quad \text{a) } 0 = n \bmod j^d \quad \text{or} \quad \text{b) } j = 2, d = 1, f = 2, n \text{ odd} \tag{19}$$

By comparison to K-cubes, our K-cube-connected cycles must satisfy *three* constraints: (15), (16), and (19). Despite this, and as illustrated by Figure 13, a $(d-2)$ -dimensional K-cube-connected cycle may be constructible where the corresponding d -dimensional K-cube is not. Note that (19) says quite a bit about the structure of K-cube-connected cycles $K_j^d(n)$ of size $\lceil n(f+1)/2 \rceil$: either $K_j^d(n)$ is a $K_{m \cdot j}^d$, or, for all n and $f = 2$, $K_j^d(n) = K_2^1(n)$ comprises two cycles, one with $\lceil n/2 \rceil$ vertices, the other consisting of $\lfloor n/2 \rfloor$ vertices. The latter holds since if n is not odd then 2 divides n ; in this case (19a) is satisfied, and we have a $K_{m \cdot 2}^1$. In particular, the size of any one-dimensional binary K-cube-connected cycle is the same as that $\lceil 3n/2 \rceil$ of a 3-connected graph with fewest edges. This is of some interest since $K_2^1(n)$ was, in fact, one of the candidate architectures proposed by [Charlan et al 11-Jun-1998] for X2000 avionics (indeed, Section 3.10 recommends $K_2^1(n)$ as a graph architecture of choice). Note also that our definition of a K-cube-connected cycle is somewhat different from that described by [Preparata and Vallemin 1981] and analyzed by [Banerjee et al 1986]. For the remainder of this section we write m in place of the integer value $\lfloor n/j^d \rfloor$.

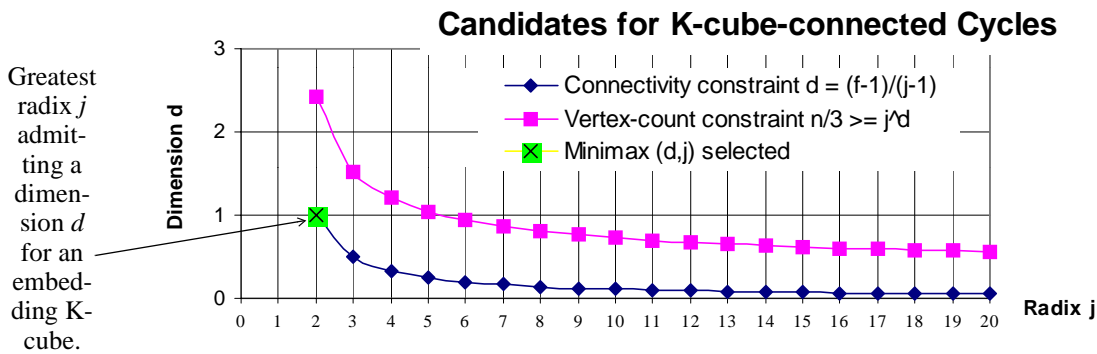


Figure 13: K-cube-connected cycle feasible where K-cube is not. Compare to Figure 12.

Theorem 12. (Connectivity, upper bound on diameter.) If $j \geq 3$ then between vertices u and v in $K_{m \cdot j}^d$ there are $d(j-1)+2$ interior-disjoint paths, none of whose length exceeds $d + m - 1$. The length of $d(j-1)+1$ of these paths is at most $d + \lfloor m/2 \rfloor + 1$. The length of $d+1$ of these paths is at most $\max(2, d) + \max(2, \lfloor m/2 \rfloor)$.

Proof. Denote by C' and C'' the respective basic cycles for u and v . Suppose that $C' \neq C''$ (implying $d \geq 1$), and the low order digits of the labels on $u = u'$ and $v = u''$ are the same. By Theorem 8, between u' and u'' there are $d(j-1)$ interior-disjoint paths $\{P_{u', u''}\}$ of length at most $d+1$, at least d of which have length at most d . The low order digit of the label on each vertex of every one of the paths in $\{P_{u', u''}\}$ is the same as the low order digit on u' and u'' . Let w' and z' be neighbors of u' in C' . By Theorem 8, there are shortest interior-disjoint paths $P_{w', w''}, P_{z', z''}$ from w' resp. $z' \in C'$ to w'' resp. $z'' \in C''$; the low order digit of the label on each vertex of $P_{w', w''}$ and $P_{z', z''}$ is the same as the low order digit on w' and w''





resp. z' and z'' ; neither $P_{w', w''}$ nor $P_{z', z''}$ traverses more than d edges. The $d(j-1)+2$ interior-disjoint paths $(u', w') \cup P_{w', w''} \cup (w'', u'')$, $(u', z') \cup P_{z', z''} \cup (z'', u'')$, and $\{P_{u', u''}\}$ have length at most $d + 2 \leq d + \lfloor m/2 \rfloor + 1 \leq d + m - 1$ (the latter is a consequence of (15)). The length of $(u', w') \cup P_{w', w''} \cup (w'', u'')$ is at most $d + 2 \leq \max(2, d) + \max(2, \lfloor m/2 \rfloor)$; the length of the d shortest paths in $\{P_{u', u''}\}$ is at most d ; therefore, at least $d+1$ of the prescribed paths have length at most $\max(2, d) + \max(2, \lfloor m/2 \rfloor)$.




Figure 14: Illustration of the second part of the proof of Theorem 12, proof of Theorem 19.

Suppose that $C' \neq C''$ (implying $d \geq 1$) and that the low order digits of the labels on $u = u'$ and $v = v''$ are different. Refer to Figure 14. By Theorem 8, there are $d(j-1)$ interior-disjoint paths $\{P_{u', u''}\}$ between $u' \in C'$ and $u'' \in C''$, all of which have length at most $d+1$, and at least d of which have length at most d . The low order digit of the label on each vertex of every one of the paths in $\{P_{u', u''}\}$ is the same as the low order digit on u' and u'' . Let (u^*, u'') be the last edge in any one of the $d(j-1)-1$ longest such paths $P_{u', u''}$, and replace this edge with a shortest path P_{u^*, v^*} in C^* followed by the edge (v^*, v'') . Applying this process to all but a shortest path in $\{P_{u', u''}\}$ yields $d(j-1)-1$ interior-disjoint paths $\{P_{u', u^*} \cup P_{u^*, v^*} \cup (v^*, v'')\}$, each of length at most $d+1 + \lfloor m/2 \rfloor$, with at least $d-1$ of these paths no longer than $d + \lfloor m/2 \rfloor$. Let $Q_{u', u''}$ be the path in $\{P_{u', u''}\}$ (of length at most d) not modified by the preceding procedure. Augment $Q_{u', u''}$ with $P_{u'', v''}$, the shortest path (of length at most $\lfloor m/2 \rfloor$) between u'' and v'' in C'' . Let $P_{u', v'}$ be the shortest path (of length at most $\lfloor m/2 \rfloor$) in C' between u' and v' . Augment $P_{u', v'}$ with $Q_{v', v''}$, a path between v' and v'' (and of length at most d) that passes through the same basic cycles as $Q_{u', u''}$. Let z' be the neighbor of v' along a path $P_{u', z'} \cup (z', v')$ from u' to v' in C' , such that $P_{u', z'}$ does not intersect the interior of $P_{u', v'}$ (since C' contains at least three vertices, z' is distinct from u'). Augment $P_{u', z'}$ (whose length is at most $m-2$) with $Q_{z', z''}$, a path between z' and z'' (and of length at most d) that passes through the same basic cycles as $Q_{u', u''}$ and $Q_{v', v''}$. The combination of $P_{u', z'} \cup Q_{z', z''} \cup (z'', v'')$, $Q_{u', u''} \cup P_{u'', v''}$, $P_{u', v'} \cup Q_{v', v''}$, and



$\{P_{u', u^*} \cup P_{u^*, v^*} \cup (v^*, v'')\}$ constitute $d(j-1)+2$ interior-disjoint paths. $P_{u', z'} \cup Q_{z', z''} \cup (z'', v'')$ has length at most $d + m - 1$. $d(j-2)$ of the paths in $\{P_{u', u^*} \cup P_{u^*, v^*} \cup (v^*, v'')\}$ have length at most $d + \lfloor m/2 \rfloor + 1$. $Q_{u', u''} \cup P_{u'', v''}$, $P_{u', v'} \cup Q_{v', v''}$, along with the $d-1$ shortest paths in $\{P_{u', u^*} \cup P_{u^*, v^*} \cup (v^*, v'')\}$, comprise $d+1$ interior-disjoint paths, each of whose length is at most $d + \lfloor m/2 \rfloor \leq \max(2, d) + \max(2, \lfloor m/2 \rfloor)$.

Suppose that $C' = C''$. By Theorem 7, there are $d(j-1)$ basic cycles $\{C^*\}$ connected to C' via edges $\{(u', u^*)\}$. Within any such C^* there is, from u^* to v^* , a shortest path P_{u^*, v^*} of length at most $\lfloor m/2 \rfloor$. Here v^* is the vertex whose low order digit is the same as that on v' , but whose d high order digits are the same as those on u^* (i.e., that differ in one digit from the d high order digits on the label of u' or v'). The $d(j-1)$ paths $\{(u', u^*) \cup P_{u^*, v^*} \cup (v^*, v')\}$ are interior-disjoint, and each has length at most $2 + \lfloor m/2 \rfloor \leq d + \lfloor m/2 \rfloor + 1$ (the inequality is satisfied for $d \geq 1$, and does not pertain at $d = 0$). To this add the shortest $P_{u', v'}$ and longest $Q_{u', v'}$ paths between u and v , strictly contained in C' . The latter two paths have lengths at most $\lfloor m/2 \rfloor \leq d + \lfloor m/2 \rfloor$ resp. $m-1 \leq d + m - 1$. If $d = 0$ then there is $d+1 = 1$ path (i.e., $P_{u', v'}$) of length $\lfloor m/2 \rfloor \leq \max(2, d) + \max(2, \lfloor m/2 \rfloor)$. Otherwise, $d \geq 1$ and any $d+1$ paths from $\{(u', u^*) \cup P_{u^*, v^*} \cup (v^*, v')\}$, traverse no more than $2 + \lfloor m/2 \rfloor \leq \max(2, d) + \max(2, \lfloor m/2 \rfloor)$ edges. Verify that $P_{u', v'}$ and $Q_{u', v'}$ are interior-disjoint with each other and with $\{(u', u^*) \cup P_{u^*, v^*} \cup (v^*, v')\}$. \square

Theorem 13. (Connectivity, upper bound on diameter.) If $j = 2$ then between vertices u and v in $K_{m,2}^d$ there are $d+2$ interior-disjoint paths. The length of each of these paths does not exceed $d + m - 1$. The length of $d+1$ of these paths is at most $d + \lfloor m/2 \rfloor + 1$; the length of d of these paths is at most $\max(2, d) + \max(2, \lfloor m/2 \rfloor)$.

Proof. Denote by C' and C'' the respective basic cycles for u and v . Suppose that $C' \neq C''$ (implying $d \geq 1$), and the low order digits of the labels on $u = u'$ and $v = u''$ are the same. By Theorem 9, between u' and u'' there are d interior-disjoint paths $\{P_{u', u''}\}$ of length at most $d+1$, at least $d-1$ of which have length at most d . The low order digit of the label on each vertex of every one of the paths in $\{P_{u', u''}\}$ is the same as the low order digit on u' and u'' . Let w' and z' be neighbors of u' in C' . By Theorem 9, there are shortest interior-disjoint paths $P_{w', w''}$, $P_{z', z''}$ from w' resp. z' in C' to w'' resp. z'' in C'' . The low order digit of the label on each vertex of $P_{w', w''}$ and $P_{z', z''}$ is the same as the low order digit on w' and w'' resp. z' and z'' ; neither $P_{w', w''}$ nor $P_{z', z''}$ traverses more than d edges. The $d+2$ interior-disjoint paths $(u', w') \cup P_{w', w''} \cup (w'', u'')$, $(u', z') \cup P_{z', z''} \cup (z'', u'')$, and $\{P_{u', u''}\}$ have length at most $d+2 \leq d + \lfloor m/2 \rfloor + 1 \leq d + m - 1$ (the latter inequality holds by (15)). At least $d+1$ of these paths have length at most $d + 2 \leq d + 1 + \lfloor m/2 \rfloor$. Among the latter, at least d paths have length at most $d + 2 \leq \max(2, d) + \max(2, \lfloor m/2 \rfloor)$.

Suppose that $C' \neq C''$ (implying $d \geq 1$), and the low order digits of the labels on $u = u'$ and $v = v''$ are different. As in the proof of Theorem 12, the essential argument is illustrated by Figure 14. By Theorem 9, there are d interior-disjoint paths $\{P_{u', u''}\}$ between u' and $u'' \in C''$, all of which have length at most $d+1$, and at least $d-1$ of which have length at most d . The low order digit of the label on each vertex of every one of the paths in $\{P_{u', u''}\}$ is the same as the low order digit on u' and u'' . Let (u^*, u'') be the last edge in any one of the $d-1$ longest such paths $P_{u', u''}$, and replace this edge with a shortest path P_{u^*, v^*} in C^* followed by the edge (v^*, v'') . Applying this process to all but a shortest path in $\{P_{u', u''}\}$ yields $d-1$ interior-disjoint paths $\{P_{u', u^*} \cup P_{u^*, v^*} \cup (v^*, v'')\}$, each of length at most $d+1 + \lfloor m/2 \rfloor$, with at least $\max(0, d-2)$ of these





paths no longer than $d + \lfloor m/2 \rfloor$. Let $Q_{u', u''}$ be the path in $\{P_{u', u''}\}$ (of length at most d) not modified by the preceding procedure. Augment $Q_{u', u''}$ with $P_{u'', v''}$, the shortest path (of length at most $\lfloor m/2 \rfloor$) between u'' and v'' in C'' . Let $P_{u', v'}$ be the shortest path (of length at most $\lfloor m/2 \rfloor$) in C' between u' and v' . Augment $P_{u', v'}$ with $Q_{v', v''}$, a path between v' and v'' (and of length at most d) that passes through the same basic cycles as $Q_{u', u''}$. Let z' be the neighbor of v' along a path $P_{u', z'} \cup (z', v')$ from u' to v' in C' , such that $P_{u', z'}$ does not intersect the interior of $P_{u', v'}$ (since C' contains at least three vertices, z' is distinct from u'). Augment $P_{u', z'}$ (whose length is at most $m - 2$) with $Q_{z', z''}$, a path between z' and z'' (and of length at most d) that passes through the same basic cycles as $Q_{u', u''}$ and $Q_{v', v''}$. The combination of $P_{u', z'} \cup Q_{z', z''} \cup (z'', v'')$, $Q_{u', u''} \cup P_{u'', v''}$, $P_{u', v'} \cup Q_{v', v''}$, and $\{P_{u', u^*} \cup P_{u^*, v^*} \cup (v^*, v'')\}$ constitutes $d+2$ interior-disjoint paths. $P_{u', z'} \cup Q_{z', z''} \cup (z'', v'')$ has length at most $d + m - 1$. One of the paths in $\{P_{u', u^*} \cup P_{u^*, v^*} \cup (v^*, v'')\}$ has length at most $d + \lfloor m/2 \rfloor + 1$. At least d interior-disjoint paths have length at most $d + \lfloor m/2 \rfloor \leq \max(2, d) + \max(2, \lfloor m/2 \rfloor)$; namely, $Q_{u', u''} \cup P_{u'', v''}$, $P_{u', v'} \cup Q_{v', v''}$, and $d-2$ of the paths in $\{P_{u', u^*} \cup P_{u^*, v^*} \cup (v^*, v'')\}$.

Suppose that $C' = C''$. By Theorem 7, there are d basic cycles $\{C^*\}$ connected to C' via edges $\{(u', u^*)\}$. Within any such C^* there is, from u^* to v^* , a shortest path P_{u^*, v^*} of length at most $\lfloor m/2 \rfloor$. Here v^* is the vertex whose low order digit is the same as that on v' , but whose d high order digits are the same as those on u^* (i.e., that differ in one digit from the d high order digits on the label of u' or v'). The d paths $\{(u', u^*) \cup P_{u^*, v^*} \cup (v^*, v'')\}$ are interior-disjoint, and each has length at most $2 + \lfloor m/2 \rfloor \leq d + \lfloor m/2 \rfloor + 1$. To this collection add the shortest $P_{u', v'}$ and longest $Q_{u', v'}$ paths between u and v , strictly contained in C' . The latter two paths have lengths at most $\lfloor m/2 \rfloor \leq d + \lfloor m/2 \rfloor$ resp. $m-1 \leq d + m - 1$. If $d = 0$ then there is $d+1 = 1$ path (i.e., $P_{u', v'}$) of length $\lfloor m/2 \rfloor \leq \max(2, d) + \max(2, \lfloor m/2 \rfloor)$. Otherwise, $d \geq 1$ and any $d+1$ paths from $\{(u', u^*) \cup P_{u^*, v^*} \cup (v^*, v'')\}$, traverse no more than $2 + \lfloor m/2 \rfloor \leq \max(2, d) + \max(2, \lfloor m/2 \rfloor)$ edges. Verify that $P_{u', v'}$ and $Q_{u', v'}$ are interior-disjoint with each other and with $\{(u', u^*) \cup P_{u^*, v^*} \cup (v^*, v'')\}$. \square

Theorem 14. (Connectivity, upper bound on diameter.) Between vertices u and v in a $K_2^1(2m+1)$ there are 3 interior-disjoint paths, each of which is no longer than $m+1$. The length of 2 of these paths is at most $2 + \lfloor m/2 \rfloor$; the length of one of these paths is at most $1 + \lfloor m/2 \rfloor$.

Proof. By (15), $m \geq 3$. Without loss of generality let C' be the basic cycle with the fewer number m of vertices. Denote by C' and C'' the respective basic cycles for u and v . We consider six cases.

Case I. Suppose that $C' \neq C''$ and that the low order digits of the labels on $u = u'$ and $v = u''$ are the same. Vertices u' and v'' are connected by an edge (u', v'') whose length is at most $1 \leq 1 + \lfloor m/2 \rfloor$. Let w'' and z'' be neighbors of u'' in C'' . If the low order digit on one of w'' or z'' (say, z'') equals m then either there is a) a path $(u', z'') \cup (z'', v'')$ (i.e., u' is vertex $m-1$ of C') or b) a path $(u', x') \cup (x', z'') \cup (z'', v'')$, (i.e., x' is vertex $m-1$ of C'). If (a) then there is also a path $(u', w') \cup (w', w'') \cup (w'', v'')$ (i.e., w' and w'' are vertices in C' resp. C'' whose low order digit equals $m-2$). If (b) then there is also a path $(u', w') \cup (w', w'') \cup (w'', v'')$ (i.e., w' and w'' are vertices in C' resp. C'' whose low order digit equals 1). If the low order digit on neither w'' or z'' equals m then, in addition to (u', v'') , trace the paths $(u', w') \cup (w', w'') \cup (w'', v'')$ and $(u', z') \cup (z', z'') \cup (z'', v'')$.






Figure 15: Illustration of Case II of Theorem 14, u and v in different cycles, v the “extra” vertex.

Case II. Suppose that $C' \neq C''$ and that the low order digit of $v = v''$ equals m (thus $u = u'$ and v'' differ in their low order digits). Refer to Figure 15A. Let x' be vertex $m-1$ of C' (with two neighbors in C'' , x'' and v''); denote by $P_{u', x'}$ the path in C' from $u = u'$ to x' , according to increasing value of the labels. With w' the zeroth vertex in C' (w' and u' may be identical), write $P_{u', w'}$ for the path in C' from u' to w' , in descending order of labels. Taking w'' to be the zeroth vertex in C'' , paths $P_{u', x'} \cup (x', v'')$ and $P_{u', w'} \cup (w', w'') \cup (w'', v'')$ are interior-disjoint. The length of $P_{u', x'}$ plus the length of $P_{u', w'}$ equals $m-1$. $P_{u', w'}$ has length $0 \leq k \leq m-1$; $P_{u', x'}$ has length $m-k-1$. That is, $P_{u', x'} \cup (x', v'')$ has length $k+1$ and $P_{u', w'} \cup (w', w'') \cup (w'', v'')$ has length $m-k+1$. If $k \leq m-k$ then the length of $P_{u', x'} \cup (x', v'')$ is at most $1 + \lfloor m/2 \rfloor$ and the length of $P_{u', w'} \cup (w', w'') \cup (w'', v'')$ is at most $m+1$. In addition, trace $(u', u'') \cup P_{u'', x''} \cup (x'', v'')$, where in C'' vertex u'' and path $P_{u'', x''}$ are counterparts to u' and $P_{u', x'}$ in C' . $(u', u'') \cup P_{u'', x''} \cup (x'', v'')$ has length $k+2 \leq 2 + \lfloor m/2 \rfloor$, and is interior-disjoint from $P_{u', x'} \cup (x', v'')$ and $P_{u', w'} \cup (w', w'') \cup (w'', v'')$. Refer to Figure 15B. If $k \geq m-k+1$ then the length of $P_{u', w'} \cup (w', x') \cup (x', v'')$ is at most $\lfloor (m+1)/2 \rfloor \leq 1 + \lfloor m/2 \rfloor$ and the length of $P_{u', y'} \cup (y', y'') \cup (y'', x'') \cup (x'', v'')$ is at most $m+1$. Here y' and y'' are the vertices of C' and C'' whose low order digit equals $m-2$. In addition, trace $(u', u'') \cup P_{u'', w''} \cup (w'', v'')$, where in C'' vertex u'' and path $P_{u'', w''}$ are counterparts to u' and $P_{u', w'}$ in C' . $(u', u'') \cup P_{u'', w''} \cup (w'', v'')$ has length $k+1 \leq \lfloor (m+1)/2 \rfloor \leq 1 + \lfloor (m+1)/2 \rfloor$, and is interior-disjoint from $P_{u', w'} \cup (w', x') \cup (x', v'')$ and $P_{u', y'} \cup (y', y'') \cup (y'', x'') \cup (x'', v'')$.

Case III. Suppose that $C' \neq C''$, that the low order digit of $v = v''$ is not equal to m (i.e., v'' is not the “extra” vertex z'' in C''), and that the value of the low order digit on the label of u' is less than that on v'' . Refer to Figure 16A. Let x' be vertex $m-1$ of C' (with two neighbors in C'' , x'' and v''); denote by $P_{u', v'}$ the path in C' from u' to x' , according to increasing value of the labels. With v' the vertex in C' whose low order digit is the same as that on v'' , write $P_{u', v'}$ for the path in C' from u' to v' , in descending order of



labels. Denote by $P_{u',x'}$ the path in C' from u' to x' , in ascending order of labels. Let $P_{z'',v''}$ be the path in C'' from z'' to v'' that includes the zeroth vertex of C'' . Paths $P_{u',v'} \cup (v', v'')$ and $P_{u',x'} \cup (x', z'') \cup P_{z'',v''}$ are interior-disjoint. The length of $P_{u',v'}$ plus the length of $P_{u',x'}$ and $P_{z'',v''}$ equals m . $P_{u',v'}$ has length $0 \leq k \leq m$; the length of $P_{u',x'}$ plus the length of $P_{z'',v''}$ is $m-k$. That is, $P_{u',v'} \cup (v', v'')$ has length $k+1$ and $P_{u',x'} \cup (x', z'') \cup P_{z'',v''}$ has length $m-k+1$. If $k \leq m-k$ then, as shown in Figure 16A, the length of $P_{u',v'} \cup (v', v'')$ is at most $1 + \lfloor m/2 \rfloor$ and the length of $P_{u',w'} \cup (w', w'') \cup (w'', v'')$ is at most $m+1$. In addition, trace $(u', u'') \cup P_{u'',v''}$, where in C'' vertex u'' and path $P_{u'',v''}$ are counterparts to u' and $P_{u',v'}$ in C' . $(u', u'') \cup P_{u'',v''}$ has length $k+1 \leq 2 + \lfloor m/2 \rfloor$, and is interior-disjoint from $P_{u',v'} \cup (v', v'')$ and $P_{u',x'} \cup (x', z'') \cup P_{z'',v''}$. If $k \geq m-k+1$ then, as shown in Figure 16B, $P_{u',x'} \cup P_{x',v'} \cup (v', v'')$ has length at most $1 + \lfloor m/2 \rfloor$, $(u', u'') \cup P_{u'',x''} \cup (x'', z'') \cup P_{z'',v''}$ has length at most $2 + \lfloor m/2 \rfloor$, and $P_{u',w'} \cup (w', w'') \cup (w'', v'')$ has length at most $m+1$. Here w' and w'' are the vertices of C' and C'' whose low order digit is one greater than those of v' and v'' (w' and w'' may be identically u' resp. u'').




