INFORMATION AND COMPUTATION **102**, 280–314 (1993)

Subcube Fault-Tolerance in Hypercubes

NIALL GRAHAM*,† AND FRANK HARARY†

Computing Research Laboratory, Department of Computer Science, New Mexico State University, Las Cruces, New Mexico 88003-0001

AND

MARILYNN LIVINGSTON‡,§ AND QUENTIN F. STOUT¶

Advanced Computer Architecture Laboratory, Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, Michigan 48109-2122

We consider the problem of determining the minimum number of faulty processors, $\kappa(n, m)$, and of faulty links, $\lambda(n, m)$, in an n-dimensional hypercube computer so that every m-dimensional subcube is faulty. Best known lower bounds for $\kappa(n, m)$ and $\lambda(n, m)$ are proved, several new recursive inequalities and new upper bounds are established, their asymptotic behavior for fixed m and for fixed $n-m$ is analyzed, and their exact values are determined for small n and m. Most of the methods employed show how to construct sets of faults attaining the bounds. An extensive survey of related work is also included, showing connections to resource allocation, k-independent sets, and exhaustive testing. © 1993 Academic Press, Inc.

1. INTRODUCTION

An *n-dimensional hypercube computer*, or n-cube, is a parallel computer with 2^n processors and network topology that of an n-dimensional binary cube. Each node of the cube is associated with a processor P while each edge (P_i, P_j) of the cube represents the direct communication link between processors P_i and P_j. Hypercube computers have been studied since 1962 [35] and have recently become the focus of intense commericial and research activity [15–19].

* Present address: Los Alamos National Laboratory, Los Alamos, New Mexico 87545.

† Research partially supported by ONR grant N00014-90-J-1860.

‡ On leave from Southern Illinois University, Edwardsville, Illinois 62026.

§ Research partially supported by NSF grant CCR-8808839.

¶ Research partially supported by NSF Grant DCR-8507851, NSF/DARPA Grant CCR-9004727, and an Incentives for Excellence award from Digital Equipment Corporation.

0890-5401/93 $5.00

One of the attractive features of the n-cube topology is its behavior in the presence of faulty processors or links. Depending on the number and location of these faults it is possible that the network still contains large subcubes which are fault-free. Since most algorithms for the n-cube specify the dimension of the network as a parameter, these algorithms can still be used in the presence of faults, although with some degradation. Assuming some minimum acceptable level of degradation, it is natural to consider the following question:

> In an n-dimensional hypercube, what is the minimum number of faulty processors (or faulty links) that cause all m-dimensional subcubes to be faulty?

This question can also be considered as part of the subcube allocation problem. In multitasking on an n-cube, the problem of dynamically assigning subcubes of a given dimension to a given task can be thought of as allocating subcubes in the presence of faults, where the busy processors and dedicated communication links can be considered "faulty".

The above question arises from problems in resource distribution [28] as well. To illustrate, suppose disks are to be attached to some of the processors of an n-cube in such a way that every m-dimensional subcube contains a processor with a disk. (We may, for example, be in a multiuser environment and want to ensure that each user has a disk in their allotted subcube.) For a given n and m, the minimum number of disks necessary is the same as the minimum number of faulty processors needed to guarantee that every m-cube is faulty. A solution to this resource distribution problem, however, requires not only the number needed, but also a construction of a minimum set of nodes of Q_n that has a node in common with each m-dimensional subcube.

In order to facilitate our discussion we need to introduce some notation. Let Q_n denote a labeled n-dimensional binary cube, where the nodes of Q_n are all the n-bit strings and two nodes are adjacent if and only if their corresponding strings differ in exactly one position. Define $\mathscr{S}(n, m)$ as the collection of all sets of nodes of Q_n whose removal leaves no Q_m, and let $\kappa(n, m)$ be the minimum size of a set in $\mathscr{S}(n, m)$. Analogously, $\mathscr{T}(n, m)$ denotes the collection of all sets of edges of Q_n whose removal from Q_n leaves no Q_m and $\lambda(n, m)$ is the minimum size of a set in $\mathscr{T}(n, m)$. When the context is clear, the informal term "fault set" will be used to mean a set in $\mathscr{S}(n, m)$ or a set in $\mathscr{T}(n, m)$. Figure 1 illustrates minimum node and edge fault sets for $n=4$ and $m=2$.