Figure 16: Case III of Theorem 14, u and v in different cycles, u 's low order digit less than that of v .

Case IV. As for Case III, only in this instance assume that the value of the low order digit on the label of u' is greater than that on v'' . Refer to Figure 17A. With v' the vertex in C' whose low order digit is the same as that on v'' , write $P_{u',v'}$ for the path in C' from u' to v' , in descending order of labels. Denote by $P_{u',x'}$ the path in C' from u' to x' , in ascending order of labels. Let $P_{z'',v''}$ be the path in C'' from z'' to v'' that passes through the zeroth vertex of C'' . Paths $P_{u',v'} \cup (v', v'')$ and $P_{u',x'} \cup (x', z'') \cup P_{z'',v''}$ are interior-disjoint. The length of $P_{u',v'}$ plus the length of $P_{u',x'}$ and $P_{z'',v''}$ equals m . $P_{u',v'}$ has length $0 \leq k \leq m$; the length of $P_{u',x'}$ plus the length of $P_{z'',v''}$ is $m-k$. That is, $P_{u',v'} \cup (v', v'')$ has length $k+1$ and $P_{u',x'} \cup (x', z'') \cup P_{z'',v''}$ has length $m-k+1$. If $k \leq m-k$ then, as shown in Figure 17A, the length of $P_{u',v'} \cup (v', v'')$ is at most $1 + \lfloor m/2 \rfloor$ and the length of $P_{u',w'} \cup (w', w'') \cup (w'', v'')$ is at most $m+1$. In addition, trace $(u', u'') \cup P_{u'',v''}$, where in C'' vertex u'' and path $P_{u'',v''}$ are counterparts to u' and $P_{u',v'}$ in





C' . $(u', u'') \cup P_{u'', v''}$ has length $k+1 \leq 2 + \lfloor m/2 \rfloor$, and is interior-disjoint from $P_{u', v'} \cup (v', v'')$ and $P_{u', x'} \cup (x', z'') \cup P_{z'', v''}$. If $k \geq m-k+1$ then, as shown in Figure 17B, $P_{u', x'} \cup P_{x', v'} \cup (v', v'')$ has length at most $1 + \lfloor m/2 \rfloor$, $(u', u'') \cup P_{u'', x''} \cup (x'', z'') \cup P_{z'', v''}$ has length at most $2 + \lfloor m/2 \rfloor$, and $P_{u', w'} \cup (w', w'') \cup (w'', v'')$ has length at most $m+1$. Here w' and w'' are the vertices of C' and C'' whose low order digit is one greater than those of v' and v'' (w' and w'' may be identically u' resp. u'').

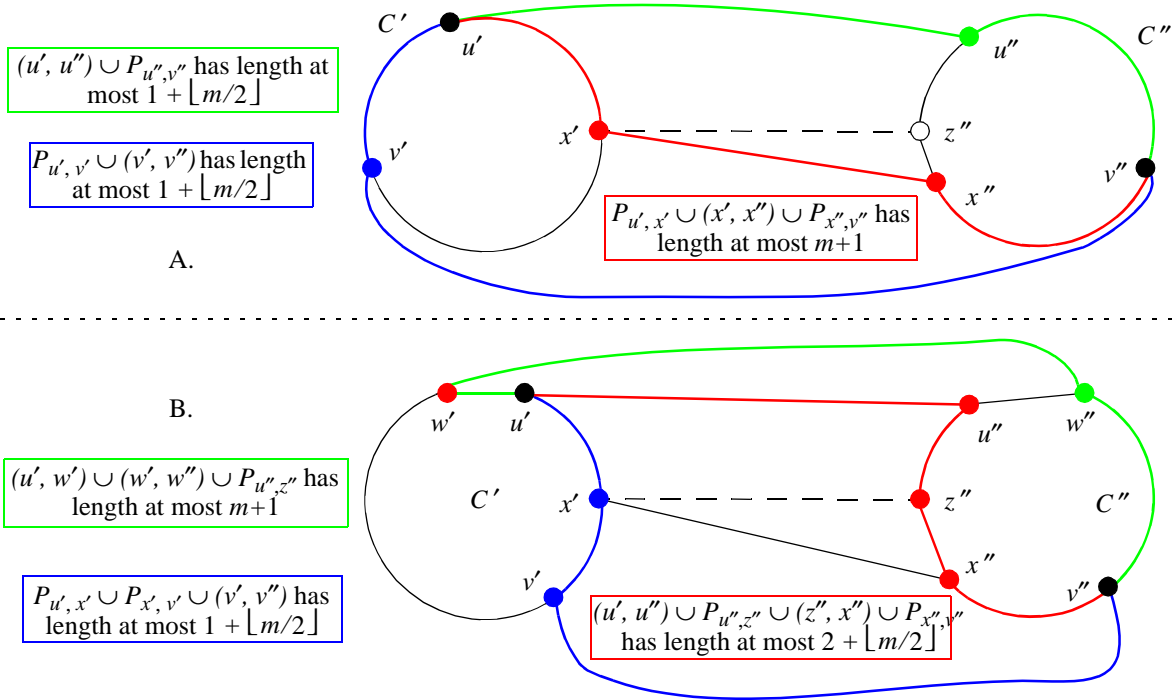


Figure 17: Case IV of Theorem 14, u and v in different cycles, u 's low order digit greater than that of v .

Case V. Suppose that both $u = u'$ and $v = v'$ are in C' . Let $P_{u', v'}$ and $Q_{u', v'}$ be interior-disjoint paths in C' between u' and v' . The sum of the lengths of the paths $P_{u', v'}$ and $Q_{u', v'}$ equals m . Without loss of generality assume that $P_{u', v'}$ is the shorter path with length k ; i.e., $1 \leq k \leq \lfloor m/2 \rfloor \leq 1 + \lfloor m/2 \rfloor$. The length of $Q_{u', v'}$ is therefore $m-k$, with $\lceil m/2 \rceil \leq m-k \leq m-1$. At most one of u' and v' can have more than one neighbor in C'' . Without loss of generality assume that v' only has one neighbor v'' in C'' . Let u'' be u' 's neighbor in C'' that is closest to v'' via a path $P_{u'', v''}$, strictly contained in C'' . The length of $P_{u'', v''}$ is $k+1$, and so $(u', u'') \cup P_{u'', v''} \cup (v'', v')$ traverses at most $k+3$ edges. If the length of $P_{u', v'}$ is at its maximum $\lfloor m/2 \rfloor \leq 1 + \lfloor m/2 \rfloor$ then $Q_{u', v'}$ has length $\lceil m/2 \rceil \leq 2 + \lfloor m/2 \rfloor$ and $(u', u'') \cup P_{u'', v''} \cup (v'', v')$ has length at most $3 + \lfloor m/2 \rfloor \leq m+1$. (The latter is established by considering separately the cases where m is even and odd, and noting that 4 is the minimum value of m even.) Otherwise, the length of $(u', u'') \cup P_{u'', v''} \cup (v'', v')$ is at most $2 + \lfloor m/2 \rfloor$, and $Q_{u', v'}$ has length at most $m-1 \leq m+1$. Verify that $P_{u', v'}$, $Q_{u', v'}$, and $(u', u'') \cup P_{u'', v''} \cup (v'', v')$ are interior-disjoint.

Case VI. Suppose that both $u = u''$ and $v = v''$ are in C'' . Let $P_{u'', v''}$ and $Q_{u'', v''}$ be interior-disjoint paths in C'' between u'' and v'' . The sum of the lengths of the paths $P_{u'', v''}$ and $Q_{u'', v''}$ equals $m+1$. Without loss of generality assume that $P_{u'', v''}$ is the shorter path with length at most $\lfloor (m+1)/2 \rfloor \leq 1 + \lfloor m/2 \rfloor$. The length of $Q_{u'', v''}$ is therefore no greater than $m \leq m+1$. Let u' be u'' 's neighbor in C' that is closest to v'' via a



path $R_{u', v'}$, strictly contained in C' . While it is possible that the low order digit on the label of u' or v' differs from that of u'' resp. v'' , in any case the length of $P_{u', v'}$ is at most $\lfloor m/2 \rfloor$. Thus $(u'', u') \cup P_{u', v'} \cup (v', v'')$ traverses at most $2 + \lfloor m/2 \rfloor$ edges.

In every case, we have three interior-disjoint paths, each of whose length is at most $m+1$; two of the paths have length at most $2 + \lfloor m/2 \rfloor$; one path is no longer than $1 + \lfloor m/2 \rfloor$. \square

Corollary 14.1. $K_j^d(n)$ has connectivity $d(j-1)+2$ and minimum size $\lceil \frac{1}{2} \cdot [n \cdot (d[j-1]+2)] \rceil = \lceil n(f+1)/2 \rceil$ if and only if conditions (15), (16), and (19) are satisfied.

Theorems 12, 13, and 14 extend Theorems 8 and 9 from K-cubes to K-cube-connected cycles. For $0 \leq i \leq f = d(j-1)+1$, the theorems enable us to bound from above the diameter of any quorum H induced by deleting i vertices from a $K_j^d(n)$ whose size, with respect to $(f+1)$ -connectivity, is minimum.

If $j > 2$ and $0 \leq i \leq d$ then from any vertex in H we can reach all other vertices by a path of length at most $\max(2, d) + \max(2, \lfloor m/2 \rfloor)$. If $j > 2$ and $d+1 \leq i \leq d(j-1)$ then from any vertex in H we can reach all other vertices by a path of length at most $d + \lfloor m/2 \rfloor + 1$. If $j > 2$ and $i = d(j-1)+1 = f$ then from any vertex in H we can reach all other vertices by a path of length at most $d + m - 1$.

If $j = 2$, n is even, and $0 \leq i \leq d-1$ then from any vertex in H we can reach all other vertices by a path of length at most $\max(2, d) + \max(2, \lfloor m/2 \rfloor)$. If $j = 2$, n is even, and $i = d$ then from any vertex in H we can reach all other vertices by a path of length at most $d + \lfloor m/2 \rfloor + 1$. If $j = 2$, n is even, and $i = d+1 = f$ then from any vertex in H we can reach all other vertices by a path of length at most $d+m-1$.

If $j = 2$, n is odd, $d = 1$, and $i = 0$ then from any vertex in H we can reach all other vertices by a path of length at most $1 + \lfloor (n-1)/4 \rfloor$. If $j = 2$, n is odd, $d = 1$, and $i = 1$ then from any vertex in H we can reach all other vertices by a path of length at most $2 + \lfloor (n-1)/4 \rfloor$. If $j = 2$, n is odd, $d = 1$, and $i = 2 = f$ then from any vertex in H we can reach all other vertices by a path of length at most $1 + (n-1)/2$.

Unlike our development for K-cubes, we do not exhibit a class of extremal counterexamples that show how the preceding bounds are best possible. Instead we pursue analogs to Theorems 10 and 11. Our development extends (8) and (10) to $B_j(d, i, m)$, the number of vertices at distance i from any vertex u in $K_{m,j}^d$. Without loss of generality assume that u 's label consists of all zeros. At $d = 0$ we have an n -vertex cycle $K_j^0(m \cdot j^d)$. For $1 \leq i < \lfloor m/2 \rfloor$ two vertices lie at distance i from u . For $i = \lfloor m/2 \rfloor = m/2$, m is even, and one vertex lies at distance $m/2$ from u ; for $i = \lfloor m/2 \rfloor = (m-1)/2$, m is odd, and two vertices lie at distance $(m-1)/2$ from u . In higher dimensions, suppose that $u = u'$ is in K' , one of j copies of $K_{m,j}^{d-1}$ comprising $K_{m,j}^d$. The number of vertices distance i from u' equals the number of vertices in K' that are distance i from u' plus $j-1$ times the number of vertices in any other copy K'' of $K_{m,j}^{d-1}$ that are distance $i-1$ from the counterpart $u'' \in K''$ of u' . Thus, the recurrence relation for $B_j(d, i, m)$ is identical to that (8) for $B_j(d, i)$, but with different boundary conditions:

$$B_j(d, 0, m) = 1, B_j(0, i, m) = 2, 1 \leq i < \lfloor m/2 \rfloor, B_j(0, \lfloor m/2 \rfloor, m) = 2 - [(m-1) \bmod 2] \quad (20)$$

Table 10 illustrates the triangular computation of $B_j(d, i, m)$, analogous to that shown for $B_j(d, i)$ in Table 8. With boundary conditions (20), we do not know how to solve the recurrence of (8) in a closed-form fashion akin to (10). However, we can shed light on the solution imposed by (20) by noting that a shortest path from the origin u to any vertex v in $K_{m,j}^d$ passes through the origin of some other basic cycle. Therefore, the number of vertices at distance i from u equals the number of vertices having a label that differs from that of u in i of the d high order digits (but whose low order digit is the same), plus (fanning out in the respective cycles) twice the number of vertices having a label that differs from that of u in $i-1$ of the d high order digits, plus ... plus twice the number of vertices having a label that differs from that of u in $i - \lfloor m/2 \rfloor - 1$ of the d high order digits, plus either twice the number of vertices having a label that differs





from that of u in $i - \lfloor m/2 \rfloor$ of the d high order digits (if m is odd) or the number of vertices having a label that differs from that of u in $i - \lfloor m/2 \rfloor$ of the d high order digits (if m is even). That is:

$$B_j(d, i, m) = \left[\begin{array}{c} \min\left(i, \lfloor \frac{m-1}{2} \rfloor\right) \\ \sum_{h=\max(0, i-d)} (j-1)^{i-h} \binom{d}{i-h} \end{array} \right] + \left[\begin{array}{c} \min\left(i, \lceil \frac{m-1}{2} \rceil\right) \\ \sum_{h=\max(1, i-d)} (j-1)^{i-h} \binom{d}{i-h} \end{array} \right] \quad (21)$$

For the summand in (21) we have used relation (10). The lower and upper indices follow by noting that $B_j(d, i)$ is nonzero if and only if $0 \leq i - h \leq d$. In two cases of particular interest, (21) reduces to

$$B_j(d, d + \lfloor m/2 \rfloor, m) = (2 - \lfloor (m-1) \bmod 2 \rfloor) \cdot (j-1)^d \quad (22)$$

$$B_j(d, d + \lfloor m/2 \rfloor - 1, m) = 2 \cdot (j-1)^d + d \cdot (2 - \lfloor (m-1) \bmod 2 \rfloor) \cdot (j-1)^{d-1} \quad (23)$$

$\downarrow d$	$j = 2, m = 5$							$j = 3, m = 4$									
	$\rightarrow i$	0	1	2	3	4	5	6	7	0	1	2	3	4	5	6	7
0		1	2	2						1	2	1					
1		1	3	4	2					1	4	5	2				
2		1	4	7	6	2				1	6	13	12	4			
3		1	5	11	13	8	2			1	8	25	38	28	8		
4		1	6	16	24	21	10	2		1	10	41	88	104	64	16	
5		1	7	22	40	45	31	12	2	1	12	61	170	280	272	144	32

Table 10: Number $B_j(d, i, m)$ of vertices at graph distance i from any other in $K_{m,j}^d$, formed by replacing each of the j^d vertices of a d -dimensional j -ary K-cube with a cycle on m vertices.

Theorem 15. Let H be any quorum induced by deleting i vertices from $K_{m,j}^d$, $0 \leq i \leq f = d(j-1)+1$. The diameter of H is at least $d + \lfloor m/2 \rfloor$.

Proof. Vertices u and v are *opposite* if they are distance $d + \lfloor m/2 \rfloor$ apart; i.e., the high order digits of their labels differ in every position and their shortest path along a corresponding m -vertex cycle is has maximum length $\lfloor m/2 \rfloor$. If m is odd then, by equation (22), any given vertex u has $2(j-1)^d$ opposites; that is, there are $2(j-1)^d$ opposite pairs that include u . Summing over all $m \cdot j^d$ vertices counts every pair of opposites twice, and the total number of opposite pairs equals $j^d(j-1)^d$. Each vertex we delete from $K_{m,j}^d$ removes at most $2(j-1)^d$ opposite pairs. Therefore, there remains at least one opposite pair as long as

$$\lfloor d(j-1)+1 \rfloor \cdot 2 \cdot (j-1)^d < 2 \cdot m \cdot j^d \cdot (j-1)^d \quad (24)$$

By (15), $m \geq 3$; since $j \geq 2$, inequality (24) is satisfied if $d+1 \leq 3 \cdot 2^d$. The latter holds if $d \leq 2^{d+1}$, which, by differentiation, is readily verified for all nonnegative d . If m is odd then we obtain inequality (24) with both sides divided by two. The theorem follows since $d+1 \leq 3 \cdot 2^d$. □

Theorem 16. Suppose that $j \geq 3$ and let H be any quorum induced by deleting i vertices from $K_{m,j}^d$, $0 \leq i \leq f = d(j-1)+1$. If $i = 0$ then the radius of H equals $d + \lfloor m/2 \rfloor$. If $d(j-1) \leq i \leq d(j-1)+1$, m is even, and $d \leq 2$ then the radius of H is at least $d + \lfloor m/2 \rfloor - 1$. Otherwise the radius of H is at least $d + \lfloor m/2 \rfloor$.





Proof. If $i = 0$ then $H = K_{m,j}^d$; by (21), the radius of $K_{m,j}^d$ equals $d + \lfloor m/2 \rfloor$. Suppose m is odd and $j \geq 3$. By (22), there remains at least one vertex opposite to any undeleted vertex u as long as

$$d(j-1)+1 < 2 \cdot (j-1)^d \quad (25)$$

(25) is satisfied if $d \leq 2^{d-1}$. The latter holds by Theorem 11. Suppose that m is even and $j \geq d \geq 3$. By equation (22), there remains at least one vertex opposite to any undeleted vertex u as long as

$$d(j-1)+1 < (j-1)^d \quad (26)$$

(26) is satisfied if $d+1 \leq 2^{d-1}$. Verify that the latter holds for $j \geq d \geq 3$. Suppose that m is even and $i \leq d(j-1)-1$. By equation (22), undeleted vertex u has an opposite as long as

$$d(j-1)-1 < (j-1)^d \quad (27)$$

which holds by Theorem 11. Suppose that m is even, $j \geq 3$, $d \leq 2$, and $d(j-1) \leq i \leq d(j-1)+1$. By equations (22) and (23), from u there remains at least one vertex at distance $d + \lfloor m/2 \rfloor$ or $d + \lfloor m/2 \rfloor - 1$ as long as

$$d(j-1)+1 < (2j + d - 2) \cdot (j-1)^{d-1} \quad (28)$$

Verify by substitution that (28) holds for $d = 0, 1$, and 2 . □

Theorem 17. Suppose that $j = 2$ and let H be any quorum induced by deleting i vertices from $K_{m,2}^d$, $0 \leq i \leq f = d+1$. If $i = 0$ then the radius of H equals $d + \lfloor m/2 \rfloor$. If m is odd and $i = 1$ then the radius of H is at least $d + \lfloor m/2 \rfloor$. Otherwise, the radius of H is at least $d + \lfloor m/2 \rfloor - 1$.

Proof. If $i = 0$ then $H = K_{m,2}^d$, with radius $d + \lfloor m/2 \rfloor$ by (21). If m is odd then, by (22), there remains at least one vertex opposite to any undeleted vertex u as long as the number of deleted vertices is less than 2. For the remaining cases note that (28) is satisfied when $j = 2$. □

Theorem 18. For positive integer $m \geq 3$ let H be any quorum induced by deleting i vertices from $K_2^1(2m+1)$, $0 \leq i \leq f = 2$. If $i = 0$ then the radius of H equals $1 + \lfloor m/2 \rfloor$. If m is odd and $i = 1$ then the radius of H is at least $1 + \lfloor m/2 \rfloor$. Otherwise, the radius of H is at least $\lfloor m/2 \rfloor$.

Proof. By the procedure on page 24, $K_2^1(2m+1)$ may be formed by inserting an $(m+1)^{\text{st}}$ vertex z'' between the m^{th} and zeroth vertices, x'' resp. w'' , in one of the cycles C'' of $K_{m,2}^d$, and connecting z'' to the m^{th} vertex x' in the other cycle C' . The distance from the edge (x'', z'') to any other vertex v of $K_2^1(2m+1)$ equals the distance from x'' to v in $K_{m,2}^d$. For any vertex $u \neq x''$, $u \neq z''$, the distance from u to other any vertex v of $K_2^1(2m+1)$ is at least as great as the respective distance in $K_{m,2}^d$. Therefore, the radius of any quorum of $K_2^1(2m+1)$ is at least as great as that of $K_{m,2}^d$. Equality follows at $i = 0$ by Theorem 14. □

For $j \geq 3$, $d \geq 3$, and $m \geq 4$, Theorems 12 and 16 imply that the radius of any quorum H induced by deleting i vertices from $K_{m,j}^d$ equals $d + \lfloor m/2 \rfloor$ whenever $0 \leq i \leq d(j-1)-1$. However, for $i = d(j-1)+1$ there is a gap of about $\lfloor m/2 \rfloor$ between the upper and lower bounds of Theorems 12 and 16. Let us narrow this gap.

Theorem 19. Suppose that $j \geq 3$ and let H be any quorum induced by deleting i vertices from $K_{m,j}^d$, $0 \leq i \leq d(j-1)+1 = f$. If $0 \leq i \leq d-1$ then the radius of H is at most $d + \lfloor m/2 \rfloor$. If $i = d$ then the radius of H is at most $d + \max(2, \lfloor m/2 \rfloor)$. If $d+1 \leq i \leq d(j-1)+1 = f$ then the radius of H is at most $d + \lfloor m/2 \rfloor + 1$.

Proof. Since $i \leq d(j-1)+1 = f$, any quorum has at least one basic cycle C' containing all of its original vertices. To see this note that





$$\lfloor d(j-1)+2 \rfloor / j^d \leq 1 \quad (29)$$

whenever $j \geq 3$ and $d \geq 1$ (by differentiation, the lefthand side of (29) decreases with increasing j and d). Let u' be any vertex of C' , and consider any other vertex v in H . If v is in C' then $\langle u, v \rangle \leq \lfloor m/2 \rfloor$ and the theorem holds. Otherwise, $v = v''$ resides in some other basic cycle C'' , some of whose vertices may have been deleted. In $K_{m,j}^d$, v'' has two neighbors w'' and z'' (one of which may be a counterpart u'' to u'). Let $v', w',$ and z' be the vertices of C' whose labels have the same low order digit as $v'', w'',$ and, respectively, z'' . Without loss of generality assume that the shortest path $P_{u', w'} \in C'$ between u' and w' is no longer than the shortest path $P_{u', z'} \in C'$ between u' and z' . By Theorem 8, there are $d(j-1)$ paths between v' and v'', w' and $w'',$ and z' and z'' , with the length of each path at most $d+1$. Moreover, these paths are pairwise interior-disjoint. For each of these three pair of vertices, Theorem 8 guarantees that d of the $d(j-1)$ paths have length no greater than d . For $0 \leq i \leq d-1$, we can always reach v'' from u' by taking a shortest path $P_{u', v'}$, of length at most $\lfloor m/2 \rfloor$, to v' , thence via one of the remaining paths between v' and v'' of length d . Thus, for $0 \leq i \leq d-1$, the radius of H is at most $d + \lfloor m/2 \rfloor$.

For $i \geq d$ note that neither $P_{u', w'}$ nor $P_{u', z'}$ has length greater than $\lfloor m/2 \rfloor$. Moreover, the length of $P_{u', w'}$ equals $\lfloor m/2 \rfloor$ if and only if $m = 3$ and $u' = v'$. The d^{th} vertex deleted may be w'' itself, or it may lie along one of d shortest interior-disjoint paths between v' and v'' ; however, these events are mutually exclusive. That is, either there is a path $P_{u', v'}$ thence to v'' , or there is a path $P_{u', w'}$, thence to w'' , thence (via a single edge) to v'' . In other words, for $i = d$ there remains between u' and v'' a path of length at most $d + \max(2, \lfloor m/2 \rfloor)$. In addition to the $d+1$ paths of length $d + \lfloor m/2 \rfloor$ or $d + \max(2, \lfloor m/2 \rfloor)$, there are $d(j-2)+1$ paths (including one that traverses $P_{u', z'}$) of length at most $d + \lfloor m/2 \rfloor + 1$. Therefore, for $d+2 \leq i \leq d(j-1)+1 = f$ the radius of H is at most $d + \lfloor m/2 \rfloor + 1$. \square

Theorem 19 tightens to $d + \lfloor m/2 \rfloor + 1$ the upper bound $d + m - 1$ obtained directly from Theorem 12 at $i = d(j-1)+1 = f$. Extending Theorem 13:

Theorem 20. Suppose that $j = 2$, $d \geq 2$ and let H be any quorum induced by deleting i vertices from $K_{m,2}^d$, $0 \leq i \leq d+2 = f$. If $0 \leq i \leq d-2$ then the radius of H is at most $d + \lfloor m/2 \rfloor$. If $i = d - 1$ then the radius of H is at most $d + \max(2, \lfloor m/2 \rfloor)$. If $d \leq i \leq d+1 = f$ then the radius of H is at most $d + \lfloor m/2 \rfloor + 1$.

Proof. Since $i \leq d+1 = f$, any quorum has at least one basic cycle C' containing all of its original vertices.

To see this note that
$$\lfloor d+2 \rfloor / 2^d \leq 1 \quad (30)$$

whenever $d \geq 2$ (by differentiation, the lefthand side of (30) decreases with increasing d). Let u' be any vertex of C' , and consider any other vertex v in H . If v is in C' then $\langle u, v \rangle \leq \lfloor m/2 \rfloor$ and the theorem holds. Otherwise, $v = v''$ resides in some other basic cycle C'' , some of whose vertices may have been deleted. In $K_{m,j}^d$, v'' has two neighbors w'' and z'' (one of which may be a counterpart u'' to u'). Let $v', w',$ and z' be the vertices of C' whose labels have the same low order digit as $v'', w'',$ and, respectively, z'' . Without loss of generality assume that the shortest path $P_{u', w'} \in C'$ between u' and w' is no longer than the shortest path $P_{u', z'} \in C'$ between u' and z' . By Theorem 9, there are d paths between v' and v'', w' and $w'',$ and z' and z'' , with the length of each path at most $d+1$. Moreover, these paths are pairwise interior-disjoint. For each of these three pair of vertices, Theorem 9 guarantees that $d-1$ of the d paths have length no greater than d . For $0 \leq i \leq d-2$, we can always reach v'' from u' by taking a shortest path $P_{u', v'}$, of length at most $\lfloor m/2 \rfloor$, to v' , thence via one of the remaining paths between v' and v'' of length d . Thus, for $0 \leq i \leq d-2$, the radius of H is at most $d + \lfloor m/2 \rfloor$.





For $i \geq d-1$ note that neither $P_{u', w'}$ nor $P_{u', z'}$ has length greater than $\lfloor m/2 \rfloor$. Moreover, the length of $P_{u', w'}$ equals $\lfloor m/2 \rfloor$ if and only if $m=3$ and $u'=v'$. The d^{th} vertex deleted may be w'' itself, or it may lie along one of d shortest interior-disjoint paths between v' and v'' ; however, these events are mutually exclusive. That is, either there is a path $P_{u', v'}$ thence to v'' , or there is a path $P_{u', w'}$, thence to w'' , thence (via a single edge) to v'' . In other words, for $i=d-1$ there remains between u' and v'' a path of length at most $d + \max(2, \lfloor m/2 \rfloor)$. In addition to the d paths of length $d + \lfloor m/2 \rfloor$ or $d + \max(2, \lfloor m/2 \rfloor)$, there are 2 paths (including one that traverses $P_{u', z'}$) of length at most $d + \lfloor m/2 \rfloor + 1$. Therefore, for $d \leq i \leq d+1 = f$ the radius of H is at most $d + \lfloor m/2 \rfloor + 1$. \square

Theorem 20 tightens to $d + \lfloor m/2 \rfloor + 1$ the upper bound $d + m - 1$ obtained directly from Theorem 13 at $i = d+1 = f$. Let us formulate analogous results for $j = 2, d = 1$. This case is relatively important since it pertains to what is arguably the most practical graph architectures for X2000 avionics (cf. Section 3.10).

Theorem 21. Let H be any quorum induced by deleting i vertices from $K_{m-2}^1, 0 \leq i \leq 2 = f$. The radius of H is at most $1 + \lfloor m/2 \rfloor$.

Proof. By Theorem 17, it suffices to consider the cases $1 \leq i \leq 2 = f$. Delete a single vertex u from the basic cycle C' that originally contains $u = u'$. For a root let $v = v''$ be any vertex in C'' (there may be another) that is opposite to u' . In the quorum formed by deleting u' , the distance from v'' to any other vertex is at most $1 + \lfloor m/2 \rfloor$ ($\lfloor m/2 \rfloor$ if m is even). This bound is preserved if we delete a second vertex $w = w'$ from C' , so it remains to consider the deletion of a second vertex $w = w''$ from C'' . In this case let the root $v = v'$ be a vertex in C' that (with respect to C') is opposite to u' . Between v' and any other vertex z' in C' there is a path of length at most $\lfloor m/2 \rfloor$ ($\lfloor m/2 \rfloor - 1$ if m is even) strictly within C' . With z'' the counterpart in C'' of z' in C' , there is a path via z' of length at most $1 + \lfloor m/2 \rfloor$ ($\lfloor m/2 \rfloor$ if m is even) between v' and z'' . This applies to every vertex of C'' , with the possible exception of $z'' = u'' \neq w''$, an undeleted vertex which (since u' is deleted), has no counterpart in C' . Vertex u'' originally has two neighbors in C'' , at most one of which w'' has been deleted, and at least one of which x'' remains undeleted. If m is even then in C' the neighbors of u' (one of which is the counterpart x' of x'') are distance $\lfloor m/2 \rfloor - 1$ from v' . Therefore, a (shortest) path in C' from v' to x' , followed by $(x', x'') \cup (x'', v'')$, has length $1 + \lfloor m/2 \rfloor$. If m is odd then x' lies at distance $\lfloor m/2 \rfloor - 1$ or $\lfloor m/2 \rfloor$ from v' . If the former then traverse the $1 + \lfloor m/2 \rfloor$ edges of the (shortest) path in C' from v' to x' , followed by $(x', x'') \cup (x'', v'')$. If the latter then, as shown in Figure 18, v' has a neighbor r' in C' that (with respect to C') is also opposite to u' (r' is also opposite to the counterpart w' of w''). Letting r' be the root reduces to the former case. In particular, traverse the $1 + \lfloor m/2 \rfloor$ edges of the (shortest) path in C' from r' to x' , followed by $(x', x'') \cup (x'', v'')$. \square




Figure 18: Illustration of the last case of the proof of Theorem 21.





Underlying K-cube $K_j^d(n)$	Number i of vertices deleted	Radius, as a function of the number i of vertices deleted, $0 \leq i \leq f = 1 + (j-1) \cdot \lfloor \log_j(n/m) \rfloor$		Number i of vertices deleted
		At least	At most	
radix $j = 2$, dimension $d = 1 = \log_2(n/m)$ number of vertices $n = 2m$ even	0, 1 (if m is odd)	equality by Equation (22), Theorems 17 and 21		0, 1 (if m is odd)
	1 (if m is even), 2	$\lfloor m/2 \rfloor$ Theorem 17	$1 + \lfloor m/2 \rfloor$ Theorem 21	1 (if m even), 2
radix $j = 2$, dimension $d = 1 = \lfloor \log_2(n/m) \rfloor$ number of vertices $n = 2m+1$ odd	0, 1 (if m is odd)	if $m = 2$ then 1; else $1 + \lfloor m/2 \rfloor$ equality for $m > 2$ by Theorems 18 and 22		0, 1 (if m is odd) 2 (if $m = 2$)
	1 (if m is even), 2	$\lfloor m/2 \rfloor$ Theorem 18	$1 + \lfloor m/2 \rfloor$ Theorem 22	1 (if m even)
			$1 + \lfloor (m+1)/2 \rfloor$ Theorem 22	2 (if $m > 2$)
radix $j = 2$, dimension $d = \log_2(n/m) \geq 2$ number of vertices $n = m \cdot 2^d$	0, 1 (if m is odd)	equality by Equation (22), Theorems 17 and 20		0, 1 (if m is odd)
	1 (if m is even), from 2 to $1 + \log_2(n/m)$	$\lfloor m/2 \rfloor - 1 + \log_2(n/m)$ Theorem 17	$\lfloor m/2 \rfloor + \log_2(n/m)$	1 (if m is even), from 2 to $\lfloor \log_2(n/m) \rfloor - 2$
			$\max(2, \lfloor m/2 \rfloor) + \log_2(n/m)$ Theorem 20	$\lfloor \log_2(n/m) \rfloor - 1$
			$1 + \lfloor m/2 \rfloor + \log_2(n/m)$ Theorem 20	$\log_2(n/m)$, $1 + \log_2(n/m)$
radix $j \geq 3$ dimension $d = \log_j(n/m)$ number of vertices $n = m \cdot j^d$	from 0 to $\lfloor \log_j(n/m) \rfloor - 1$	equality by Equation (22), Theorems 16 and 19		from 0 to $\lfloor \log_j(n/m) \rfloor - 1$
	from $\log_j(n/m)$ to $\lfloor (j-1) \cdot \log_j(n/m) \rfloor - 1$	$\lfloor m/2 \rfloor + \log_j(n/m)$ Theorem 16	if $d = \log_j(n/m) = 1$ then $1 + \lfloor m/2 \rfloor$ else $\max(2, \lfloor m/2 \rfloor) + \log_j(n/m)$ Theorems 19 and 27	$\log_j(n/m)$
	$(j-1) \cdot \log_j(n/m)$, $1 + (j-1) \cdot \log_j(n/m)$ (m odd)		if $d = \log_j(n/m) = 1$ then $1 + \lfloor m/2 \rfloor$ else $1 + \lfloor m/2 \rfloor + \log_j(n/m)$ Theorems 19 and 27	from $1 + \log_j(n/m)$ to $1 + (j-1) \cdot \log_j(n/m)$
	$(j-1) \cdot \log_j(n/m)$, $1 + (j-1) \cdot \log_j(n/m)$ (m even)	$\lfloor m/2 \rfloor - 1 + \log_j(n/m)$ Theorem 16		

Table 11: Radius of quorums induced from d -dimensional j -ary K-cube-connected cycles $K_j^d(n)$.

Theorem 22. Let H be any quorum induced by deleting i vertices from $K_2^1(2m+1)$, $0 \leq i \leq 2 = f$. If $i \leq 1$ then the radius of H is at most $1 + \lfloor m/2 \rfloor$. Otherwise, the radius is at most $1 + \lfloor (m+1)/2 \rfloor$.

Proof. By Theorem 18, it suffices to consider the cases $1 \leq i \leq 2 = f$. Suppose we delete a single vertex u from the basic cycle C' that originally contains m vertices, including $u = u'$. This splits C' , but leaves





intact the other basic cycle C'' . For a root let $v=v''$ be a vertex in C'' (there may another) whose counterpart in C' is, with respect to C' , opposite to u' . In the quorum formed by deleting u' , the distance from v'' to any other vertex z'' in C'' is at most $\lfloor (m+1)/2 \rfloor$. Otherwise, by traversing the shortest path in C'' from v'' to z'' , and thence the edge (v'', z') we can reach any vertex z'' in C' by a path whose length is at most $1 + \lfloor m/2 \rfloor$. These bounds on distance are preserved if we delete a second vertex from C' .

Suppose instead that we delete a single vertex $v = v''$ from the basic cycle C'' that originally contains $m+1$ vertices, including $u=u''$. For a root let $v = v'$ be the m^{th} vertex (whose label has low order digit $m-1$) in C' . Noting that v' has two neighbors in C'' , it follows that the distance from v' to any other vertex is at most $1 + \lfloor m/2 \rfloor$. Hence when $i = 1$ the radius of $K_2^1(2m+1)$ is at most $1 + \lfloor m/2 \rfloor$. These bounds on distance on distance are preserved if we delete a second vertex from C'' . It remains to consider the case where we delete one vertex u' from C' and one vertex w'' from C'' .

Let the root $v = v'$ be a vertex in C' that (with respect to C') is opposite to u' , and suppose that u' is the not the m^{th} vertex in C' (*i.e.*, prior to being deleted, u' has only one neighbor in C''). Between v' and any other vertex z' in C' , there is a path of length at most $\lfloor m/2 \rfloor$ ($\lfloor m/2 \rfloor - 1$ if m is even) strictly within C' . With z'' the counterpart in C'' of z' in C' , there is a path via z' of length at most $1 + \lfloor m/2 \rfloor$ ($\lfloor m/2 \rfloor$ if m is even) between v' and z'' . As in Theorem 21, this applies to every vertex of C'' , with the possible exception of $z'' = u'' \neq w''$, an undeleted vertex which (since u' is deleted), has no counterpart in C' . Vertex u'' has two neighbors x'', y'' in C'' , at least one of which x'' remains undeleted. If m is even then in C' the two neighbors of u' (one of which is the counterpart x' of x'') are distance $\lfloor m/2 \rfloor - 1$ from v' . Therefore, a (shortest) path in C' from v' to x' followed by $(x', x'') \cup (x'', v'')$ has length $1 + \lfloor m/2 \rfloor$. If m is odd then x' lies at distance $\lfloor m/2 \rfloor - 1$ or $\lfloor m/2 \rfloor$ from v' . If the former then traverse the $1 + \lfloor m/2 \rfloor$ edges of the (shortest) path in C' from v' to x' , followed by $(x', x'') \cup (x'', v'')$. If the latter then, as shown in Figure 18, v' has a neighbor r' in C' that (with respect to C') is also opposite to u' (r' is also opposite to the counterpart y' of y''). Letting r' be the root reduces to the former case. In particular, we traverse the $1 + \lfloor m/2 \rfloor$ edges of the (shortest) path in C' from r' to x' , followed by $(x', x'') \cup (x'', v'')$.

Finally, suppose that in C' we delete the m^{th} vertex u' , (*i.e.*, the low order digit on the label of u' equals $m-1$ and, prior to being deleted, u' has two neighbors in C''). Delete arbitrary vertex y'' in C'' , and let u' be a vertex in C' that (with respect to C') is opposite v' . Between v' , and any other vertex z' in C' , there is a path of length at most $\lfloor m/2 \rfloor$ ($\lfloor m/2 \rfloor - 1$ if m is even), strictly within C' . With z'' the counterpart in C'' of z' in C' , there is a path via z' of length at most $1 + \lfloor m/2 \rfloor$ ($\lfloor m/2 \rfloor$ if m is even) between v' and z'' . This applies to every vertex of C'' , with the exception of the m^{th} or $(m+1)^{\text{st}}$ vertices u'' *resp.* s'' , whenever one or both of u'' and s'' remain undeleted. If in C'' neither x'' , the predecessor of u'' , nor y'' , the successor of s'' , is undeleted then take as a root a vertex v' in C' that (with respect to C') is opposite to u' (*i.e.*, the low order digit of v' is $\lfloor m/2 \rfloor$ or $\lceil m/2 \rceil$). To reach u'' from v' follow a shortest path in C' to x' , traverse the edge from x' to its counterpart x'' , and trace the edge (x'', u'') , a total length of $1 + \lceil m/2 \rceil$ *resp.* $1 + \lfloor m/2 \rfloor$. Suppose on the other hand that either x'' or y'' is deleted, and without loss of generality assume that x'' remains in the quorum. Refer to Figure 19. Take as a root the vertex r' whose label has a low order digit equal to $1 + \lceil m/2 \rceil$. To reach u'' *resp.* s'' from v' follow a shortest path in C' to x' , traverse the edge from x' to its counterpart x'' , and trace the edge (x'', u'') , a total length of $\lfloor m/2 \rfloor$ *resp.* $1 + \lfloor (m+1)/2 \rfloor$. This leaves the successor to y'' at distance $1 + \lfloor (m+1)/2 \rfloor$ from v' , with all other vertices in the quorum at a distance from v' less than $1 + \lfloor (m+1)/2 \rfloor$. □






Figure 19: Illustration of the last case of the proof of Theorem 22.

Tables 11 and 12 synopsise our results for K-cube-connected cycles. With respect to radius, we see from Table 11 that the difference between our upper and lower bounds is typically one or zero, and in no case exceeds two. As with the case of K-cubes, we obtain lower and upper bounds on the value of $\rho(n, d(j-1)+2)$ by taking the maximum of the lower and upper bounds on the radius, as a function of the number of vertices deleted. For example, if $j \geq 3$, $d \geq 2$, and $m \geq 4$ then $0 \leq \rho(n, d(j-1)+2) - (d + \lfloor m/2 \rfloor) \leq 1$; i.e., our estimate is within one of the maximum radius. Independent of the choice of G , Theorem 6 tends to underestimate the value of ρ . In at least one case, however, the lower bound of Theorem 6 is exact: $\rho(6,2) \geq 2$, and this bound is achieved by $K_{m,2}^1$, shown in Figure 20.

Underlying K-cube $K_j^d(n)$	Number i of vertices deleted	Diameter, as a function of the number i of vertices deleted, $0 \leq i \leq f = 1 + (j-1) \cdot \lfloor \log_j(n/m) \rfloor$	
		At least	At most
radix $j = 2$, dimension $d = \log_2(n/m) = 1$ number of vertices $n = 2m+1$ odd	0	$1 + \lfloor m/2 \rfloor$, equality by Equation (22), Theorem 15	
	1	$1 + \lfloor m/2 \rfloor$ Theorem 15	$2 + \lfloor m/2 \rfloor$ Theorem 14
	2		$m+1$ Theorem 14
radix $j = 2$, dimension $d = \log_2(n/m)$ number of vertices $n = m \cdot 2^d$	0	$\lfloor m/2 \rfloor + \log_2(n/m)$, equality by Equation (22)	
	from 1 to $\lfloor \log_2(n/m) \rfloor - 1$	$\lfloor m/2 \rfloor + \log_2(n/m)$ Theorem 15	$\max(2, \lfloor m/2 \rfloor) + \max[2, \log_2(n/m)]$ Theorem 13
	$\log_2(n/m)$		$1 + \lfloor m/2 \rfloor + \log_2(n/m)$ Theorem 13
	$1 + \log_2(n/m)$		$m - 1 + \log_2(n/m)$ Theorem 13
radix $j \geq 3$, dimension $d = \log_j(n/m)$ number of vertices $n = m \cdot j^d$	0	$\lfloor m/2 \rfloor + \log_j(n/m)$, equality by Equation (22)	
	from 1 to $\log_j(n/m)$	$\lfloor m/2 \rfloor + \log_j(n/m)$ Theorem 15	$\max(2, \lfloor m/2 \rfloor) + \max[2, \log_j(n/m)]$ Theorem 12
	from $1 + \log_j(n/m)$ to $(j-1) \cdot \log_j(n/m)$		$1 + \lfloor m/2 \rfloor + \log_j(n/m)$ Theorem 12
$1 + (j-1) \cdot \log_j(n/m)$	$m - 1 + \log_j(n/m)$ Theorem 12		

Table 12: Diameter of quorums induced from d -dimensional j -ary K-cube-connected cycles $K_j^d(n)$.





3.5 Quorums from K-cube-connected Edges

This section complements the preceding by giving results for graphs whose structure lies between that of K-cubes and K-cube-connected cycles. Refer to Figure 20. A d -dimensional j -ary K-cube-connected edge of order n , denoted $K_{2,j}^d$, is the result of replacing each of the j^d vertices of K_j^d with an edge. For a basis, a zero-dimensional K-cube-connected edge $K_{2,j}^0$ is an edge connecting two vertices. The high order d digits of the label on a vertex u in edge h of $K_{2,j}^d$ are identical to the d digits on the label of vertex h of the corresponding K_j^d . The low order digit on u is its label in the corresponding $K_{2,j}^0$. Vertex u shares an edge with vertex v if and only if i) u and v are neighbors in a basic edge $K_{2,j}^0$, or ii) the low order digits of u and v are identical, and the high order digits differ in exactly one position. This definition gives rise to a development analogous to that of the Section 3.4. For example, the respective counterparts to (15) and (16) are

$$n/2 = j^d \tag{31}$$

and

$$d(j-1)+1 = f+1 \tag{32}$$

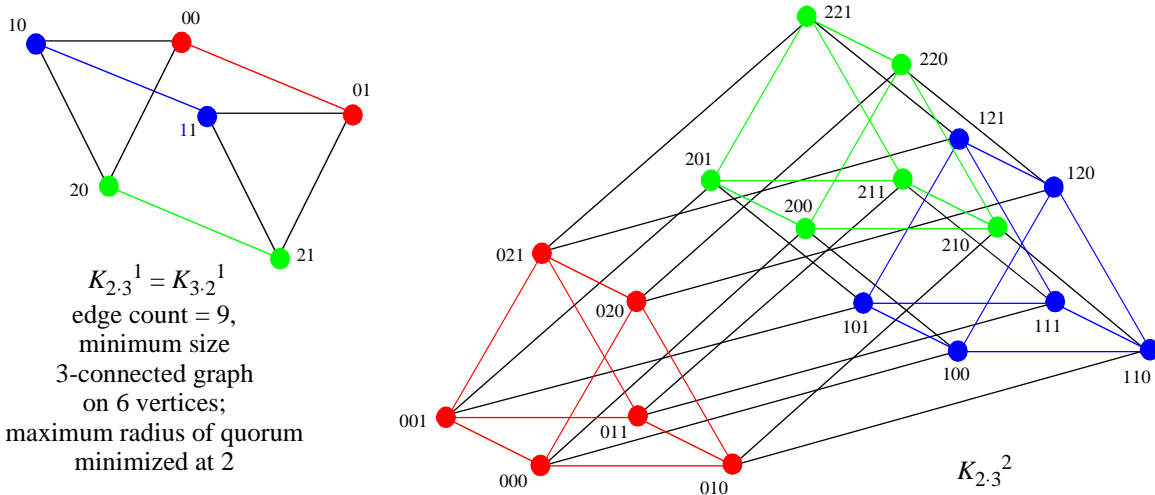


Figure 20: K-cube-connected edges $K_{2,j}^d$, radix $j = 3$. At $j = 2$ $K_{2,2}^d$ reduces to a binary K-cube K_2^{d+1} .

On the other hand, (19) pertains intact. Except for the case $n = 5$, therefore, a K-cube-connected graph with given connectivity and minimum count of edges structure cannot have as its basis a mixture of edges and cycles. It is for this reason that we have equality in (31), and are freed from having to consider analogs to Theorems 14, 18, and 22. If $n = 5$ then $j = 2$. When we delete $i = 0, 1, 2$ vertices from $K_2^1(5)$, the radius of the resulting quorum is at least 1, 2, resp. 1 and at most 2, 2, resp. 2. The minimum diameter of a spanning tree of a quorum of $K_2^1(5)$ always equals 2. For $j \geq 3$ we have a counterpart to Theorems 8 and 12:

Theorem 23. (Connectivity, upper bound on diameter.) If $j \geq 3$ then between vertices u and v in a $K_{2,j}^d$ there are $d(j-1)+1$ interior-disjoint paths, none of whose length exceeds $d + 2$. The length of $d+1$ of these paths is at most $d + 1$.

The proof of Theorem 23 is similar to those for Theorems 8 and 12; in the interest of shortening the exposition we omit the details. Note that the definition of a d -dimensional binary K-cube-connected edge coincides with that of a $(d+1)$ -dimensional binary K-cube. That is, $K_{2,2}^d = K_2^{d+1}$, and without loss of generality we may neglect K-cube-connected edges based on binary K-cubes. In particular, we are freed from having to consider counterparts to Theorems 13, 17, 20, and 21.





The recurrence of (8) pertains intact. That is, the recurrence relation for $B_j(d,i,2)$, the number of vertices at distance i from any vertex u in $K_{2,j}^d$, is identical that for K-cubes ($m = 1$) or K-cube-connected cycles ($m \geq 3$). By comparison to (8) or (20), the case $m = 2$ implies slightly different boundary conditions:

$$B_j(d,0,2) = 1, \quad B_j(0,1,2) = 1 \tag{33}$$

For $1 \leq i \leq d+1$, equation (21) provides the solution to the recurrence (8), with boundary conditions (33):

$$B_j(d,i,2) = B_j(d,i) + B_j(d,i-1) = (j-1)^i \binom{d}{i} + (j-1)^{i-1} \binom{d}{i-1} \tag{34}$$

where $\binom{d}{d+1} = 0$. At $j=2$ equation (34) reduces to the familiar $B_2(d+1,i) = \binom{d+1}{i}$. Table 13 illustrates.

$\downarrow d$	$j = 2$							$j = 3$							$j = 4$							
	$\rightarrow i$	0	1	2	3	4	5	6	0	1	2	3	4	5	6	0	1	2	3	4	5	6
0	1	1						1	1						1	1						
1	1	2	1					1	3	2					1	4	3					
2	1	3	3	1				1	5	8	4				1	7	15	9				
3	1	4	6	4	1			1	7	18	20	8			1	10	36	54	27			
4	1	5	10	10	5	1		1	9	32	56	48	16		1	13	66	162	189	81		
5	1	6	15	20	15	6	1	1	11	50	120	160	112	32	1	16	105	360	675	648	243	

Table 13: Number $B_j(d,i,2)$ of vertices at graph distance i from any other in a d -dimensional j -ary K-cube-connected edge $K_{2,j}^d$. The table may be verified or extended using (8) and (33), or (34).

Equation (34) also enables proofs of analogs to Theorems 10 and 11:

Theorem 24. Let H be any quorum induced by deleting i vertices from $K_{2,j}^d$, $0 \leq i \leq f = d(j-1)$. The diameter of H is at least $d+1$.

Theorem 25. Let H be any quorum induced by deleting i vertices from $K_{2,j}^d$, $0 \leq i \leq f = d(j-1)$, $j \geq 3$. The radius of H is at least $d+1$.

Table 14 summarizes our results for K-cube-connected edges.

Number i of vertices deleted, $0 \leq i \leq f$ $f = (j-1) \cdot \log_j(n/2)$	Radius		Diameter	
	At least	At most	At least	At most
from 0 to $\log_j(n/2)$	1 + $\log_j(n/2)$ Theorems 23 and 25		1 + $\log_j(n/2)$ Theorems 23 and 24	
from 1 + $\log_j(n/2)$ to $(j-1) \cdot \log_j(n/2)$	1 + $\log_j(n/2)$ Theorem 23	if $d = \log_j(n/2) = 1$ then 2 else 2 + $\log_j(n/2)$ Theorems 23 and 27	1 + $\log_j(n/2)$ Theorem 24	2 + $\log_j(n/2)$ Theorem 23

Table 14: Properties of quorums induced by deleting vertices from K-cube-connected edges $K_{2,j}^d$, $j \geq 3$.



3.6 Chordal Graphs and Cycles of K-cubes

This section by gives a partial characterization of the chordal graphs alluded to on page 9. Our results will be sufficient to motivate modification of these chordal graphs. The modified graphs turn out to be a cycles of cliques, a class of graphs that we have already studied.

Refer to Figure 21A. For odd $2q-1 = f \geq 3$, the chordal graphs $C_{n,2q}$ of [Hayes 1976] prescribe that each vertex is connected to $2q \leq n-1$ closest vertices along a cycle C_n . By Theorem 4 of [Hayes 1976], $C_{n,2q}$ is $2q$ -connected, and hence yields a quorum in the presence of any $2q-1$ partitioning faults. But what are the radius and diameter of such a quorum? As illustrated by Figure 21B, $C_{n,2q}$ can be modified in a fashion that appears to reduce the radius of an induced quorum. For the sake of clarity we call these modified structures *secant graphs*. For $n = m \cdot f$, a secant graph $C(m:f)$ is formed as follows. Divide the vertices into m classes, labeled from 0 to $m-1$. The index of a vertex's class is the high order digit on its label. Within each class, number the vertices from 0 to $f-1$. The index of a vertex within any class is the low order digit on its label. Connect vertex i to vertex $[i+1 \bmod n]$. Connect two vertices whenever their low order digit is the same. Let us confirm that, in fact, the maximum radius of a quorum induced by deleting up to f vertices of $C(m \cdot [2q-1])$ is less than that for $C_{m \cdot (2q-1), 2q}$, where $f = 2q-1$.




Figure 21: Chordal graph architecture $C_{n,2q}$ versus secant graph architecture $C(m:f)$; $n=9, q=2, f=3, m=3$.

First note that the definition of $C(m:f)$ coincides with that for $K_{m \cdot f}^1$; that is, $C(m:f)$ is a clique of cycles of length m . As drawn in Figure 21B, $C(m:f)$ appears to the eye as a (three) cycle of (three-vertex) cliques. No matter how we draw a graph, however, its adjacency remains unchanged. More generally, if we make m copies of a d -dimensional j -ary K-cube K_j^d , and then connect corresponding vertices of these K-cubes into a single cycle, we have the same result as if we had replaced every vertex in K_j^d by a C_m and connected the cycles according to the procedure on page 24. Let us record this observation as

Theorem 26. A cycle of m K-cubes K_j^d is identical to a K-cube of cycles $K_{m \cdot f}^d$. In particular, a secant graph $C(m:f)$ is a cycle of m cliques K_f .

To effectively compare $C_{m \cdot (2q-1), 2q}$ with $C(m \cdot [2q-1])$ we refine the result of Theorems 19 and 20 via what amounts to an extension of Theorem 5.

Theorem 27. For $m \geq 2$ and $0 \leq i \leq f = j-1 + \min(1, \lfloor (m-1)/2 \rfloor)$, let H be any quorum induced by deleting i vertices from $C(m:f) = K_{m \cdot f}^1$. The radius of H is at most $1 + \lfloor m/2 \rfloor$.

Proof. Suppose first that $m = 2$; that is, we have a quorum induced from a K-cube-connected edge. Since $i \leq j-1$, there remains at least one edge that spans the two j -ary cliques. Pick any vertex u from such an edge. To reach any other vertex v traverse at most one edge to reach the K_j (perhaps with one of more of its vertices deleted) to which v belongs. By Theorem 5, whatever remains of K_j affords an edge from the counterpart of u in v 's K_j to v . The length of the path between u and v is at most $2 = 1 + \lfloor m/2 \rfloor$.

Suppose on the other hand that $m \geq 3$; that is, we have a quorum induced from a K-cube-connected cycle. If there is a basic cycle from which a vertex has not been deleted then pick any vertex u from such a basic



cycle. To reach a destination vertex v traverse the basic cycle of which u is a member, in at most $\lfloor m/2 \rfloor$ edges, arriving at the j -ary clique K_j to which v belongs. By Theorem 5, whatever remains of K_j affords an edge that joins the counterpart of u in v 's K_j to v . The path between u and v has length at most $1 + \lfloor m/2 \rfloor$. If there is no basic cycle from which a vertex has not been deleted then $i = f = j$ and every basic cycle has exactly one vertex deleted from it. Pick any central vertex u from such a basic cycle. If the destination vertex v is in a clique other than that to which u belongs then assume that the counterpart of v in the clique of which u is a member has not been deleted. To reach v traverse the basic cycle of which u is a member, in at most $\lfloor m/2 \rfloor$ edges, arriving at the j -ary clique K_j to which v belongs. By Theorem 5, whatever remains of K_j affords an edge that joins the counterpart of u in v 's K_j to v . The path between u and v has length at most $1 + \lfloor m/2 \rfloor$. Finally, assume that v is in a clique other than that to which u belongs, and that the counterpart of v in the clique of which u is a member has been deleted. To reach v traverse the basic cycle of which u is a member, in at most $\lfloor m/2 \rfloor - 1$ edges arriving at w , the neighbor of v 's closest counterpart to u in this cycle (that we can do this is assured since the counterpart of v is the only vertex deleted from the cycle). If in v 's clique the counterpart to w has not been deleted then trace the edge from w to its counterpart, thence the edge to v . If in v 's clique the counterpart to w has been deleted then, as in Theorems 21 and 22, go from u to z , the other neighbor of the counterpart of v in the basic cycle of which u is a member, to the counterpart of w in v 's clique, to v . The length of this path between u and v is at most $1 + \lceil m/2 \rceil$, which equals $1 + \lfloor m/2 \rfloor$ if m is even. If m is odd then there is a second central vertex x in u 's basic cycle; x is closer by one edge to w . In this case the distance from x to any other vertex is at most $1 + \lfloor m/2 \rfloor$. \square

Theorem 28. In $C_{m \cdot (2q-1), 2q}$ the distance between vertex 0 and vertex i equals $\min(\lceil i/q \rceil, \lceil (n-i)/q \rceil)$.

Proof. Let u and v be the vertices in $C_{m \cdot (2q-1), 2q}$ whose respective labels are 0 and i . Assume, as in Figure 21A, that the vertices of $C_{m \cdot (2q-1), 2q}$ are labeled clockwise in ascending order. Any shortest path between u and v is either strictly clockwise or strictly counterclockwise. Therefore, any minimum length path from u to v is a minimal length path, in the sense described by [Bollabás, 1978], p. xvi. From u to v let P be a path of clockwise length i or counterclockwise length $n-i$ along the perimeter of $C_{m \cdot (2q-1), 2q}$. In a greedy fashion, replace a longest subpath of P with a single edge. Iterating on this procedure yields shortest clockwise and counterclockwise paths. Clockwise, we replace q edges with a single edge a total of $\lfloor i/q \rfloor$ times, and, if $i \neq 0 \pmod q$, substitute a single edge for $i \pmod q$ edges. Counterclockwise, we replace q edges with a single edge a total of $\lfloor (n-i)/q \rfloor$ times, and, if $n-i \neq 0 \pmod q$, substitute a single edge for $n-i \pmod q$ edges. Therefore the length of the shortest path between u and v equals $\min(\lceil i/q \rceil, \lceil (n-i)/q \rceil)$. \square

The distance prescribed in Theorem 28 is maximized when $i = n/2$ or, if n is odd, when $i = \lfloor n/2 \rfloor$ or $\lceil n/2 \rceil$. Since $\max_{0 < i < n} \min(\lceil i/q \rceil, \lceil (n-i)/q \rceil) = \min[\max_{0 < i < n}(\lceil i/q \rceil), \max_{0 < i < n}(\lceil (n-i)/q \rceil)]$, we have

Corollary 28.1. The radius of $C_{m \cdot (2q-1), 2q} = C_{n, 2q}$ is $\lceil \lfloor m \cdot (2q-1)/2 \rfloor / q \rceil = \lceil \lfloor n/2 \rfloor / q \rceil$.

Recall from the beginning of Section 3.1 that we are interested in minimizing the maximum radius of a quorum. By Theorem 27, the radius of $C(m \cdot \lfloor 2q-1 \rfloor)$, and of any quorum of $C(m \cdot \lfloor 2q-1 \rfloor)$, equals $1 + \lfloor m/2 \rfloor$. To establish that $C(m \cdot \lfloor 2q-1 \rfloor)$ is preferred to $C_{n, 2q}$ it therefore suffices to exhibit a quorum of $C_{m \cdot (2q-1), 2q}$ whose radius is greater than $1 + \lfloor m/2 \rfloor$. In particular, this applies in the case of zero faults, whence Corollary 28.1 pertains and the quorum is itself $C_{m \cdot (2q-1), 2q}$. When $m \geq 4$ and $q \geq 3$, we have

$$(q+1)/(q-1) \leq m/2 \quad (35)$$

$$\text{Multiply (35) by } q-1 \text{ and add } mq-1: \quad q(1 + \lfloor m/2 \rfloor) \leq q(1 + m/2) \leq m(q + 1/2) - 1 \leq \lfloor m \cdot (2q-1)/2 \rfloor \quad (36)$$

$$\text{Multiply (36) by } 1/q: \quad 1 + \lfloor m/2 \rfloor \leq \lfloor m \cdot (2q-1)/2 \rfloor / q \leq \lceil \lfloor m \cdot (2q-1)/2 \rfloor / q \rceil = \lceil \lfloor n/2 \rfloor / q \rceil \quad (37)$$

Thus when $m \geq 4$ and $q \geq 3$ the radius of $C(m \cdot \lfloor 2q-1 \rfloor)$ is no greater than that of $C_{m \cdot (2q-1), 2q}$. Similar manipulations reveal that (37) holds when $m \geq 6$ and $q = 2$, when $m = 3$ and $q \geq 2$, when $m = 4$ and $q = 2$, and when $m = 5$ and $q = 2$. Refined estimates and substitutions establish that the lefthand side of (37) is, in fact, *strictly* less than the righthand side whenever $m \geq 3$ and $q \geq 3$, or when $m \geq 6$ and $q = 2$. In summary:





Corollary 28.2. Over the range 0 ≤ i ≤ f = 2q - 1, the maximum radius of a quorum obtained by deleting i vertices from the secant graph C(m·[2q-1]) = K_{m,f}¹ never exceeds the maximum radius of a quorum obtained by deleting i vertices from the chordal graph C_{m·(2q-1),2q}. In particular, when i=0 and either m ≥ 3 and q ≥ 3 or m ≥ 6 and q = 2, the radius of C(m·[2q-1]) = K_{m,f}¹ is strictly less than that of C_{m·(2q-1),2q}.

In this section we have examined and contrasted chordal graphs along with secant graphs. A secant graph is constructible whenever the number of faults f divides the total number n of nodes and is, in fact, a one-dimensional K-cube-connected cycle. From the viewpoint of minimizing the maximum radius of quorum, a secant graph is at least as good, and generally better, than a chordal graph. We have not carried our comparison to the case of an even number of faults. Moreover, chordal graphs are constructible for any value of f < n,¹⁹ while secant graphs are constructible if and only if f divides n. Nevertheless, our analysis provides a basis for preferring K-cube-connected cycles to chordal graphs, and this is our recommendation.

3.7 Quorums from C-cubes

Often referred to in the literature as a "hypercube" or simply a "cube", a labeled d-dimensional j-ary C-cube C_j^d is constructed as follows.²⁰ For j = 2: C₂^d is a d-dimensional binary K-cube K₂^d (equivalently, a (d-1)-dimensional binary K-cube-connected edge K_{2,2}^{d-1}); for j = 4: C₄^d is a K₂^{2d} (proof by induction); binary cubes are characterized by Section 3.3. For j > 2: C_j⁰ is a single unlabeled vertex. C_j¹ is a cycle (Section 3.1) on j vertices, numbered circularly from 0 to j-1; two vertices are joined by an edge if and only if the modulo j difference in their labels equals ±1. Note that a one-dimensional j-ary C-cube C_j¹ is the same as a j-vertex zero-dimensional j-ary K-cube-connected cycle K_{j,j}⁰. In general, to construct C_j^d we i) make j copies of C_j^{d-1}; ii) prepend i to the label of each vertex of the ith copy of C_j^{d-1}; iii) connect with an edge vertices u and v (from different copies of C_j^{d-1}) if and only if the modulo j difference in the high order digits of the labels on u and v equals ±1, and the low order d-1 digits are identical. Alternatively, we can reserve d digits for the label on each vertex, thus giving to rise a construction that is independent of the order in which dimensions are populated. Figure 7 illustrates 4-ary and ternary C-cubes in 2 resp. 3 dimensions. Note that, since a cycle on three vertices is also a three-vertex clique, C₃^d = K₃^d (equivalently, a (d-1)-dimensional ternary K-cube-connected cycle K_{3,3}^{d-1}); these are characterized by Section 3.3. It suffices therefore to consider dimensions d ≥ 2 and radices j ≥ 5, and such is the focus of this section.




Figure 22: Labeling and connectivity for a C₄-cube and C₃-cube = K₃³ in two resp. three dimensions.

19. K-cube-connected cycles and chordal graphs of order n and having the least resp. greatest number of edges per vertex are the same as cycles resp. cliques. I.e., K_j⁰(n) = C_{n,2} = C_n resp. K_n¹(n) = C_{n,n-1} = K_n.

20. We use a "C" to preface the term for a cube C_j^d that is based on cycles, as opposed to a clique-based (K-)cube; with respect to the latter, the K derives (cf. Section 3.7) from notation for a j-vertex clique K_j.





As with K-cubes, it is useful to know salient properties of C-cubes. Some (but not all) of these properties are listed in [Zargham 1996] (p. 204). Recalling that the radix j is greater than four, let us establish results pertaining to these properties. By step (i) on the preceding page, C_j^d contains j copies of C_j^{d-1} ; therefore the order $n_C(d, j)$ of C_j^d equals $j \cdot n_C(d-1, j)$. Subject to the initial condition $n_C(0, j) = 1$, verify that the unique solution of this recurrence relation is the same as that (5) for the number of vertices in a j -ary K-cube:

$$n_C(d, j) = j^d \quad (38)$$

By step (iii) on the preceding page, the degree of a vertex in C_j^d equals its degree in C_j^{d-1} plus 2, the number of edges that connect it to vertices with the same labels in neighboring copies of C_j^{d-1} . Subject to the initial condition of zero edges in C_j^0 , the degree of each vertex in C_j^d is therefore $2d$ (39)

Summing (39) over all j^d vertices counts every edge twice. Hence the number $e_C(d, j)$ of edges in C_j^d is

$$e_C(d, j) = d \cdot j^d \quad (40)$$

As is the case with K-cubes (as well as edges and cycles of K-cubes), C-cubes are *vertex symmetric*.¹⁴ Moreover, and as illustrated by Figure 22, the vertices of C_j^d are in one-to-one correspondence with ordered d -tuples, each of whose coordinates is a nonnegative integer. This suggests that, if two vertices $u = (u_{d-1}, \dots, u_0)$ and $v = (v_{d-1}, \dots, v_0)$ are sufficiently close, their distance should be given by the L_1

metric (also known as the *city block*, or *Manhattan metric*): $\langle u, v \rangle_1 = \sum_{k=0}^{d-1} |u_k - v_k|$ (41)

This tendency is born out by the L_1 "modulo j " metric of (42). By analogy with Theorem 7:

Theorem 29. If u and v are vertices of C_j^d , labeled according to steps (i) – (iii) on page 44, then

$$\langle u, v \rangle_{\text{mod } j} = \sum_{k=0}^{d-1} \min(|u_k - v_k|, j - |u_k - v_k|) \quad (42)$$

Proof. Regard arbitrary vertices u and v in C_j^d . Since C_j^d is vertex symmetric, we can assume without loss of generality that $u = (0, \dots, 0) = \mathbf{0}$. By step (iii) on page 44, we must traverse at least $\min(v_k, j - v_k)$ edges along the i^{th} axis. Thus the distance from $\mathbf{0}$ to v is at least (42). Further, and again by the construction on page 44, this bound is achieved by traversing v_k edges in the positive direction of the i^{th} axis (if $v_k \leq j - v_k$) or (if $v_k > j - v_k$) by traversing $j - v_k$ edges in the negative direction of the i^{th} axis. \square

Equation (42) is maximized when the respective terms in the summation are maximized. That is, when $v_k = \lfloor j/2 \rfloor$, for all k ranging between 0 and $d - 1$. It immediately follows:

Corollary 29.1. The radius and diameter of C_j^d are identically $d \cdot \lfloor j/2 \rfloor$.

Corollary 29.1 addresses the case of a C_j^d without faults. To derive a lower bound on radius, consider the number $B_j^C(d, i)$ of integer lattice points on the surface of, as well as the total number $V_j^C(d, i)$ in, a closed ball of L_1 modulo j radius i . By Corollary 29.1 and equation (42), we know that $V_j^C(d, d \cdot \lfloor j/2 \rfloor) = j^d$ (43)

For the sake of visualization assume that j is odd; translate the labels of C_j^d so that the point $((j-1)/2, \dots, (j-1)/2)$ becomes the origin. By (42), any point v in the ball of interest belongs to an L_1 ball centered at the new origin, as long as all of the (translated) coordinates of v satisfy $v_k \leq (j-1)/2$. Let us establish the volume and surface area of such a ball. If the radius i equals 0 then the ball contains just the origin, which is





also on the surface in the sense that it is the number of points exact distance 0 from the center. Adopting the latter definition:

$$B_j^C(0,0) = V_j^C(d,0) = 1 \tag{44}$$

At the outset it is not clear what meaning we should accord the surface area of zero-dimensional ball with positive radius. However, if we hold strictly to the definition used for (44) then the surface area of a zero-dimensional ball equals zero whenever $i > 0$:

$$B_j^C(0, i > 0) = 0 \tag{45}$$

whence
$$V_j^C(0, i) = 1 \tag{46}$$

Refer to Figure 23. Equations (44), (45), and (46) are consistent with the one-dimensional case $B_j^C(0, i) = 2$ and $V_j^C(0, i) = 2i+1$ (which could have served as boundary conditions) as well as with the respective recurrences:

$$B_j^C(d,i) = B_j^C(d-1,i) + 2 \sum_{k \leq 0 \leq i-1} B_j^C(d-1,k) = B_j^C(d-1,i) + B_j^C(d-1,i-1) + B_j^C(d,i-1) \tag{47}$$

$$V_j^C(d,i) = V_j^C(d-1,i) + 2 \sum_{k \leq 0 \leq i-1} V_j^C(d-1,k) = V_j^C(d-1,i) + V_j^C(d-1,i-1) + V_j^C(d,i-1) \tag{48}$$

To obtain the righthand relation we have recursively applied the lefthand side to a split sum. Table 15 illustrates computation of B_j^C and V_j^C , analogous to that depicted by Tables 8, 10, and 13.

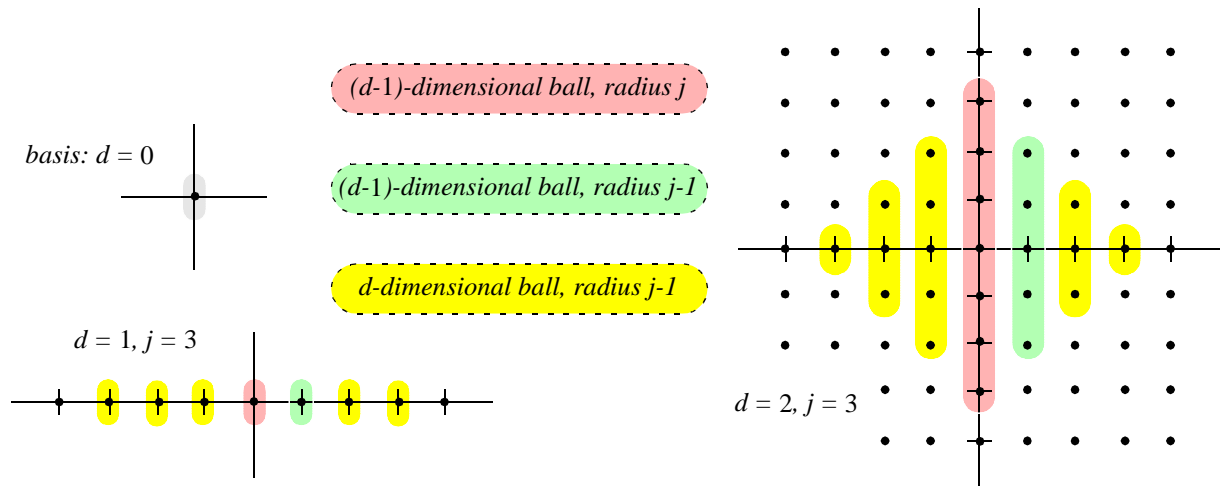


Figure 23: Balls in the L_1 metric: recursive composition and enumeration of volume and surface area.

Notice that the recurrence (47) for B_j^C is the same as that (48) for V_j^C , but boundary condition (45) for B_j^C differs from that (46) for V_j^C . As a result, and as illustrated in Table 15, B_j^C is asymmetric, while V_j^C is a symmetric function of d and j . Let us use combinatorial means to solve for V_j^C . Again we focus on balls centered at $u = \mathbf{0}$ in the translated coordinate system, and restrict the absolute value of each coordinate of v to a value no greater than $(j-1)/2$.

Consider the 2^d -tant comprising all strictly positive coordinates included in a ball of L_1 radius i . The number B_j^{C+} of positive integer lattice points on the surface of this ball equals the number of solutions to

$$i = \sum_{k=0}^{d-1} v_k \tag{49}$$





↓ d	$B_j^C(d,i)$								$V_j^C(d,i)$							
	0	1	2	3	4	5	6	7	0	1	2	3	4	5	6	7
0	1	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1
1	1	2	2	2	2	2	2	2	1	3	5	7	9	11	13	15
2	1	4	8	6	12	16	20	24	1	5	13	25	41	61	85	113
3	1	6	18	38	66	102	146	198	1	8	25	63	129	231	377	575
4	1	8	32	88	192	360	608	952	1	10	41	129	321	681	1289	2241
5	1	10	50	170	450	1002	1970	3530	1	12	61	231	681	1683	3653	7183
6	1	12	72	292	912	2364	5336	10836	1	12	85	377	1289	3653	8989	19825
7	1	14	98	462	1666	4942	12642	28814	1	12	113	575	2241	7183	19825	48639

Table 15: B_j^C and V_j^C count the number of vertices on the surface of, *resp.* included in, a closed ball encompassing integer lattice points, each of whose distance from the center is no greater than $\lfloor j/2 \rfloor = 7$. The ball has integer L_1 radius i , and is centered at a point whose coordinates correspond to a label in C_j^d .

Equation (49) has ordinary generating function:
$$\left[\frac{x}{1-x} \right]^d = \left[x \sum_{k=0}^{\infty} x^k \right]^d = x^d \sum_{k=1}^{\infty} \binom{d+k-1}{k} x^k \quad (50)$$

By Chapter 6 of [Tucker 1984], $B_j^{C^+}(d,i)$ is the coefficient of x^i in (50): $B_j^{C^+}(d,i) = \binom{i-1}{i-d} = \binom{i-1}{d-1}$ (51) where the righthand side makes use of the symmetry of binomial coefficients (*cf.* page 17). Summing over all i yields the volume of intersection of the ball with the strictly positive 2^d -tant:

$$\begin{aligned} V_j^{C^+}(d,i) &= \sum_{k=d}^i \binom{k-1}{d-1} = 0 + \binom{d-1}{d-1} + \sum_{k=d+1}^i \binom{k-1}{d-1} = \binom{d-1}{d} + \binom{d-1}{d-1} + \sum_{k=d+1}^i \binom{k-1}{d-1} \quad (52) \\ &= \binom{d}{d} + \binom{d}{d-1} + \sum_{k=d+2}^i \binom{k-1}{d-1} = \binom{d+1}{d} + \binom{d+1}{d-1} + \sum_{k=d+3}^i \binom{k-1}{d-1} = \dots \\ &= \binom{i-3}{d} + \binom{i-3}{d-1} + \sum_{k=i-1}^i \binom{k-1}{d-1} = \binom{i-2}{d} + \binom{i-2}{d-1} + \sum_{k=i}^i \binom{k-1}{d-1} = \binom{i-1}{d} + \binom{i-1}{d-1} = \binom{i}{d} \end{aligned}$$

The iterative simplification in (52) makes use of the recurrence relation (*cf.* page 17) for Pascal's triangle.

Again recalling that each coordinate is restricted to a value no greater than $(j-1)/2$, let us establish (51) and (52) by way of arguments which, unlike the preceding derivation, avoid generating functions and binomial identities. For $B_j^{C^+}(d,i)$, label i tally marks with the integers from 1 to i . Tag each of d tallies, in ascending order of tallies. Tagging the q^{th} tally with the k^{th} tag signifies that the value of the k^{th} coordinate equals the number of tallies after $(k-1)^{\text{st}}$ tag, up to, and including, the q^{th} tally. Note that there an implicit tag prior to the first tally, and that this construction assures that all coordinates are positive. For the sum of the coordinates to equal i , we must tag the i^{th} tally. This leaves $B_j^{C^+}(d,i) = \binom{i-1}{d-1}$ ways to distribute $d-1$ indistinguishable tags among $i-1$ distinguishable tallies. Since $V_j^{C^+}(d,i)$ corresponds to the case where the sum of the d coordinates is at most i , we are no longer required to tag the i^{th} tally. There are $\binom{i}{d}$ ways to distrib-





ute the d tags among the i tallies, and this is the number of positive integer vertices in a d -dimensional ball of L_1 radius i centered at the origin.

Write $V_j^{C^\pm}(d,i)$ and $B_j^{C^\pm}(d,i)$ for the number of vertices in *resp.* on a d -dimensional ball of L_1 radius i centered at the origin, such that no coordinate is zero. The number of ways of ordering d signs (plus or minus) equals 2^d ; each ordering corresponds to a 2^d -tant in d -dimensional space. In consequence,

$$V_j^{C^\pm}(d,i) = 2^d \binom{i}{d} \quad B_j^{C^\pm}(d,i) = 2^d \binom{i-1}{d-1} \quad (53)$$

For any k coordinates set to zero, we have $2^{d-k} \binom{i}{d-k}$ *resp.* $2^{d-k} \binom{i-1}{d-1-k}$ vertices in or on a d -dimensional ball of L_1 radius i centered at the origin. Since there are $\binom{d}{k}$ ways of setting k coordinates to zero, the volume is given by

$$V_j^C(d,i) = \sum_{k=0}^d 2^{d-k} \binom{d}{k} \binom{i}{d-k} = \sum_{k=0}^d 2^k \binom{d}{k} \binom{i}{k} = \sum_{k=0}^i 2^k \binom{d}{k} \binom{i}{k} \quad (54)$$

The righthand side of (54) explicates how is $V_j^C(d,i)$ is symmetric with respect to i and d . This is in accordance with boundary conditions (44) and (46), recurrence (48), and Table 15, but is to be contrasted with the asymmetric solution to (53):

$$B_j^C(d,i) = \sum_{k=0}^d 2^{d-k} \binom{d}{k} \binom{i-1}{d-1-k} = \sum_{k=1}^d 2^k \binom{d}{k} \binom{i-1}{k-1} = \sum_{k=1}^i 2^k \binom{d}{k} \binom{i-1}{k-1} \quad (55)$$

When the radius i exceeds $(j-1)/2$, a ball centered at the origin of C_j^d (translated) no longer includes all of the points encompassed by the analogous ball (of identical L_1 radius i) in the d -dimensional space of points whose coordinates are integers. For j odd, the ball of interest in C_j^d excludes those points having a coordinate whose absolute value exceeds $(j-1)/2$; analogous to (50), the ordinary generating function is

$$\left(\frac{x}{1-x}\right)^d \left(1-x^{\frac{j-1}{2}}\right)^d = \sum_{k=1}^{\infty} \binom{d+k-1}{k} x^{d+k} \sum_{q=1}^d \binom{d}{q} (-1)^q x^{\frac{(j-1)q}{2}} \quad (56)$$

wherein for $B_j^{C^+}(d,i)$ we extract the coefficient of x^i . Though somewhat more complicated, the case for j even is essentially similar. Rather than pursue this line, we focus on enumerating those points of interest: *i.e.*, those most distant, or most nearly distant, from any given vertex in C_j^d .

Consider points at maximum distance from the origin in an (untranslated) C_j^d , where j is even. Vertex v is maximally distant from the origin if and only if each of the terms in (42) equals $j/2$. This is possible if and only if each coordinate of v equals $j/2$. Thus $(j/2, \dots, j/2)$ is the unique point at maximum distance $dj/2$ from the origin:

$$B_j^C(d,dj/2) = 1 \quad j \text{ even} \quad (57)$$

Again for the case of j even, vertex v is distance $(dj/2)-1$ from the origin if and only if and only if $d-1$ terms in (42) equal $j/2$, and one term equals $(j/2)-1$. The coordinate corresponding to the term whose value equals $(j/2)-1$ has two possible values: $(j/2)-1$ and $(j/2)+1$. There are d ways of choosing this term, in which case the remaining $d-1$ terms are determined. Thus the points at distance one less than the maximum from the origin are those having $d-1$ coordinates equal to $j/2$ and one coordinate equal to $(j/2) \pm 1$:





$$B_j^C(d, \lfloor dj/2 \rfloor - 1) = 2d \quad j \text{ even} \tag{58}$$

Suppose that j is odd. Vertex v is maximally distant from the origin if and only if each of the terms in (42) equals $(j-1)/2$. Thus the points at maximum distance $d(j-1)/2$ from the origin have coordinates of the form $((j\pm 1)/2, \dots, (j\pm 1)/2)$. That is:

$$B_j^C(d, d(j-1)/2) = 2^d \quad j \text{ odd} \tag{59}$$

Let us apply the notion of opposite pairs to the case of C-cubes: u and v are *opposite* if their distance equals the diameter (alternatively, the radius) $d \lfloor j/2 \rfloor$ of C_j^d . Vertices u and v are *nearly opposite* if their distance is $d \lfloor j/2 \rfloor - 1$, one less than the diameter (alternatively, the radius) of C_j^d .

Theorem 30. Let H be any quorum induced by deleting i vertices from C_j^d , $0 \leq i \leq f = 2d-1$, $j \geq 5$. The diameter of H is at least $d \lfloor j/2 \rfloor$.

Proof. Suppose j is even. By (57), any given vertex u belongs to one opposite pair. Summing over all j^d vertices counts every pair of opposites twice, and the total number of opposite pairs equals $\frac{1}{2} j^d$. Each vertex we delete from C_j^d removes at most one opposite pair. Therefore, there remains at least one opposite pair as long as

$$4d \leq j^d \tag{60}$$

which follows by noting that $d \leq 2^{d-1} \leq 5^{d-1} \leq j^{d-1}$. Suppose that j is odd. By (57), any given vertex u belongs to 2^d opposite pairs. Summing over all j^d vertices counts every pair of opposites twice, and the total number of opposite pairs equals $2^{d-1} j^d$. Each vertex we delete from C_j^d removes at most 2^d opposite pairs. Therefore, there remains at least one opposite pair as long as

$$(2d-1)2^d < 2^{d-1} j^d \tag{61}$$

which reduces to (60). □

Theorem 31. Let H be any quorum induced by deleting i vertices from C_j^d , $0 \leq i \leq f = 2d-1$, $j \geq 5$. If $i = 0$ or j is odd then the radius of H is at least $d(j-1)/2$. For $i \geq 1$ and j is even, the radius of H is at least $(dj/2) - 1$.

Proof. The case $i = 0$ is covered by Corollary 29.1. Suppose that j is odd. By (59), undeleted vertex u has at least one opposite as long as

$$2d-1 < 2^d \tag{62}$$

which follows by remarks following (13). Suppose that j is even. By (58), there is at least one vertex nearly opposite to undeleted vertex u as long as

$$2d-1 < 2d \tag{63}$$

which follows since zero is less than one. □

path length →	1	2	3	4	5	6	7	8	9	10
000	001	002	003	013	023	033	133	233	333	
000	010	020	030	130	230	330	331	332	333	
000	100	200	300	301	302	303	313	323	333	
000	006	005	004	014	024	034	134	234	334	333
000	060	050	040	140	240	340	341	342	343	333
000	600	500	400	401	402	403	413	423	433	333

stage →	m		
path →	0	1	2
	1	2	0
	2	0	1

permutation matrix, cyclic group of order 3

Table 16: Illustration of Theorem 32: $2d = 6$ paths from the origin (0,0,0) to opposite (3,3,3) vertex in a three-dimensional 7-ary C-cube. Swingback paths are listed in the bottom three rows.





Theorem 32. (C-cube connectivity, upper bound on diameter, $j \geq 5$.) If v lies at distance $i > 0$ from vertex u of C_j^d then between u and v there is a set of $2d$ interior-disjoint paths. Let q be the number of coordinates where u and v are identical. i) $d-q$ of these paths $P(0) \dots P(d-q-1)$ have length i ; ii) $2q$ of these paths $P(d-q) \dots P(d+q-1)$ have length $i+2$. For $0 \leq r \leq d-q-1$, let c_r^+ denote the value of $\max(|u_k - v_k|, j - |u_k - v_k|)$ that is no larger than any set of $d-1-r$ other such c^+ 's, (cf. (64)) with the ordering ranging over $0 \leq k \leq d-1$. iii) Of the remaining $d-q$ paths $P(d+q) \dots P(2d-1)$, path $P(d+q+r)$ traverses $i+2c_r^+ - j$ edges.

Proof. By induction on d . As a basis take $d = 1$. Since C_j^d is vertex symmetric we can, without loss of generality, suppose that $u_0 = 0$ and $v_0 = i$. For property (i), trace from u to v a path $P(0)$ of minimum length i by traversing i edges along the cycle. Property (ii) holds since q is necessarily 0. For (iii), trace from u to v a path $P(1)$ in a direction opposite to, and interior-disjoint with, $P(0)$; note that $c_0^+ = j-i$, and that $P(1+0-1+1) = P(1)$ has length $j-i = i+2j-2i-j = i+2c_0^+ - j$. The theorem holds at $d = 1$.

Assume that the theorem holds in $0, \dots, (d-1)$ dimensions, and regard arbitrary vertices u and v in C_j^d , $d > 1$, $j \geq 4$. Suppose that $q = 0$; i.e., the coordinates of u and v differ in all d dimensions.

i) For the 0^{th} coordinate, trace a shortest path $P'(0)$, of length $\min(|u_0 - v_0|, j - |u_0 - v_0|)$, from u to (u_{d-1}, \dots, v_0) . By induction, the C_j^{d-1} prescribed by setting the 0^{th} coordinate to v_0 contains a path $P''(0)$ from (u_{d-1}, \dots, v_0) to v , and this path traverses $i - \min(|u_0 - v_0|, j - |u_0 - v_0|)$ edges. Catentating $P'(0)$ with $P''(0)$ gives an i -edge path $P(0)$ from u to v . For $h = 1, \dots, d-1$, iterate this process to synthesize path $P(h)$: at the start of the h^{th} iteration rotate each coordinate value by adding it to h , and converting the sum to its principal value mod d . As illustrated by the righthand side of Table 16, this completes a symmetric permutation matrix for the cyclic group of order d [Artin 1975] (VII:1.4). At $h = 2$, for example, coordinates along the path change in the order $2, \dots, d-1, 0, 1$. With respect to any vertex along a path, define the *stage* to be the number m of different coordinates that have changed; $q = 0$ implies $0 \leq m \leq d-1$. Entry (h, m) of the permutation matrix equals $(h+m) \bmod d$. Consider any two paths $P(h_1)$ and $P(h_2)$, for any stage $m < d-1$. Since entries 0 through m of any row map to successive elements of the cyclic group of order d , at least one of the values in columns 0 through m of row h_1 (resp. h_2) must not be in columns 0 through m of row h_2 (resp. h_1). But this means that, through stage m , the set of coordinates of $P(h_1)$ that are unchanged from their original values in u differ from the coordinates of $P(h_2)$ that are unchanged from their original values in u . Thus, the only possible intersection of $P(h_1)$ and $P(h_2)$ is at stage $d-1$. But this is also impossible: the $(h_1 + d-1 \bmod d-1)^{\text{th}}$ coordinate in $P(h_1)$ increments, in a monotone fashion modulo $d-1$, toward the coordinate value of v in that dimension, while the remaining paths have already attained the coordinate value of v in that dimension. Therefore, any path so constructed is interior-disjoint with any other.

iii) Continuing the case for $q = 0$, construct an additional d paths by substituting a *swingback* at the 0^{th} stage of the preceding procedure. For stages 0 through $d-1$, that is, begin by tracing a path $P'(d-1+h)$ of length $\max(|u_h - v_h|, j - |u_h - v_h|)$ from u to $(u_{d-1}, \dots, v_{h \pm 1 \bmod j}, \dots, u_0)$; if $\max(|u_h - v_h|, j - |u_h - v_h|) = j - |u_h - v_h|$ then the zeroth stage path stops at $v_h + 1 \bmod j$; otherwise it stops at $v_h - 1 \bmod j$.

This construction results in a swingback path $P(h)$ passing through a neighbor of v , with the h^{th} coordinate equal to $v_h \pm 1$. As illustrated by the bottom three rows of Table 16, the final step in the path traverses an edge to v . Note that the total length of $P(h)$ is $i + j - 2|u_h - v_h|$ if $\min(|u_h - v_h|, j - |u_h - v_h|) = |u_h - v_h|$; otherwise, $\min(|u_h - v_h|, j - |u_h - v_h|) = j - |u_h - v_h|$ and the path length is $i - j + 2|u_h - v_h|$. In any case, sorting the swingback paths by their lengths yields a set of $d-q = d-0 = d$ paths $P(d) \dots P(2d-1)$, with $P(d+r)$ traversing $i + 2c_r^+ - j$ edges, and $0 \leq r \leq d-q-1 = d-1$.





By an argument similar to that pertaining to paths without swingback, any path with swingback intersects no other path (with or without swingback), at least up to the next-to-last edge in the path. As remarked previously, the next-to-last edge advances to a unique neighbor of v (*i.e.*, one which has not been traversed by any other path, with or without swingback). For $q = 0$, that is, any two paths constructed in steps (i) or (iii) are interior-disjoint.

Now suppose that the integer q is positive. With u as source and v as destination, inductively apply the preceding procedure for $q = 0$ to the $d-q$ coordinates not shared by u and v . i) The C_j^{d-1} prescribed by the q coordinates whose values are the same in u and v contains $2(d-q)$ pairwise interior-disjoint $u-v$ paths, $d-q$ of which traverse i edges.

ii) Construct $2q$ bypass paths as follows. If k is the index of a coordinate such that $u_k = v_k$, then traverse to a neighbor of u by crossing one edge in the k^{th} dimension; *i.e.*, by incrementing or decrementing u_k . From this neighbor construct a path to the neighbor of v obtained by incrementing *resp.* decrementing the k^{th} coordinate of v . From u 's neighbor to v 's neighbor, a single path of length i is guaranteed by applying the procedure for $q = 0$ to the $d-q$ coordinates not shared by u and v . Traversing from v 's neighbor to v completes a path of length $i+2$. For each such k we obtain two paths (one by incrementing u_k and the other by decrementing u_k), with the k^{th} coordinate unique for every path so constructed. As a result, any bypass path is interior-disjoint with any other bypass path, as well as with any of the $2(d-q)$ paths (with or without swingback) whose vertex labels vary only in the coordinates not shared by u and v . The bypass procedure constructs $2q$ paths $P(d-q) \dots P(d+q-1)$ between u and v ; each bypass path traverses $i+2$ edges.

iii) By induction, the C_j^{d-1} prescribed by the q coordinates whose values are the same in u and v contains $d-q$ paths $P(d+q) \dots P(2d-1)$, pairwise interior-disjoint among themselves as well as with those constructed in steps (i) and (ii). Path $P(d+q+r)$ traverses $i+2c_r^+ - j$ edges. The theorem holds for $d > 1$. \square

Corollary 32.1. C_j^d is $2d$ -connected, and guarantees a quorum in the presence of any $2d-1$ faults.

Let us use our results to formulate upper bounds on quorum diameter at $i = 0, 1, \dots, 2d-1 = f$ faults. Since $i = 0$ is covered by Corollary 29.1, we focus on $1 \leq i \leq d-1$. Although q may assume any value in the range 0 to $d-1$, the distances of Theorem 32 attain a maximum only if $q = 0$; *i.e.*, for paths constructed according to procedure (i). To see this, and without loss of generality, note that any two opposites attain the diameter $d \lfloor j/2 \rfloor$ with $i = 0$. By contrast, the source and destination of a type (ii) bypass path must be identical in at least one of the coordinates. Therefore, any path constructed according to procedure (ii) has length at most $(d-1) \lfloor j/2 \rfloor + 2 \leq d \lfloor j/2 \rfloor$, where the latter follows since $j \geq 5$. For values $1 \leq i \leq d-1$, where paths of type (i) or type (ii) apply, it is the type (i) paths which realize the greatest number $d \lfloor j/2 \rfloor$ of edges.

For a number i of faults in the range $d \leq i \leq 2d-1$, consider the length of paths constructed by procedure (iii), with $q = 0$. For $0 \leq r \leq d-q-1$, define c_r as the value $j - c_r^+$; that is, c_r is the $(r+1)^{\text{st}}$ greatest addend in $\langle u, v \rangle$, the distance (42). Since $c_0^+ \leq \dots \leq c_r^+ \leq \dots \leq c_{d-1}^+$, it follows that $c_0 \geq \dots \geq c_r \geq \dots \geq c_{d-1}$ (64)

Writing $\langle P \rangle$ for the length of path P , express the length of the paths constructed by step (iii) as:

$$\langle P(d+r) \rangle = \sum_{k=0}^{r-1} c_k + (j - c_r) + \sum_{k=r+1}^{d-1} c_k \quad (65)$$

Consistent with (64), and by the remark preceding Corollary 29.1, the righthand side of (65) is at most

$$rc_0 + (j - c_r) + (d - r - 1)c_{r+1} \leq r \lfloor j/2 \rfloor + (j - c_r) + (d - r - 1)c_{r+1} \quad (66)$$

If $r < d - 1$ then the righthand side of (66) is bounded from above by

$$r \lfloor j/2 \rfloor + (d - r - 2)c_r \leq (d - 1) \lfloor j/2 \rfloor + \lceil j/2 \rceil \quad (67)$$

If $r = d - 1$ then the righthand side of (66) is at most $(d - 1) \lfloor j/2 \rfloor + j - c_{d-1} \leq (d - 1) \lfloor j/2 \rfloor + j - 1$ (68)





To complete our analysis of these paths, note that the righthand side of (67) is achieved between any vertices u and v , all d of whose coordinates differ by an absolute value of $\lfloor j/2 \rfloor$ or $j - \lfloor j/2 \rfloor$; for illustration: $u = \mathbf{0}$, $v = (\lfloor j/2 \rfloor, \dots, \lfloor j/2 \rfloor)$. Further, the righthand side of (68) is achieved between any vertices u and v , $d-1$ of whose coordinates differ by an absolute value of $\lfloor j/2 \rfloor$ or $j - \lfloor j/2 \rfloor$, and one of whose coordinates differs by ± 1 ; for illustration: $u = \mathbf{0}$, $v = (\lfloor j/2 \rfloor, \dots, \lfloor j/2 \rfloor, 1)$. Furthermore, and by remarks following Corollary 32.1, these pathlengths exceed those of paths constructed by procedure (ii). In summary:

Corollary 32.2. Let H be any quorum induced by deleting i vertices from C_j^d , $0 \leq i \leq f = 2d-1$, $j \geq 5$, $d \geq 2$. If $i \leq d-1$ then the diameter of H is at most $d \lfloor j/2 \rfloor$. If $d \leq i \leq 2d-2$ then the diameter of H is at most $(d-1) \lfloor j/2 \rfloor + \lceil j/2 \rceil$. If $i = 2d-1$ then the diameter of H is at most $d \lfloor j/2 \rfloor + \lceil j/2 \rceil - 1$.

Number i of vertices deleted, $0 \leq i \leq f$ $f = 2 \cdot \lceil \log_j n \rceil - 1$	Radius		Diameter	
	At least	At most	At least	At most
0	$\lfloor j/2 \rfloor \cdot \log_j n$ Corollary 29.1			
from 1 to $\lceil \log_j n \rceil - 1$	if j is odd then $\frac{1}{2} \cdot (j-1) \cdot \log_j n$ else $\frac{1}{2} \cdot j \cdot \lceil \log_j n \rceil - 1$ Theorem 31	$\lfloor j/2 \rfloor \cdot \log_j n$ Theorem 30, Corollary 32.2		
from $\lceil \log_j n \rceil$ to $2 \cdot \lceil \log_j n \rceil - 2$		$\lfloor j/2 \rfloor \cdot (\lceil \log_j n \rceil - 1) + \lceil j/2 \rceil$ Corollary 32.2	$\lfloor j/2 \rfloor \cdot \log_j n$ Theorem 30	$\lfloor j/2 \rfloor \cdot (\lceil \log_j n \rceil - 1) + \lceil j/2 \rceil$ Corollary 32.2
$2 \cdot \lceil \log_j n \rceil - 1$		$\lfloor j/2 \rfloor \cdot (\log_j n) + \lceil j/2 \rceil - 1$ Corollary 32.2		$\lfloor j/2 \rfloor \cdot (\log_j n) + \lceil j/2 \rceil - 1$ Corollary 32.2

Table 17: Properties of quorums induced by deleting vertices from C-cubes C_j^d , $j \geq 5$, $d \geq 2$.

3.8 Choosing a Graph Architecture

Sections 3.1 through 3.7 provide a taxonomy that includes minimum size graph architectures whose quorum radii are, a technical sense, optimum. In question is how to choose from among these architectures, so as to minimize the maximum radius of any quorum. Refer to Table 18. At either end of the range of fault tolerance our choice is both prescribed (stars, cycles, and cliques) and optimum. Between these extremes we may choose from regular graphs $K_{m,j}^d$ whose parameters m , j , and d , are independent: d -dimensional j -ary K-cubes ($m = 1$), K-cube-connected edges ($m = 2$), and K-cube-connected cycles ($m \geq 3$).

For given fault tolerance f , what values of $n = m \cdot j^d$ allow us to build a $K_{m,j}^d$? At $m = 1$ we take all ordered pairs (d, j) of positive integers such that $d \cdot (j-1) = f + 1$. Each distinct (j, d) determines a K-cube K_j^d , with $n = j^d$ and quorum radius at most $d + 1$. Similarly, for $m = 2$ we take all ordered pairs (j, d) such that $d \cdot (j-1) = f$. Each distinct (j, d) determines a K-cube-connected edge $K_{2,j}^d$, with $n = 2 \cdot j^d$ and quorum radius at most $d + 2$. For $m \geq 3$ we take all ordered triples (m, j, d) such that $d \cdot (j-1) = f - 1$. Each distinct (m, j, d) determines a K-cube-connected cycle $K_{m,j}^d$, with $n = m \cdot j^d$ and quorum radius at most $d + 1 + \lfloor m/2 \rfloor$. Since minimizing the dimension d is equivalent to maximizing the radix j , for fixed value of m we minimize the quorum radius by letting j take on the greatest possible value. Figure 24 demonstrates how, for fixed values of n and f , it is possible to have both a K-cube and a K-cube-connected cycle. Figure 25 shows how we may also have both a K-cube-connected edge and a K-cube-connected cycle. Except for the trivial case $K_{2,2}^0 = K_2^1$ ($n = 2$, $f = 0$), we know of no values of n and f for which a K-cube and a K-cube-connected edge may exist simultaneously.





Figures 24 and 25 also illustrate how we have embodied, in *executable* form, the theorems and corollaries of Sections 3.1 through 3.7. GRAFT (GRAph Architecture Fault Tolerance) calculates both graph architectures and their salient properties. GRAFT is implemented as a Microsoft Excel workbook, and accompanies this report as file GRAFT.xls. The main worksheet summarizes the quorum radius by taking the maximum of our lower and upper bounds on $\rho(n,i)$, as the number i of faults ranges between 0 and f . Underlying the summary are detailed worksheets for stars, cycles, cliques, K-cubes, K-cube-connected cycles, K-cube-connected edges, and C-cubes. As a function of the number of vertices deleted, the underlying worksheets give bounds on the radius of the quorum induced; by Theorem 2, these bounds apply as well to the radius of a tree that spans the quorum. The underlying worksheets also detail lower and upper bounds on the quorum diameter, as well as the minimum diameter of a tree spanning the quorum. Figures 26 and 27 illustrate detailed worksheets corresponding to the summary of Figure 25, at $(n, f) = (54, 6)$.

GRAFT: GRaph Architecture Fault Tolerance Calculator, Version 2.0. Computes n -node f -fault tolerant graph architectures having minimum number of point-to-point connections, bounded radius ρ and diameter.			Copyright 1999 by Laurence E. LaForge, NASA/ASEE Summer Faculty Fellow. 10-Oct-1998, 18-Oct-1999. Reprint rights granted to NASA and to the ASEE for research and educational purposes. Based on theory developed in my report: <i>Fault Tolerant Physical Interconnection of X2000 Computational Avionics</i> .	
Input:	n = number of nodes 64	f = maximum number of partitioning faults 8	e = minimum number of point-to-point connections: 288	Average number of point-to-point connections per node (number of ports per node) 9.00
	Feasible graph architecture(s) with minimum number of point-to-point connections:		Graph radius $\rho(n,f)$ of quorum and of tree spanning the quorum	
			At least	At most
Recommended:	3-dimensional 4-ary K-cube		3	4
Feasible, but not recommended:	1-dimensional 8-ary K-cube-connected cycle with 8 cycles, each containing 8 vertices		5	6

Figure 24: GRAFT’s main worksheet summarizes properties of feasible graph architectures.

GRAFT: GRaph Architecture Fault Tolerance Calculator, Version 2.0. Computes n -node f -fault tolerant graph architectures having minimum number of point-to-point connections, bounded radius ρ and diameter.			Copyright 1999 by Laurence E. LaForge, NASA/ASEE Summer Faculty Fellow. 10-Oct-1998, 18-Oct-1999. Reprint rights granted to NASA and to the ASEE for research and educational purposes. Based on theory developed in my report: <i>Fault Tolerant Physical Interconnection of X2000 Computational Avionics</i> .	
Input:	n = number of nodes 54	f = maximum number of partitioning faults 6	e = minimum number of point-to-point connections: 189	Average number of point-to-point connections per node (number of ports per node) 7.00
	Feasible graph architecture(s) with minimum number of point-to-point connections:		Graph radius $\rho(n,f)$ of quorum and of tree spanning the quorum	
			At least	At most
Recommended:	3-dimensional 3-ary K-cube-connected edge		4	5
Feasible, but not recommended:	1-dimensional 6-ary K-cube-connected cycle with 6 cycles, each containing 9 vertices		5	6

Figure 25: GRAFT facilitates exploration of alternative fault tolerant graph architectures.

GRAFT ameliorates the burden of remembering and applying the bulk of the more than 30 theorems and corollaries contained in this report. The designer uses GRAFT by adjusting the values of n and f . As Figure 26 shows, if no candidate architecture is feasible then GRAFT displays instructions summarizing relations between n and f that achieve one of the candidate architectures. As a practical matter, the designer





can often arrive at feasible values for n and f by simply trying different combinations. Section 3.9 illustrates this process by applying it to the proposed adjacency for X2000 core avionics.

K-cube-connected edge?	i = number of partitioning faults $\leq f$	Lower bound on radius $\rho(n,i)$, in general independent of graph architecture	Radius of quorum and of tree spanning quorum, as a function of i		Diameter of quorum, as a function of i		Minimum diameter of spanning tree, as a function of i	
			At least	At most	At least	At most	At least	At most
TRUE								
Structure:	0	3.00	4	4	4	4	7	8
3-dimensional 3-ary K-cube-connected edge	1	3.00	4	4	4	4	7	8
	2	3.00	4	4	4	4	7	8
	3	3.00	4	4	4	4	7	8
	4	2.00	4	5	4	5	7	10
	5	2.00	4	5	4	5	7	10
	6	2.00	4	5	4	5	7	10

Figure 26: Detailed worksheet corresponding to recommended architecture of Figure 25.

K-cube-connected Cycle?	i = number of partitioning faults $\leq f$	Lower bound on radius $\rho(n,i)$, in general independent of graph architecture	Radius of quorum and of tree spanning quorum, as a function of i		Diameter of quorum, as a function of i		Minimum diameter of spanning tree, as a function of i	
			At least	At most	At least	At most	At least	At most
TRUE								
Structure:	0	3.00	5	5	5	5	9	10
1-dimensional 6-ary K-cube-connected cycle with 6 cycles, each containing 9 vertices	1	3.00	5	5	5	4	9	10
	2	3.00	5	6	5	6	9	12
	3	3.00	5	6	5	6	9	12
	4	2.00	5	6	5	6	9	12
	5	2.00	5	6	5	6	9	12
	6	2.00	5	6	5	9	9	12

Figure 27: Detailed worksheet corresponding to architecture of Figure 25. Feasible, but not recommended.

We conclude this section by addressing the theoretical optimality of K-cubes, K-cube-connected edges, K-cube-connected cycles, and C-cubes. In a ratioed asymptotic sense, the K-cube constructions can deliver the best possible value $\Theta(\log n)$ of $\rho(n, f)$; *i.e.*, a quorum radius that, within a constant factor (perhaps equal to one) matches the lower bounds of Theorem 6. Moreover, K-cubes and their relatives are preferred to C-cubes for two reasons: 1) the radius of a C-cube quorum exceeds the diameter of the comparable K-cube having identical fault tolerance; 2) there is *no* relation between j and d such that, as $n_C = j^d \rightarrow \infty$, the ratio of the C-cube quorum radius to the general lower bound of Theorem 6, does *not* diverge; *i.e.*, this ratio must approach infinity. With respect to both criteria, that is, C-cubes are sub-optimal.

For real x , the sign of x is indicated by the function $\text{signum}(x)$. If $x > 0$ then $\text{signum}(x) = 1$; if $x < 0$ then $\text{signum}(x) = -1$; if $x = 0$ then $\text{signum}(x) = 0$. Refer to Table 18. The signum function allows us to conveniently encapsulate the fault tolerance of $K_{m,j}^d$ as

$$f = (j-1) \cdot d + \text{signum}(m-2) \tag{69}$$





GRAFT: Graph Architecture Fault Tolerance Calculator, Version 2.0. Computes n -node f -fault tolerant graph architectures having minimum number of point-to-point connections, bounded radius ρ and diameter.			Copyright 1999 by Laurence E. LaForge, NASA/ASEE Summer Faculty Fellow, 10-Oct-1998, 18-Oct-1999. Reprint rights granted to NASA and to the ASEE for research and educational purposes. Based on theory developed in my report: <i>Fault Tolerant Physical Interconnection of X2000 Computational Avionics</i> .																	
Input:	n = number of nodes 54	f = maximum number of partitioning faults 7	e = minimum number of point-to-point connections:	Average number of point-to-point connections per node (number of ports per node)																
Adjust n or f.	No feasible architecture computed.			Graph radius $\rho(n, f)$ of quorum and of tree spanning the quorum																
				<table border="1"> <tr> <th>At least</th> <th>At most</th> </tr> <tr> <td>not feasible</td> <td>with given n and f</td> </tr> <tr> <td>not feasible</td> <td>with given n and f</td> </tr> <tr> <td>not feasible</td> <td>with given n and f</td> </tr> <tr> <td>not feasible</td> <td>with given n and f</td> </tr> <tr> <td>not feasible</td> <td>with given n and f</td> </tr> <tr> <td>not feasible</td> <td>with given n and f</td> </tr> <tr> <td>not feasible</td> <td>with given n and f</td> </tr> </table>	At least	At most	not feasible	with given n and f	not feasible	with given n and f	not feasible	with given n and f	not feasible	with given n and f	not feasible	with given n and f	not feasible	with given n and f	not feasible	with given n and f
At least	At most																			
not feasible	with given n and f																			
not feasible	with given n and f																			
not feasible	with given n and f																			
not feasible	with given n and f																			
not feasible	with given n and f																			
not feasible	with given n and f																			
not feasible	with given n and f																			
For a star:	make $f = 0$																			
For a cycle:	make $f = 1$																			
For a clique:	make $n = f+1$ or $n = f+2$																			
For a K-cube:	make $n = [(f-1)/d-1]^d$, for positive integer d																			
For a K-cube-connected edge:	make $n = 2[f/d+1]^d$, for positive integer $d \leq f/2$																			
For a K-cube-connected cycle:	make $n = m [(f-1)/d+1]^d$, for integers $d \geq 1, f \geq 3$, and $m \geq 3$, or choose $f = 2$ and any integer $n \geq 5$																			
For a C-cube:	make $n = j^{(f+1)/2}$, for integers $j \geq 5$ and odd $f \geq 3$																			

Figure 28: GRAFT offers suggestions whenever it cannot construct a minimum size graph architecture.

Theorem 33. Denote by $\rho_{\text{Thm 6}}^-$ the lower bound on the radius of any quorum, as given by Theorem 6. If

$$\rho_{m,j,d}^- = \log_j(n/m) + \lfloor m/2 \rfloor \quad \text{and} \quad \rho_{m,j,d}^+ = 1 + \log_j(n/m) + \lfloor m/2 \rfloor \quad (70)$$

are the minimum *resp.* maximum radius of quorums of $K_{m,j}^d$, as listed in Table 18, then

$$\frac{\left[d + \left\lfloor \frac{m}{2} \right\rfloor \right] [\ln(j-2) + \ln d]}{\ln m + d \ln j} \leq \frac{\rho_{m,j,d}^-}{\rho_{\text{Thm 6}}^-} \leq \frac{\rho_{m,j,d}^+}{\rho_{\text{Thm 6}}^-} \leq \frac{\left[d + \left\lfloor \frac{m}{2} \right\rfloor + 1 \right] [\ln j + \ln d]}{\ln m + d \ln j - 1.4} \quad (71)$$

Proof. Explicate $\rho_{m,j,d}^-, \rho_{m,j,d}^+$, and $\rho_{\text{Thm 6}}^-$, the latter without the ceiling function. Making use of (69), substitute $j = 1 + \lfloor f - \text{signum}(m-2) \rfloor / d$. For the lower bound invoke the inequalities $-1 \leq \text{signum}(m-2)$, $n(f-1)+3 \leq nf$, and $\ln[(f-1)/(f+2)] < 0$. For the upper bound observe that $\text{signum}(m-2) \leq 1$, $(j-1)d+2 \leq jd$, and $-1.4 < \ln[(f-1)/(f+2)]$. The result follows by algebraic manipulation. \square

It is interesting to note that, in the large, the fault tolerance (69) of $K_{m,j}^d$ is dominated by j and d , and grows in a fashion that is independent of m . By contrast, the radius of $K_{m,j}^d$ is dominated by m and d , and is independent of j . Our conclusions about the optimality of the quorum radius of $K_{m,j}^d$ depend on how m, j , and d tend to infinity. If the left and right sides of (71) tend to some limit λ then, in the large, $\rho_{m,j,d}^-, \rho_{m,j,d}^+$, and $\rho_{\text{Thm 6}}^-$ are within a factor λ of $\rho(m \cdot j^d, (j-1) \cdot d + \text{signum}[m-2])$, the minimum value (over all graphs) of the maximum quorum radius. Abbreviating the latter quantity as $\rho_{m,j,d}$, we obtain the following result.

Corollary 33.1. If, for all $n_K = m \cdot j^d \geq k$, q and r are least upper bounds such that $\left\lfloor \frac{m}{2} \right\rfloor + 1 \leq qd$ and $\ln d \leq r \ln j$, then $\rho_{m,j,d}^+ \leq (\rho_{m,j,d}) \cdot (1 + q + qr + r)$.

Under the conditions of Corollary 33.1, that is, the maximum quorum radius of $K_{m,j}^d$ approaches a value that is within a factor $1 + q + qr + r$ of the minimum. Several special cases of Corollary 33.1 are of particular interest: a) $d \in o(j)$; b) $m \in o(d)$; c) both (a) and (b). In this instance the maximum radius of quorums induced from $K_{m,j}^d$ is asymptotically within a factor a) $1 + q$, b) $1 + r$, or c) 1 of $\rho_{m,j,d}$.





Fault tolerance f	Graph architectures	Maximum of quorum radii $\rho(n, i), 0 \leq i \leq f$		Maximum radius of quorum divided by lower bound $\rho^-_{\text{Thm 6}}$	References
		At least	At most		
0	$\mathcal{G}_{n,0}$ uniquely the set of n -vertex stars S_n	1		Exactly best possible	Table 7
1	$\mathcal{G}_{n,1}$ uniquely the set of n -vertex cycles C_n	$\lfloor n/2 \rfloor$		Exactly best possible	Table 7
2	$\mathcal{G}_{n,2}^+$ includes 1-dimensional binary K-cube-connected cycles $K_2^1(n = 2m+1), m \geq 2$	1 if $n = 5$ else $1 + \lfloor \lfloor n/2 \rfloor / 2 \rfloor$	$1 + \lfloor \lfloor n/2 \rfloor / 2 \rfloor$	Don't know	Table 11, Theorem 6, discussion on p. 40
$2 \cdot \lfloor \log_j n \rfloor - 1 = 2d - 1,$	$\mathcal{G}_{n,2d-1}^+$ includes d -dimensional j -ary C-cubes $C_j^d, j \geq 5$	$\lfloor j/2 \rfloor \cdot \log_j n$	$\lfloor j/2 \rfloor \cdot (\log_j n) + \lceil j/2 \rceil - 1$	Definitely <i>not</i> best possible: ratio diverges to ∞ as $n \rightarrow \infty$	Table 17, Theorems 6, 34, 35 Corollary 35.1
$\lfloor (j-1) \cdot \log_j n \rfloor - 1 = (j-1) \cdot d - 1$	$\mathcal{G}_{n,(j-1)d-1}^+$ includes d -dimensional j -ary K-cubes K_j^d	$\log_j n$	$1 + \log_j n$	As $n \rightarrow \infty$: within a factor of 1 of best possible whenever $d \in o(j)$ and $m \in o(d)$ or d and j bounded. Within $1+q+qr+r$ of best possible whenever $\lfloor \frac{m}{2} \rfloor + 1 \leq qd$ and $\ln d \leq r \ln j$, for least upper bounds q, r .	Tables 9, 11, and 14 Corollary 33.1, Theorem 6
$(j-1) \cdot \log_j(n/2) = (j-1) \cdot d$	$\mathcal{G}_{n,(j-1)d}^+$ includes d -dimensional j -ary K-cube-connected edges $K_{2,j}^d, j \geq 3$	2 if $d = 1$			
$1 + (j-1) \cdot \log_j(n/m) = (j-1) \cdot d + 1$	$\mathcal{G}_{n,(j-1)d+1}^+$ includes d -dimensional j -ary K-cube-connected cycles $K_{m,j}^d, m \geq 3$	$1 + \log_j(n/2)$	$2 + \log_j(n/2)$		
		2 if $d = 1$			
		$1 + \lfloor m/2 \rfloor$	$1 + \lfloor m/2 \rfloor$		
		$\lfloor m/2 \rfloor + \log_j(n/m)$	$1 + \lfloor m/2 \rfloor + \log_j(n/m)$		
$n-2, n-1$	$\mathcal{G}_{n,n-2,1}, \mathcal{G}_{n,n-1,1}$ uniquely the set of n -vertex cliques K_n	1		Exactly best possible	Table 7, Theorem 6

Table 18: Radius of quorums induced by deleting vertices from n -vertex graph architectures.

If both m and d are bounded then the only way for the number of vertices to approach infinity is for the radix j to increase. In this case we can improve Corollary 33.1 to best possible.

Corollary 33.2. If $d, m \in \Theta(1)$ then $\lim_{n \rightarrow \infty} \frac{\rho_{m,j,d}^+}{\rho_{m,j,d}} = \frac{\rho_{m,j,d}^-}{\rho_{m,j,d}} = 1$.

In the ratioed asymptotic sense of Corollaries 33.1 and 33.2, both the lower bounds of Theorem 6 and the quorum radius of $K_{m,j}^d$ are best possible. In other cases it may be that one of these bounds is best possible, but this remains to be proved. We also stress that $\rho_{m,j,d}^+ / \rho^-_{\text{Thm 6}}$ and $\rho_{m,j,d}^- / \rho^-_{\text{Thm 6}}$ approach one quite slowly. The reason for this appears to be the $\ln j$ factors in the expressions of (71). As computed by GRAFT, for example, at $(n, f) = (121, 19)$ and $(n, f) = (512, 20)$ we have $(m, j, d) = (1, 11, 2)$ and $(m, j, d) = (1, 8, 3)$; the corresponding ratios are $3/2$ and $4/3$.





Before presenting the last two theorems of this section, let us review our terminology. Refer to the two middle columns of Table 18, as well as to the introductory material on page 9. By the *maximum radius* $\rho(n, f)$ we mean the largest radius of any quorum induced by f or fewer faults. Thus, for example, to obtain a *lower* bound on the maximum radius of a K-cube (*resp.* C-cube), we take the largest of the lower bounds on radii as listed in Table 9 (*resp.* Table 17); for an *upper* bound on the maximum radius of a K-cube or C-cube quorum, we take the largest of the upper bounds on radii as listed in Table 9 *resp.* Table 17. Similarly, introduce the *maximum diameter* $\Delta(n, f)$ as the largest diameter of any quorum induced by f or fewer faults. Thus, for example, to obtain a *lower* bound on the maximum diameter of a K-cube (*resp.* C-cube) quorum, we take the largest of the lower bounds on diameter as listed in Table 9 (*resp.* Table 17); for an *upper* bound on the maximum diameter of a K-cube or C-cube, we take the largest of the upper bounds on diameter as listed in Table 9 *resp.* Table 17. Finally, note that if f is the worst-case fault tolerance of an n -vertex graph architecture, then the *fractional fault (worst-case) tolerance* is simply $f_{\text{frac}} = f/n$. With these notions in hand, we can quantify relative merit of K-cubes and C-cubes.

Theorem 34. If the worst-case fault tolerance f of K_j^d equals that of C_J^D then, for $j, J \geq 5, d, D \geq 2$:

$$\text{The maximum diameter } \Delta_K \text{ of } K_j^d \text{ is less than the maximum radius } \rho_C \text{ of } C_J^D: \quad \Delta_K < \rho_C \quad (72)$$

$$\text{The order } n_K(j, d) \text{ of } K_j^d \text{ is less than the order } n_C(J, D) \text{ of } C_J^D: \quad n_K < n_C \quad (73)$$

Proof. By hypothesis, and by Corollaries 9.1 and 32.1: $f + 1 = d(j - 1) = 2D$ (74)

$$\text{By Table 9:} \quad \Delta_K \leq d + 1 \quad (75)$$

$$\text{By Table 17, and by inequalities (74) and (75):} \quad \frac{1}{2}D(J - 1) = \frac{1}{4}d(j - 1)(J - 1) \leq \rho_C \quad (76)$$

$$\text{For (72) it therefore suffices to show} \quad d + 1 < \frac{1}{4}d(j - 1)(J - 1) \quad (77)$$

$$\text{But (77) holds since } j, J \geq 5, d, D \geq 2, \text{ and} \quad 1 + 1/d \leq 2 < 4 \leq j - 1 \quad (78)$$

Now note that, for integers $r > q \geq 5$, we have $r/q < 6/5 < 1.7 < 2 < 5^{1/2}$. Hence $5^{1/2(q-1)}/q < 5^{1/2(r-1)}/r$ and the value of $j/5^{1/2(j-1)}$ decreases strictly with increasing integer $j \geq 5$. In particular, since $5 < 5^2 = 25$, and since $J \geq 5, d \geq 2$, we can make use of (74): $n_K = j^d < 5^{1/2d(j-1)} \leq J^{1/2d(j-1)} = J^D = n_C$ (79)

Thus, (72) and (73) hold. □

Inequality (72) of Theorem 34 says that, for given fault tolerance, the maximum diameter of K-cube quorums is less than than the maximum radius of C-cube quorums. Moreover, (73) establishes that the worst-case *fractional* fault tolerance of K-cubes is superior to that of C-cubes. Recalling the discussion at the beginning of Section 3.7, Theorem 34 focuses on radices greater than 4 and dimensions greater than 1 since, for $j \leq 4$ or $d = 1$, C-cubes are isomorphic to K-cubes or cycles. But in how many cases can the fault tolerance of a C-cube equal that of a K-cube? That is, for what constructions is the degree of each vertex in a K-cube equal to that $f+1$ of any vertex in a C-cube? By inspection of (74), such a construction is realized if and only the degree of every vertex of the K-cube is an even integer no less than eight. In other words, for $j > 4$ and $d > 1$, Theorem 34 applies to all C-cubes; moreover, Theorem 34 applies to a subset of K-cubes (loosely speaking, "half" of them) that map many-to-one onto the set of C-cubes.

Despite Theorem 34's quantitative preference for K-cubes over C-cubes, it seems plausible that, when divided by $\rho_{\text{Thm 6}}^-$, the maximum radius of C-cube quorums attains a limit, akin to that expressed by Corollaries 33.1 and 33.2. That is, we still do not know whether, for some scaling of j and d , the maximum radius of quorums induced from C_j^d is asymptotically within a constant factor of $\rho_{\text{Thm 6}}^-$. Alas, such scalability is impossible, as the next theorem shows.

Theorem 35. As $n_C(j, d) = j^d$ tends to infinity, the ratio $\rho_C(j^d, 2d-1)/\rho_{\text{Thm 6}}^-$ grows without bound.

Proof. Suppose to the contrary that $\rho_C/\rho_{\text{Thm 6}}^- \in \Theta(1)$. Then for some j, d , and k corresponding to all



$n_C(j,d) \geq k$, the ratio is bounded from above by a least constant $b \geq 1$. As with Theorem 33, we employ

simplifying substitutions to consider
$$\frac{(j-1)\ln d}{\ln j} \leq 2 \left[\frac{\lfloor \frac{j}{2} \rfloor d \ln d}{d \ln j} \right] = \frac{2\rho_C^-}{\rho_{\text{Thm 6}}^-} \leq 2b \quad (80)$$

for such sufficiently large $n_C(j,d) \geq k$. The scaling condition $n \rightarrow \infty$ implies that $(j-1) \cdot \ln d \rightarrow \infty$. Hence, for the upper bound b to exist, the denominator on the lefthand side of (80) must approach infinity: $\ln j \rightarrow \infty$. But this means that $j \rightarrow \infty$. As $j \rightarrow \infty$, $(j-1)/\ln j$ grows without bound; hence there can be no k such that, for all $n_C(j,d) \geq k$, (80) is satisfied. That is, $\rho_C(j^d, 2d)/\rho_{\text{Thm 6}}^-$ grows without bound. \square

Theorem 35 says that the bound of Theorem 6 (a relative of the *Moore bound* mentioned on page 15) cannot be achieved by C-cubes, even in the sense of asymptotic ratios. This is *not* the same as a wholesale assertion about the ratio of C-cube quorum radii to the optimum value of the maximum radius $\rho(n, f)$, and we are not in a position to advance such a claim. However, for scaling trends that enable K-cubes to come within a constant factor of $\rho(n, f)$, we *can* be certain that the ratio $\rho_C/\rho(n_C, f)$ diverges. More precisely:

Corollary 35.1. For $(j-1)d$ even, let j and d be the radix and dimension of the class of K_j^d such that $d \in \Theta(1)$ or, with r the least upper bound such that, for all $n_C = j^d \geq k$, $\ln d \leq r \ln j$. Let $\{C_j^{1/2(j-1)d=D}\}$ be the class of C-cubes corresponding to such K_j^d 's, as prescribed by the discussion following Theorem 34. If $n_K(j,d) = j^d$ tends to infinity then, by equation (73) of Theorem 34, n_C tends to infinity; moreover, by Theorem 35, the ratio $\rho_C/\rho(n_C, f)$ grows without bound.

3.9 Underware for Distributed Configuration

When combined with breadth-first search, the proof of Theorem 2 provides a $\Theta(n+e)$ algorithm for constructing a tree from a graph of order n and size e ([Chartrand and Lesniak 1986]). In particular, if u is a central vertex of a quorum H induced by deleting up to f vertices of G , then applying this algorithm to a central vertex of H gives a spanning tree T whose distances are the same as for H . In particular, the radius of T equals the radius of H . To find the central vertices of H , it suffices to compute the (symmetric) all-distances matrix (use breadth-first search from every vertex in H , running time $O(n(n+e))$), [Cormen, Leiserson, Rivest 1993], Sec. 23.2). Column j of row i in the all-distances matrix gives the distance between vertex i and vertex j . Sort the columns of each row in, say, descending order of the value of the entries (using COUNTING-SORT this can be done in time $\Theta(n)$ per row, $\Theta(n)^2$ overall, [Cormen, Leiserson, Rivest 1993], Sec. 9.2). As a result, the eccentricity of vertex i is now in column 1 of row i . In time $\Theta(n)$, find the radius of the graph by extracting the minimum value in column 1.²¹ Rescan each row of column 1; if the distance in $(i,1)$ equals the radius then insert i in the list of central vertices. To recap:

Theorem 36. For any graph of order n and size e , we can use breadth-first search to compute, on a Turing machine equivalent and in time $O(n(n+e))$, a spanning tree having minimum radius.

The root of the tree T computed by Theorem 36 is a central vertex of both T and the quorum H that T spans. Were we to have a known fault-free node that could control configuration (via, say, an I²C bus) then we could make use of Theorem 36 to compute the root (and consequently, the rest of) a 1394 bus with minimum radius. Unfortunately, this question begs the question of worst-case fault tolerance. In consequence, we need to provide for *distributed* diagnosis and configuration. The attendant algorithms and implementations are *underware*: that is, they underlie and enable successful configuration of a 1394 bus.

21. For an alternative way of computing the radius, (absent proofs of correctness or running time), see Algorithm 12.4 of [Chacra et al 1979].



Under the hypothesis that the number of faults is no greater than the worst-case maximum f , how can we design individual nodes to cooperatively perform diagnosis and configuration? To solve this problem we propose a separate algorithm for each graph architecture. For a star the value of f equals 0, so by assumption the star configures properly as long as it meets other 1394 requirements. To shorten this exposition we give details only in the case of cycles and one-dimensional binary K-cube-connected cycles.

For cycles, Figures 29 and 30 illustrate the action and timing of configuration algorithm A_{cycle} in instances with and without faults. The key idea is to partition C_n into overlapping paths $P_{0, \lfloor n/2 \rfloor}$ and $\bar{P}_{\lfloor n/2 \rfloor, 0}$, with the former spanning $\lfloor n/2 \rfloor + 1$ successive vertices $u_0, u_1, \dots, u_{\lfloor n/2 \rfloor}$, and the latter traversing $\lfloor n/2 \rfloor + 1$ vertices $u_{\lfloor n/2 \rfloor}, u_{\lfloor n/2 \rfloor + 1}, \dots, u_0$. Let us explain A_{cycle} as we establish its correctness and efficiency. An initial bus reset is either a power-up bus reset, or a node insertion or deletion reset, or a software initiated bus reset which is preceded by filling the port disable bits of each node with their default values (period (T_1, T_0) of Figure 30, [Anderson 1998] pp. 244-247, Chap. 14; [P1394 1995] Table 4-28). For A_{cycle} the default values of the port disable bits are set at lines 1 and 2, either in ROM for power-up, or by software after power-up but prior to line 3. At line 3, and as depicted by the red and blue solid lines of Figure 29A(i), nodes simultaneously configure one or two (disjoint) bus(es) from $P_{0, \lfloor n/2 \rfloor}$. If there is no bus fault in $P_{0, \lfloor n/2 \rfloor}$ then by the end of line 11 the path configures as a tree (period (T_0, T_3) of Figure 30).




Figure 29: Distributed configuration by A_{cycle} from a cycle: A. with no faults, B. with one fault.

Recalling the discussion at the beginning of Section 3, the nodes of $P_{0, \lfloor n/2 \rfloor}$ perform mutual test in order to diagnose more completely the health of each node. As implied by line 5 of A_{cycle} , this diagnosis involves all layers of protocols, including application-level exchange via tasks running on each processor. Such high-level diagnosis is in keeping with the spirit of Bob Rasmussen's approach of software lock and key, though in our case the I²C bus need not be involved. An advantage of high-level diagnosis is that much more than just the physical layer of the 1394 is exercised, and the probability of fault detection is increased. A disadvantage is that it may be difficult or impossible for software to distinguish various types of low-level faults, and the probability of fine-grained fault isolation is decreased. However, and as Savio Chau has pointed out, it would be both expensive and risky to modify the VERILOG or VHDL sources for 1394 bus controllers. Our recommendations are consistent with Savio's goal of avoiding this exposure.





Distributed Diagnosis and Configuration Algorithm A_{cycle} % Configure a tree from a cycle

- 1) Enable all ports % Initial port disable values in ROM
- 2) except those between $(u_{n-1}, u_0), (u_{\lfloor n/2 \rfloor}, u_{\lfloor n/2 \rfloor + 1} \bmod n)$
- 3) Initial bus reset % **Bus** reset (not command reset)
- 4) For each of u_i 's enabled ports in $P_{0, \lfloor n/2 \rfloor}$ % Have at most one fault;
- 5) If test $(u_i, u_{\lfloor i-1 \rfloor \bmod n})$ fails % hence at most 2 buses are formed
- 6) then u_i marks $u_{\lfloor i-1 \rfloor \bmod n}$, disables its port to $u_{\lfloor i-1 \rfloor \bmod n}$ % Record results of failed test
- 7) u_i issues a bus reset % and disable immediately
- 8) If test $(u_i, u_{\lfloor i+1 \rfloor \bmod n})$ fails % Performing a bus reset
- 9) then u_i marks $u_{\lfloor i+1 \rfloor \bmod n}$, disables its port to $u_{\lfloor i+1 \rfloor \bmod n}$ % guarantees two leaves
- 10) u_i issues a bus reset
- 11) Propagate the marked status of each node throughout bus % Get info to least one of $u_0, u_{\lfloor n/2 \rfloor}$
- 12) u_0 disables its port to u_1 ; u_1 disables its port to u_0 % Switch to complementary bus
- 13) $u_{\lfloor n/2 \rfloor}$ disables its port to $u_{\lfloor n/2 \rfloor - 1} \bmod n$; $u_{\lfloor n/2 \rfloor - 1} \bmod n$ disables its port to $u_{\lfloor n/2 \rfloor}$
- 14) u_0 enables its port to u_{n-1} ; u_{n-1} enables its port to u_0
- 15) $u_{\lfloor n/2 \rfloor}$ enables its port to $u_{\lfloor n/2 \rfloor + 1} \bmod n$; $u_{\lfloor n/2 \rfloor + 1} \bmod n$ enables its port to $u_{\lfloor n/2 \rfloor}$
- 16) u_0 and $u_{\lfloor n/2 \rfloor}$ issue bus reset % Node insertion/ deletion
- 17) For each of u_i 's enabled ports in $\bar{P}_{\lfloor n/2 \rfloor, 0}$ % Have at most one fault;
- 18) If test $(u_i, u_{\lfloor i-1 \rfloor \bmod n})$ fails % hence at most 2 buses are formed
- 19) then u_i marks $u_{\lfloor i-1 \rfloor \bmod n}$, disables its port to $u_{\lfloor i-1 \rfloor \bmod n}$ % Record results of failed test
- 20) u_i issues a bus reset % and disable immediately
- 21) If test $(u_i, u_{\lfloor i+1 \rfloor \bmod n})$ fails % Performing a bus reset
- 22) then u_i marks $u_{\lfloor i+1 \rfloor \bmod n}$, disables its port to $u_{\lfloor i+1 \rfloor \bmod n}$ % guarantees two leaves
- 23) u_i issues a bus reset
- 24) Propagate the marked status of each node throughout bus % Get info to least one of $u_0, u_{\lfloor n/2 \rfloor}$
- 25) If u_1 is not marked by u_0
- 26) then u_0 enables its port to u_1
- 27) If $u_{\lfloor n/2 \rfloor - 1} \bmod n$ is not marked by $u_{\lfloor n/2 \rfloor}$ % Have at most one fault;
- 28) and some other node is marked % hence $u_0, u_{\lfloor n/2 \rfloor}$, if not faulty,
- 29) then $u_{\lfloor n/2 \rfloor}$ enables its port to $u_{\lfloor n/2 \rfloor - 1} \bmod n$ % has status of marked nodes
- 30) If u_0 is not marked by u_1
- 31) then u_1 enables its port to u_0
- 32) If $u_{\lfloor n/2 \rfloor}$ is not marked by $u_{\lfloor n/2 \rfloor - 1} \bmod n$
- 33) and some other node is marked
- 34) then $u_{\lfloor n/2 \rfloor - 1} \bmod n$ enables its port to $u_{\lfloor n/2 \rfloor}$
- 35) u_0 and $u_{\lfloor n/2 \rfloor}$ issue bus reset % Final configuration




Figure 30: Parallel-series event timeline for distributed diagnosis and configuration algorithm A_{cycle} .






Figure 31: The shaded area represents feasible specifications for mutual test and diagnosis for 0% of faulty called good (solid lines) with the feasible region for at most 2% of good called faulty (dotted lines). Data obtained using DWI simulator. [LaForge and Korver 1997].




Figure 32: Mutual test and diagnosis with 75% of all nodes faulty; h measures the test redundancy.

By contrast to configuration, results for diagnosis are sufficiently well-established to merit direct application to X2000. As Figures 31 and 32 suggest, the *theoretical* basis for mutual test is well-founded [LaForge et al 1994], [LaForge and Korver 1997]. In addition, *practical* experiments with similar approaches suggest that the coverage of high-level diagnosis is very close to 100% [Bianchini and Buskens 1992]. In keeping with this, we assume that tests applied by good nodes are accurate.

Refer to Figure 29B(i). If the directed point-to-point test (u_i, u_{i-1}) fails then (lines 5 and 6) u_i disables its port to u_{i-1} (period (T_1, T_2) of Figure 30). This includes, but is not limited to, the case where u_i sees no response from u_{i-1} prior to arbitration timeout or where u_i sees erroneous output (such as babbling noise) from u_{i-1} . In such an instance u_i becomes a leaf and (line 7) resignals bus reset. The bus returns to the state at line 4. By hypothesis, at most one node is faulty. If u_{i-1} is good then u_i is faulty and we do not want (u_i, u_{i-1}) to be included in the final configuration. If u_{i-1} is faulty then u_i is good and we still do not want (u_i, u_{i-1}) to be included in the final configuration. Hence, the correct action is for u_i to disable its port to u_{i-1} and broadcast this action (lines 6 and 11). A similar argument establishes the correctness of the action (lines 8 through 11) in response to the failure of the directed point-to-point test (u_i, u_{i+1}) . Under a fault model that, strictly speaking, is outside the one we have adopted, it is possible for one of u_i 's neighbors, say u_{i-1} , to disable its connection with u_i , while the other neighbor u_{i+1} maintains an enabled connection to u_i . This is a result of mutual tests that point to a healthy connection between u_i and u_{i+1} , but an unhealthy connection between u_{i-1} and u_i . In this case the only logical possibility is that u_i is good except for its ability to communicate with u_{i-1} . Again, the correct action is for u_i to disable its port to u_{i-1} .

If one of the nodes $u_1, \dots, u_{\lfloor n/2 \rfloor - 1}$ is marked faulty then lines 5 through 10 give rise to two buses from $u_0, u_1, \dots, u_{\lfloor n/2 \rfloor}$. If either u_0 or $u_{\lfloor n/2 \rfloor}$ is marked as faulty then there is only one bus formed. In either case, at the end of line 11 the status of the all the nodes of $u_0, u_1, \dots, u_{\lfloor n/2 \rfloor}$ is known by the fault-free nodes (v, w) with lowest *resp.* highest index between 0 and $\lfloor n/2 \rfloor$. Either u_0 is good and $v = u_0$ or u_0 is





faulty and $v = u_1$. Either $u_{\lfloor n/2 \rfloor}$ is good and $w = u_{\lfloor n/2 \rfloor}$ or $u_{\lfloor n/2 \rfloor}$ is faulty and $v = u_{\lfloor n/2 \rfloor} - 1$. Since at most one node is faulty, however, at least one of the overlapping nodes u_0 or $u_{\lfloor n/2 \rfloor}$ will be included in (v, w) (and if only one then the other overlapping node is faulty). Nodes $u_0, u_{\lfloor n/2 \rfloor}$ and their neighbors rewrite their port registers to (lines 12 through 15) enable configuration (line 16) of $\bar{P}_{\lfloor n/2 \rfloor, 0}$ on the next bus reset (period (T_3, T_4) of Figure 30). To allow for the possibility of a single fault, this reset is initiated by both u_0 and $u_{\lfloor n/2 \rfloor}$ (line 16, period (T_4, T_5) of Figure 30). As shown in Figure 29A(ii) and 29B(ii), the mutual test, diagnosis, and configuration of $\bar{P}_{\lfloor n/2 \rfloor, 0}$ (lines 17 through 24) proceeds as for $P_{0, \lfloor n/2 \rfloor}$.

The hypothesis of at most one faulty node assures that at least one of $P_{0, \lfloor n/2 \rfloor}$ or $\bar{P}_{\lfloor n/2 \rfloor, 0}$ (both, if there are no faults) forms a single bus. If there are no faults then (lines 25 through 35) nodes form a tree rooted at u_0 , with leaves $u_{\lfloor n/2 \rfloor}$ and $u_{\lfloor n/2 \rfloor + 1}$ (Figure 29A(iii), period (T_7, T_9) of Figure 30). If there is a "one-sided" fault (only one of a node's neighbors declares it faulty) then we form a tree whose leaves are the faulty node and the neighbor declaring it faulty. Otherwise, two neighbors of some node agree that the node is faulty, and these neighbors become leaves of the 1394 bus. In summary:

Theorem 37. In the presence of any one faulty node, and in at most four bus reset periods (three resets altogether), A_{cycle} configures a tree of diameter at most $n-1$ from an n -vertex cycle (the only 1-tolerant architecture with minimum count n of point-to-point connections).

Let us mention a few points concerning the implementation of A_{cycle} . The 1394 specification prescribes that 1394 topology maps are not preserved across bus resets, and so this information cannot be used ([Anderson 1998] p. 254). For this reason the results of mutual test are recorded in the memory of each node, and survive subsequent software-initiated bus resets (note that software bus resets do not require rebooting of the operating system on each node). Except where power is lost, however, port disable bits are preserved across resets ([Anderson 1998] p. 261). Implementing A_{cycle} (or its analog for cliques, K-cubes, or K-cube-connected edges or cycles) will require a careful estimate of each T_i in Figure 30. For example, consider the beginning of (T_4, T_5) . The $167\mu\text{s}$ minimum reset duration, if adhered to, represents an upper bound on the windows of tolerance (which must account for setup and hold on drivers and receivers) as u_0 and $u_{\lfloor n/2 \rfloor}$ switch from one sub-bus to another ([Anderson 1998] p. 262).

The preceding explanation and proof applies to diagnosis and configuration of an architecture based on C_n , where diagnosis is carried out after a bus reset. What about the case of faults which appear after bus reset, and in the course of nominal bus operation? Under the assumption of no "one-sided faults", detection and configuration in the presence of a single fault u can be readily carried out by periodically performing the same point-to-point tests used in A_{cycle} . If u is a leaf in the current tree then the parent of u disconnects itself from u , signals a bus reset, and A_{cycle} configures the truncated tree beginning at line 4. If u is an interior node in the current tree then the neighbors of u command the two leaves to enable their ports to each other, and then issue bus resets. A_{cycle} configures the truncated tree beginning at line 4. Similarly, the case of "one-sided" faults can be handled by adding tests destined two links from the testing node. Note that isolation of a faulty node does not depend on obtaining its cooperation, an advantage over the proposed "back-door" scheme using the I²C bus.

Let us now consider diagnosis and configuration for $K_2^1(n)$, the architecture perhaps most pertinent to X2000. As in the case of cycles, if we have a fault-free, global means of diagnosis and configuration then we can achieve the bounds of Table 18 using Theorem 36. We can also achieve these bounds in a distributed, parallel fashion. For positive integer q , algorithm $A_2^1(4q)$ configures a quorum in the case $n = 2m = 2 \cdot 2 \cdot q \geq 12$; the remaining two cases ($n = 2m+1$ or $n = 2m$, m odd) are similar. For the sake of brevity we omit pseudocode for $A_2^1(4q)$.

Refer to Figure 34. The key idea is to consider *mated* pairs of nodes (u, v) , one from each of the two cycles comprising $K_2^1(4q)$; the low order digit on the label of u equals the low order digit on the label of v , and so





(u, v) is an edge in $K_2^1(4q)$. A mated pair is *available* if at least one of its nodes is good; otherwise it is *unavailable*. Under the hypothesis that any instance of $K_2^1(4q)$ contains at most two faults, at most one mated pair is unavailable. An instance of $K_2^1(4q)$ maps to an instance of C_{2q} as follows. Each available mated pair in $K_2^1(4q)$ corresponds to a good node in C_{2q} ; an unavailable mated pair in $K_2^1(4q)$ (there is at most one) corresponds to a fault in C_{2q} . The problem of configuring an instance of $K_2^1(4q)$ reduces to that of configuring an instance of C_{2q} , as long as adjacent available mated pairs can be connected. This is illustrated by comparing Figure 29 with steps e through g of Figures 33 through 36.

To assure that adjacent available mated pairs can be interconnected, we schedule four preliminary steps. Refer to steps a through d of Figures 33 through 37. At each step, $K_2^1(4q)$ is partitioned into q disjoint paths; the nominal length of each path (in the absence of faults) equals 4. Similar to the use of paths in algorithm A_{cycle} , each path scheduled by $A_2^1(4q)$ forms a disjoint 1394 bus or, in the presence of faults, at most two disjoint buses. Also as in A_{cycle} , nodes perform point-to-point tests on their neighbors. In the event of an arbitration timeout induced by an unresponsive neighbor, a node disables the port to that neighbor and signals a bus reset on the bus formed so far. Thus, each of steps a through d of Figures 33 through 37 may take two bus resets to stabilize. This is a consequence of the 1394 specifications with respect to PARENT_NOTIFY and CHILD_NOTIFY (Chapters 13, 14, 15, [Anderson 1998]).




Figure 33: Distributed configuration of a minimum radius tree via $A_2^1(16)$. Compare with Figure 29A.

At the end of step d, each good node contains the status of its mate, as well as the status of the nodes contained in its neighboring mated pairs. If both nodes in a mated pair (u, v) are good then, in steps e, f, and g, the respective nodes enable the connection (u, v) . Also in steps e, f, and g, either u or v enables the connection to a node w in its neighboring mated pair, if and only if i) w is good; ii) the label on w is greater than that of any other good node (there is at most one other) in the mated pair to which w belongs; iii) A_{cycle} would schedule the corresponding nodes in C_{2q} to be connected. Suppose that only one node, say x , in a mated pair is good. Node x enables its connections to a node y in a neighboring mated pair if and only if i) y is good; ii) A_{cycle} would schedule the corresponding nodes in C_{2q} to be connected.






Figure 34: Action of distributed configuration algorithm $A_2^1(16)$ in the presence of a single fault.




Figure 35: Action of $A_2^1(16)$ in the presence of two faults occurring in the same pair of mated nodes.

Under the preceding conditions, verify by enumeration that adjacent available mated pairs can be interconnected, unless we have the (local) fault pattern illustrated by Figure 37. As illustrated by Figures 33 through 36, the former cases are handled by A_{cycle} , as previously proved. The pattern depicted in Figure 37 amounts to adjacent faults in C_{2q} ; although not proved previously, A_{cycle} successfully configures this instance as well. Note that in steps e and f we need not perform any point-to-point tests, but instead execute just those portions of A_{cycle} which propagate node status to the "overlapping" mated pairs. Observe also that, by ensuring that any fault occurs in a mated pair corresponding to a leaf in C_{2q} , any path between two good nodes in the configured tree traverses at most two times between the basic cycles. Therefore:



Theorem 38. In the presence of any two faulty nodes, and in at most 11 bus reset periods, $A_2^1(4q)$ configurations from $K_2^1(4q)$ a tree of diameter at most $2q+1 = n/2 + 1$.




Figure 36: Action of $A_2^1(16)$ in the presence of two faults occurring in separated pairs of mated nodes.




Figure 37: Action of $A_2^1(16)$ in the presence of two faults occurring in neighboring pairs of mated nodes.

3.10 Application to X2000

We conclude our technical development by illustrating how our results for architectures and algorithms apply in the case of sparse tolerance ($f \leq 3$) to node failures. If we take $f = 1$ then by Table 7 the unique minimum size architecture is a C_n . By Table 7, the maximum diameter of a tree spanning a quorum of C_n equals $n-1$. Since the maximum diameter of a 1394 bus is 16, the maximum number of nodes in single fault-tolerant architecture with fewest point-to-point interconnects equals 17. As illustrated in Figure 38, we could as well come to this conclusion by using GRAFT, whose logic incorporates Table 7. We obtain a





conservative estimate on the diameter of the spanning tree by doubling the upper bound on quorum radius. Alternatively, we can examine the detailed worksheet, in this case the sheet entitled "Cycle". For n > 17 we must either resort to architectures that are not of minimum size, or we must increase the fault tolerance. While the former is beyond the scope of this report, the latter may be reasonable, especially considering space shuttle requirements for tolerance to two faults.

GRAFT: GRaph Architecture Fault Tolerance Calculator, Version 2.0. Computes n-node f-fault tolerant graph architectures having minimum number of point-to-point connections, bounded radius p and diameter.			Copyright 1999 by Laurence E. LaForge, NASA/ASEE Summer Faculty Fellow. 10-Oct-1998, 18-Oct-1999. Reprint rights granted to NASA and to the ASEE for research and educational purposes. Based on theory developed in my report: <i>Fault Tolerant Physical Interconnection of X2000 Computational Avionics.</i>	
Input:	n = number of nodes	f = maximum number of partitioning faults	e = minimum number of point-to-point connections:	Average number of point-to-point connections per node (number of ports per node)
	18	1	18	2.00
Recommended:	Cycle on 18 vertices			9 9
Input:	n = number of nodes	f = maximum number of partitioning faults	e = minimum number of point-to-point connections:	Average number of point-to-point connections per node (number of ports per node)
	17	1	17	2.00
Recommended:	Cycle on 17 vertices			8 8

Figure 38: GRAFT computes the maximum number 17 of nodes in a minimum size single fault tolerant architecture whose quorums are all spanned by a tree having diameter within the 1394 bus's limit of 16.

Refer to Figure 39. For f = 2 GRAFT is able to construct a K2^1(n) for all n > 2, and furthermore tells us that the diameter remains within our limit of 16 as long as n ≤ 30. GRAFT's upper bound on a maximum diameter equals 16 for n = 30, 29, 28, and 27, but decreases to 14 at n = 26. Refer to Figure 40. For f = 3 the upper bound 16 on diameter, as computed by GRAFT, attains the limits imposed by the 1394 bus at n = 44, 40, 39, and 36. Within the range 3 < n ≤ 44, GRAFT is able to find only Km,j^d(n)'s whose dimension d equals 1 or 2, and whose radix j equals 2 or 3. As mentioned in Section 3.8, the radius of Km,j^d(n) may not be a monotone function of n. This is born out at n = 42, wherein GRAFT identifies a K14,3^1 whose quorum diameter may be as much as 18.

Figure 41 illustrates the proposed architecture of [Charlan et al 11-Jun-1998]. According to the diagram, the degree of each node is either 4 or 2. By the discussion at the top of page 9, this renders the architecture tolerant to at most one fault, and furthermore leaves either 2 or 4 ports per node unused. In a 2-Sep-1998 conversation, Carl Steiner and Don Hunter explained that the intention is to maximize the number of connected ports. With this clarification, the architecture of Figure 41 can be redrawn as the multigraph of Figure 42A. Let us analyze this multigraph architecture. First note that having two sets of wires ("for redundancy") between pairs of nodes does not increase the tolerance to nodes whose failure acts to partition the bus. Evidently, the duplicate sets of wires account for the possibility of faulty ports on nodes which otherwise function properly. However, this reasoning is contrary to our understanding of the X2000 fault model, whereby each node is itself a fault containment region. We have been unable to identify any document that points to a bus controller, or a portion of a bus controller, as a fault containment region

Further, Carl Steiner and Don Hunter accord negligible probability to the event of a break in the wires between ports. In the interest of conserving both circuit area and y-axis connector pins, I recommend dispensing with the duplicate 1394 bus. Moreover, by pre-designating two roots, the architectures of Figures 41 and 42A unnecessarily reduce from 2 to 1 the tolerance to partitioning faults (i.e., if each of the designated roots is faulty then we cannot form a tree that spans the quorum). I recommend not pre-designating any pair of nodes as candidates for the root of the tree to be configured. If this tact is taken, then the number of such pre-designated roots should be no less than one plus the number of faults tolerated.





GRAFT: Graph Architecture Fault Tolerance			Copyright 1999 by Laurence E. LaForge, NASA/ASEE Summer Faculty Fellow. 10-Oct-1998, 18-Oct-1999. Reprint rights granted to NASA and to the ASEE for research and educational purposes. Based on theory developed in my report: <i>Fault Tolerant Physical Interconnection of X2000 Computational Avionics.</i>	
Calculator, Version 2.0. Computes n -node f -fault tolerant graph architectures having minimum number of point-to-point connections, bounded radius ρ and diameter.				
Input:	n = number of nodes	f = maximum number of partitioning faults	e = minimum number of point-to-point connections:	Average number of point-to-point connections per node (number of ports per node)
	31	2	47	3.03
	Feasible graph architecture(s) with minimum number of point-to-point connections:			Graph radius $\rho(n,f)$ of quorum and of tree spanning the quorum
				At least At most
Recommended:	1-dimensional 2-ary K-cube-connected cycle with 1 cycle containing 15 vertices, along with 1 cycle containing 16 vertices			8 9
Input:	n = number of nodes	f = maximum number of partitioning faults	e = minimum number of point-to-point connections:	Average number of point-to-point connections per node (number of ports per node)
	30	2	45	3.00
Recommended:	1-dimensional 2-ary K-cube-connected cycle with 2 cycles, each containing 15 vertices			8 8
Input:	n = number of nodes	f = maximum number of partitioning faults	e = minimum number of point-to-point connections:	Average number of point-to-point connections per node (number of ports per node)
	29	2	44	3.03
Recommended:	1-dimensional 2-ary K-cube-connected cycle with 1 cycle containing 14 vertices, along with 1 cycle containing 15 vertices			8 8
Input:	n = number of nodes	f = maximum number of partitioning faults	e = minimum number of point-to-point connections:	Average number of point-to-point connections per node (number of ports per node)
	28	2	42	3.00
Recommended:	1-dimensional 2-ary K-cube-connected cycle with 2 cycles, each containing 14 vertices			8 8
Input:	n = number of nodes	f = maximum number of partitioning faults	e = minimum number of point-to-point connections:	Average number of point-to-point connections per node (number of ports per node)
	27	2	41	3.04
Recommended:	1-dimensional 2-ary K-cube-connected cycle with 1 cycle containing 13 vertices, along with 1 cycle containing 14 vertices			7 8
Input:	n = number of nodes	f = maximum number of partitioning faults	e = minimum number of point-to-point connections:	Average number of point-to-point connections per node (number of ports per node)
	26	2	39	3.00
Recommended:	1-dimensional 2-ary K-cube-connected cycle with 2 cycles, each containing 13 vertices			7 7

Figure 39: GRAFT computes a maximal number 30 of nodes in a minimum size 2-fault tolerant graph architecture whose quorums are all spanned by a tree having diameter within the 1394 bus's limit of 16.





GRAFT: Graph Architecture Fault Tolerance			Copyright 1999 by Laurence E. LaForge, NASA/ASEE Summer Faculty Fellow. 10-Oct-1998, 18-Oct-1999. Reprint rights granted to NASA and to the ASEE for research and educational purposes. Based on theory developed in my report: <i>Fault Tolerant Physical Interconnection of X2000 Computational Avionics.</i>	
Calculator, Version 2.0. Computes n -node f -fault tolerant graph architectures having minimum number of point-to-point connections, bounded radius ρ and diameter.				
Input:	n = number of nodes	f = maximum number of partitioning faults	e = minimum number of point-to-point connections:	Average number of point-to-point connections per node (number of ports per node)
	45	3	90	4.00
	Feasible graph architecture(s) with minimum number of point-to-point connections:			Graph radius $\rho(n,f)$ of quorum and of tree spanning the quorum
				At least At most
Recommended:	1-dimensional 3-ary K-cube-connected cycle with 3 cycles, each containing 15 vertices			8 9
Input:	n = number of nodes	f = maximum number of partitioning faults	e = minimum number of point-to-point connections:	Average number of point-to-point connections per node (number of ports per node)
	44	3	88	4.00
Recommended:	2-dimensional 2-ary K-cube-connected cycle with 4 cycles, each containing 11 vertices			7 8
Input:	n = number of nodes	f = maximum number of partitioning faults	e = minimum number of point-to-point connections:	Average number of point-to-point connections per node (number of ports per node)
	42	3	84	4.00
Recommended:	1-dimensional 3-ary K-cube-connected cycle with 3 cycles, each containing 14 vertices			8 9
Input:	n = number of nodes	f = maximum number of partitioning faults	e = minimum number of point-to-point connections:	Average number of point-to-point connections per node (number of ports per node)
	40	3	80	4.00
Recommended:	2-dimensional 2-ary K-cube-connected cycle with 4 cycles, each containing 10 vertices			7 8
Input:	n = number of nodes	f = maximum number of partitioning faults	e = minimum number of point-to-point connections:	Average number of point-to-point connections per node (number of ports per node)
	39	3	78	4.00
Recommended:	1-dimensional 3-ary K-cube-connected cycle with 3 cycles, each containing 13 vertices			7 8
Input:	n = number of nodes	f = maximum number of partitioning faults	e = minimum number of point-to-point connections:	Average number of point-to-point connections per node (number of ports per node)
	36	3	72	4.00
Feasible, but not recommended:	2-dimensional 6-ary C-cube			6 8

Figure 40: GRAFT computes a maximal number 44 of nodes in a minimum size 3-fault tolerant graph architecture whose quorums are all spanned by a tree having diameter within the 1394 bus's limit of 16.



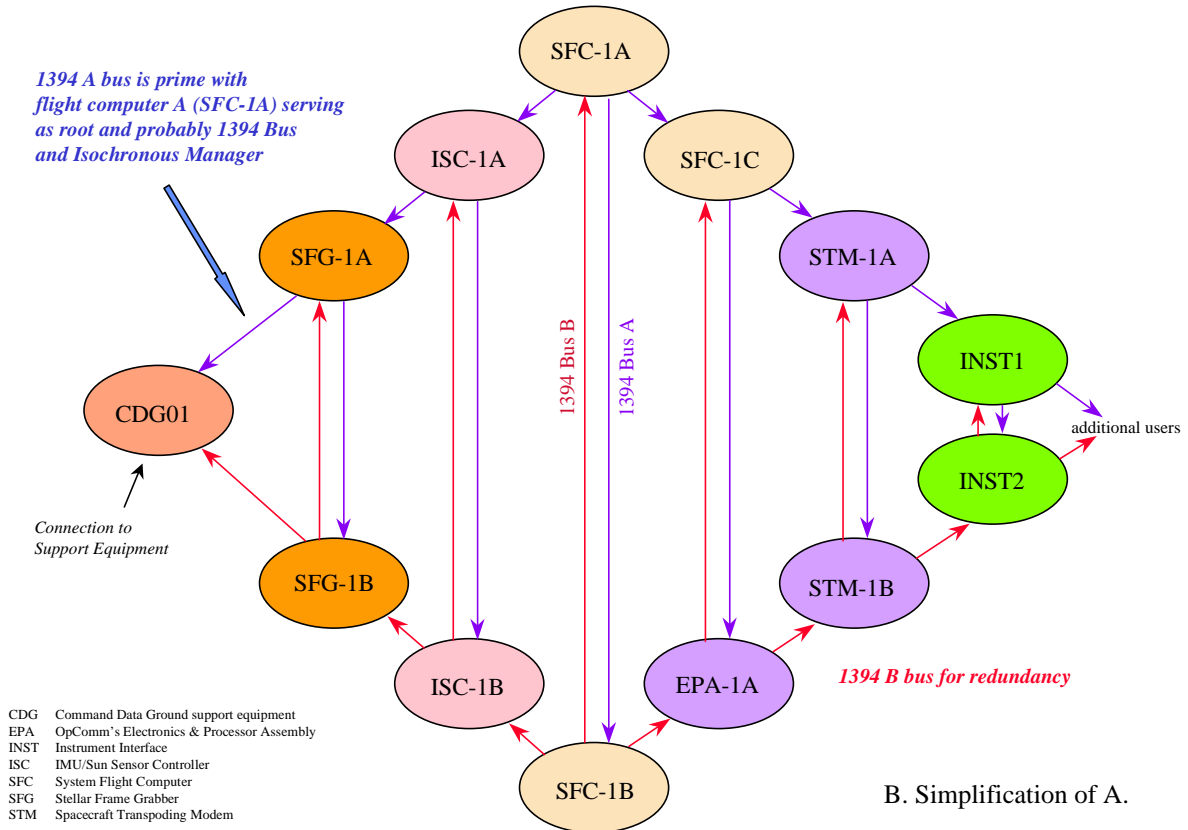



Figure 41: 1394 bus architecture proposed for X2000 ([Steiner 11-Mar-1997]). Not shown is the "back-door" I²C bus ([Charlan et al 11-Jun-1998], Option D).






Figure 42: Refinements of architecture of Figure 41.

A. As explained by Carl Steiner and Don Hunter. B. $K_2^1(n)$ as recommended by this report.

Prior to this report, the most recent recommendation for X2000 bus fault tolerance was Option D of [Charlan et al 11-Jun-1998]. Option D combines a "back-door" I²C bus with the architecture of either Figure 41 or 42A. However, and as pointed out in Section 2, it takes only one faulty node to defeat the multidrop I²C. Suppose that this faulty node happens to be one of the pre-designated roots of Figures 41 and 42A, and that a second pre-designated root fails. In such a case we cannot form a tree that spans the quorum. Therefore, even with a "back-door" I²C bus, the worst-case fault tolerance of Option D is at most one. At a cost of six 1394 ports (36 wires) and one I²C connection (two wires) per node, we are substantially overpaying for single fault tolerance. The architecture that I recommend is a refinement of Option C as described in [Charlan et al 11-Jun-1998]. Assuming that the avionics package is populated by at least 18 but no more than 44 nodes, an economical solution is the $K_2^1(n)$ depicted by Figure 42B. In this case we halve the number of 1394 ports and eliminate the I²C bus. Doing this recovers 20 input/output pins per node, and at the same time increases the fault tolerance from 1 to 2. Alternatively, we can keep six ports per node, eliminating only the I²C. In this case, and as computed by GRAFT, we can tolerate five faults in as many as 96 nodes, all the while staying within the 16 hop limit imposed by the 1394 bus.

Certainly, we have not considered every detail of X2000 avionics. As with any model, the applicability of our results is properly tested as details are factored in. For example, let us review the extent to which our analysis is consistent with considerations of power ([Anderson 1998], chapter 20) and flight computers.

Power is sourced to the bus through *switch slices*. In the worst case, the number of dead switch slices that can be tolerated is no greater than the number of switch slices minus the minimum number of switch slices that can support the bus. Similarly, the number of faults tolerated is no greater than the number of flight computers minus the minimum number of flight computers necessary to complete the mission (for first delivery, one working flight computer). For example, if the bus can complete its mission with a single flight computer and a single switch slice then building three flight computers and three switch slices, as part of a $K_2^1(n)$, maintains 2 fault tolerance. Dropping to (say) two switch slices or two flight computers reduces the tolerance to 1, even though there exists a quorum in the presence of any two faulty nodes.

The caveats, of course, are that the margins at the switch slices, combined with capacitive buffering at each node, are sufficient to accommodate the RLC transient associated with a bus power disconnect. To guard against over and under voltage, each node's PHY layer should tie together, with breakers, all power inputs. The basis for recovering from overvoltage was suggested in [Charlan et al 11-Jun-1998]. With these caveats, maintaining connectivity among all the working nodes (including at least one working switch slice) suffices to maintain a working PHY layer in each of the good nodes.

In this section we have reinforced the use of GRAFT for deciding on an architecture, and applied GRAFT to sparse fault tolerance for X2000. Having settled on an architecture, it remains to develop an algorithm for distributed diagnosis and configuration. Section 3.9 spells out algorithms for the cases $f = 1$ and $f = 2$. We leave as future work the extension of these algorithms to cliques, K-cubes, and K-cube-connected edges and cycles. At the outset, I estimate that development of a repertoire of such algorithms would take 160 hours.





A. References

A.1 Related NASA Documents

- [Barry 22-Jan-1998] R. C. Barry. X2000 /TMOD fault scenario. Viewgraph presentation. Jet Propulsion Laboratory, Pasadena, CA, 22-Jan-1998.
- [Barry 26-Feb-1998] R. C. Barry. X2000 fault protection approach, plan. Viewgraph presentation. Jet Propulsion Laboratory, Pasadena, CA, 26-Feb-1998.
- [Chau 17-Apr-1998] S. Chau. X2000 core avionics peer review: Fault protection. Viewgraph presentation. Jet Propulsion Laboratory, Pasadena, CA, 17-Apr-1998. Online at <http://knowledge.jpl.nasa.gov/adssdlib/>.
- [Chau and Holmberg 17-Apr-1998] S. Chau and E. Holmberg. X2000 core avionics peer review: CDH slices description. Viewgraph presentation. Jet Propulsion Laboratory, Pasadena, CA, 17-Apr-1998. Online at <http://knowledge.jpl.nasa.gov/adssdlib/>.
- [Charlan et al 11-Jun-1998] W. Charlan, S. Chau, J. Donaldson, D. Geer, C. Guiar, H. Luong, N. Palmer, V. Randolph, R. Rasmussen, C. Steiner, S. Woods. X2000 core avionics peer review: bus tiger team. Viewgraph presentation. Jet Propulsion Laboratory, Pasadena, CA, 11-Jun-1998.
- [Chau 11-Jun-1998] S. Chau. X2000 core avionics peer review: characteristics of a reliable bus. Viewgraph presentation. Jet Propulsion Laboratory, Pasadena, CA, 11-Jun-1998.
- [Chau 18-Aug-1998] S. Chau. Backup: upstream connection failure. X2000 preliminary design review. Viewgraph presentation. Jet Propulsion Laboratory, Pasadena, CA, 18, 19, 20-Aug-1998.
- [Guiar 11-Jun-1998] C. Guiar. X2000 core avionics peer review: design approach/ priorities. Viewgraph presentation. Jet Propulsion Laboratory, Pasadena, CA, 11-Jun-1998.
- [Guiar 23-Jul-1998] C. Guiar. X2000 level 3 requirements. Jet Propulsion Laboratory, Pasadena, CA, 17-Apr-1998. Online at <http://knowledge.jpl.nasa.gov/adssdlib/>.
- [Guiar Jun-1998] C. Guiar. X2000 fault containment regions. Viewgraph presentation. Jet Propulsion Laboratory, Pasadena, CA, post 11-Jun-1998.
- [Hunter 11-Mar-1997] D. J. Hunter. X2000 design, implementation, and cost review: first delivery project: electronic packaging. Viewgraph presentation. Jet Propulsion Laboratory, Pasadena, CA, 11 through 13-Mar-1997. Online at <http://knowledge.jpl.nasa.gov/adssdlib/>.
- [JPL D-5703 23-Jul-1998] Management and quality policy. JPL document 5703. 23-Jul-1998. Online at <http://dmie.jpl.nasa.gov/>.
- [JPL 4-11 1-Apr-1984] Reliability assurance, standard practice instruction. JPL policy reference 4-11. 1-Apr-1984. Online at <http://dmie.jpl.nasa.gov/>.
- [Kemski 14-Jul-1998] R. P. Kemski. X2000 mission assurance plan. Document JPL-D-15516. Jet Propulsion Laboratory, Pasadena, CA, 14-Jul-1998. Online at <http://knowledge.jpl.nasa.gov/adssdlib/>.
- [Rasmussen 11-Jun-1998] R. Rasmussen. X2000 core avionics peer review: symmetric architecture. Viewgraph presentation. Jet Propulsion Laboratory, Pasadena, CA, 11-Jun-1998.
- [Steiner 11-Mar-1997] C. Steiner. X2000 design, implementation, and cost review: first delivery project: avionics system engineering. Viewgraph presentation. Jet Propulsion Laboratory, Pasadena, CA, 11 through 13-Mar-1997. Online at <http://knowledge.jpl.nasa.gov/adssdlib/>.
- [Steiner 17-Apr-1998] C. Steiner. X2000 core avionics peer review: avionics overview. Viewgraph presentation. Jet Propulsion Laboratory, Pasadena, CA, 17-Apr-1998. Online at <http://knowledge.jpl.nasa.gov/adssdlib/>.
- [Woerner, Spear, Parker 7-Aug-1998] D. Woerner, A. Spear, G. Parker. X2000 first delivery implementation plan. Draft. Document JPL-D-15438. Jet Propulsion Laboratory, Pasadena, CA, 7-Aug-1998. Online at <http://knowledge.jpl.nasa.gov/adssdlib/>.

A.2 Related Books and Dissertations

- [Anderson 1998] D. Anderson. *Firewire System Architecture*. Reading, MA: Addison Wesley, 1998.
- [Artin 1975] M. Artin. *Linear Algebra and Group Theory*. Massachusetts Institute of Technology: notes for course 18.701, 1975.
- [Bollabás, 1978] B. Bollabás. *Extremal Graph Theory*. London: Academic Press, 1978.
- [Blough 1988] D. M. Blough. *Fault Detection and Diagnosis in Multiprocessor Systems*. Ph.D. dissertation. Baltimore: Johns Hopkins University, 1988.
- [Chacra et al 1979] V. Chacra, P. M. Ghare, and J. M. Moore. *Applications of Graph Theory Algorithms*. New York: North Holland, 1979.





- [Chartrand and Lesniak 1986] G. Chartrand and L. Lesniak. *Graphs and Digraphs*. Belmont, CA: Wadsworth, Inc, 1986. 2nd ed.
- [Comtet 1974] L. Comtet. *Advanced Combinatorics*. Dordrecht, Holland: R. Reidel Publishing, 1974. 2nd ed.
- [Cormen, Leiserson, Rivest 1993] T. H. Cormen, C. E. Leiserson, and R. L. Rivest. *Introduction to Algorithms*. Cambridge: MIT Press. Tenth printing, 1993.
- [LaForge 1991] L. E. LaForge. *Fault Tolerant Arrays*. PhD dissertation. Montreal: McGill University, 1991.
- [Ore 1962] O. Ore. *Theory of Graphs*. Providence: American Mathematical Society Publications, 1962.
- [P1394 1995] *P1394: Standard for a High Performance Serial Bus*. Institute of Electrical and Electronics Engineers, Inc. New York: Draft 8.0v2, July, 1995.
- [Paret and Fenger 1997] D. Paret and C. Fenger. *The P^2C Bus: From Theory to Practice*. New York: John Wiley and Sons, 1997.
- [Thomas 1969] G. B. Thomas. *Calculus and Analytic Geometry*. Reading, MA: Addison Wesley, 1969. 4th ed.
- [Tucker 1984] Alan Tucker. *Applied Combinatorics*. New York: John Wiley and Sons, 1984.
- [Wakerly 1990] J. F. Wakerly. *Digital Design: Principles and Practice*. Englewood Cliffs, NJ: Prentice-Hall. 2nd ed.
- [Zargham 1996] M. R. Zargham. *Computer Architecture: Single and Parallel Systems*. Upper Saddle River, NJ: Prentice-Hall, 1996.

A.3 Related Papers and Articles

- [Armstrong and Gray 1981] J. R. Armstrong and F. G. Gray. *Fault diagnosis in a boolean n cube of microprocessors*. *IEEE Transactions on Computers*. Vol. C-30, No. 8, August, 1981.
- [Banerjee et al 1986] P. Banerjee, S-Y Kuo, and W. K. Fuchs. *Reconfigurable cube-connected cycles architectures*. *Proceedings, 16th International Symposium on Fault Tolerant Computing*. July, 1986, pp. 286-291.
- [Bianchini and Buskens 1992] R. P. Bianchini, Jr. and R. W. Buskens. *Implementation of on-line distributed systems-level diagnosis theory*. *IEEE Transactions on Computers*. Vol. 41, No. 5, May, 1992.
- [Chen and Upadhyaya 1993] Y-Y Chen and S. J. Upadhyaya. *Reliability, reconfiguration, and spare allocated issues in binary tree architectures based on multiple-level redundancy*. *IEEE Transactions on Computers*. Vol. 42, No. 6, June, 1993.
- [Hayes 1976] J. P. Hayes. *A graph model for fault tolerant computing systems*. *IEEE Transactions on Computers*. Vol. C-25, No. 9, September, 1976. pp. 875-884.
- [Koren and Pradhan 1986] I. Koren and D. K. Pradhan. *Yield and performance enhancement in VLSI and WSI multi-processor systems*. *Proceedings of the IEEE*. Vol. 74, No. 5, May, 1986. pp. 699-711.
- [Kwan and Toida 1981] C.-L. Kwan and S. Toida. *Optimal fault-tolerant realizations of some classes of hierarchical tree systems*. *11th Annual Symposium on Fault-Tolerant Computing*. 24-26 June, 1981. New York: IEEE Press. pp. 176-178.
- [LaForge et al 1994] L. E. LaForge, K. Huang, and V. K. Agarwal. *Almost sure diagnosis of almost every good element*. *IEEE Transactions on Computers*. Vol. 43, 3, March, 1994. pp. 295-305.
- [LaForge 1994] L. E. LaForge. *Feasible regions quantify the configuration power of arrays with multiple fault types*. *Proceedings, European Dependable Computing Conference I*. K. Echtler, D. Hammer, and D. Powell, eds. Berlin: Springer-Verlag, 1994. pp. 453-469.
- [LaForge 1997] L. E. LaForge. *Configuration for fault tolerance*. *Reconfigurable Architectures: High Performance by Configware* (Proceedings, 1997 International Workshop on Reconfigurable Architectures). R. W. Hartenstein and V. K. Prasanna, eds. Bruchsal, Germany: ITpress Verlag. pp. 117-120. Online portable document format (PDF) at <http://ec.db.erau.edu/cce/centers/faculty/laforge/Refereed/>.
- [LaForge 1998] L. E. LaForge. *Configuration of locally spared arrays in the presence of multiple fault types*. *IEEE Transactions on Computers*. Vol. 48, No. 4, April, 1999. pp 398-416.
- [LaForge and Korver 1997] L. E. LaForge and K. F. Korver. *Mutual test and diagnosis*. *IEEE Design and Test*. Submitted 7-Feb-1997. Online at <http://ec.db.erau.edu/cce/centers/faculty/laforge/Publications/Manuscripts/Mutual-Test-and-Diagnosis.pdf>.
- [McDermid and Talbot 1998] J. McDermid and N. Talbert. *The cost of COTS*. *Computer*. June, 1998. pp. 46-52.
- [Murty and Vijayan 1964] U. S. R. Murty and K. Vijayan. *On accessibility in graphs*. *Sakhya Ser. A*, Vol 26, 1964. pp. 299 - 302.
- [Preparata et al 1967] F. Preparata, G. Metzger, and R. Chien. *On the connection assignment problem of diagnosable systems*. *IEEE Transactions on Computers*, **EC-16**, 6, December, 1967. pp. 848-854.
- [Preparata and Vallemin 1981] F. P. Preparata and J. Vallemin. *The cube-connected cycles, a versatile network for parallel computation*. *Communications of the ACM*, May 1981. pp. 30-39.





- [Sampels 1997] M. Sampels. [Large networks with small diameter](#). Graph-Theoretic Concepts in Computer Science, 23rd International Workshop. Berlin: 1997. R. H. Möhring, editor. pp 288-302.
- [Scheinerman 1987] E. R. Scheinerman. [Almost sure fault tolerance in random graphs](#). *SIAM Journal of Computing*. Vol. 16, No. 6, December, 1987. pp. 1124-1134.
- [Somani and Agarwal 1987] A. K. Somani and V. K. Agarwal. [Distributed diagnosis algorithms for regular interconnected structures](#). *IEEE Transactions on Computers*. Vol 41, 7, July, 1992. pp. 899-906.
- [Turán 1954] P. Turán. [On the theory of graphs](#). *Colloquium Mathematicum*. **III**, 1954. pp. 19-30.