There are many alternative methods of fault tolerance not measured by the κ and λ functions. Two basic graph-theoretic approaches are to provide additional edges and/or nodes, or to weaken the notion of a subcube. In the former, hardware is added so that the system still has a Q_n as a subsystem after a fault occurs [8, 33, 36]. This approach must be taken at the time of hardware design, and can tolerate relatively few faults without inordinate expense. In the latter approach, the notion of edge is weakened to allow paths of length greater than one in order to route around faults. This implementation is via software, perhaps together with hardware modifications to permit use of links to and from faulty processors. Generally, many more faults can be tolerated with this approach and it is frequently possible to provide a reconfigured subcube of the desired size in the presence of several faults [14]. This solution suffers a performance penalty, however, because each communication step in a reconfigured subcube takes longer than a communication step in the original hypercube. Neither of these approaches has yet been implemented in any commercial hypercube, and we will not pursue these methods here.

The fault tolerance approach we analyze assumes no hardware modification, incurs no communication penalties, and can be easily utilized on all

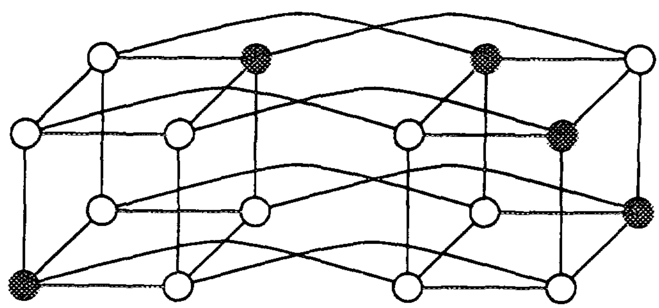

5 faulty nodes destroying every 2-cube

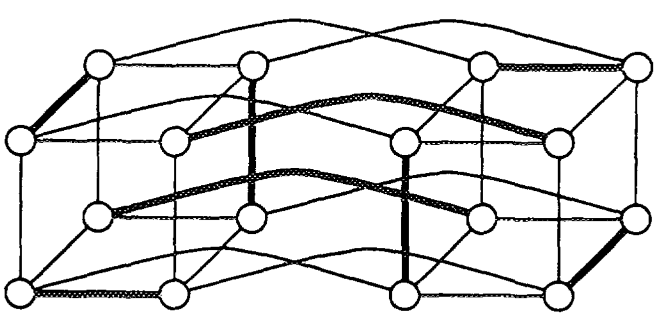

8 faulty edges destroying every 2-cube

FIG. 1. Minimum fault sets for $n = 4$, $m = 2$.

current commerical hypercubes. Further, the κ and λ functions are of interest for a variety of reasons beyond simple fault tolerance. This will be shown below, where ties are exhibited between these functions and problems in resource allocation, exhaustive testing, and κ-independent sets.

1.1. *Prior Work*

A family F of sets is *k-independent* if for every pair of disjoint subsets S_1 and S_2 of F such that $|S_1| + |S_2| = k$ there is at least one element common to all the sets in S_1 but which is in none of the sets in S_2. In Section 2 we show the direct relationship between k-independent sets and κ. The earliest published work relevant to evaluating κ and λ is apparantly that of Schönheim [34] and Brace and Daykin [6], who determined the maximum size of a 2-independent family, and Kleitman and Spencer [26], who considered the general problem of determining the maximum size of families of k-independent sets. Kleitman and Spencer used a probabilistic argument to establish a lower bound for the maximum size of a family of k-independent subsets of a set, proved an upper bound for this maximum, and determined the maximum size of 2-independent sets by constructive means. These results yield the value of $\kappa(n, n-2)$ and bounds for $\kappa(n, n-k)$. Chandra *et al.* [7] studied the problem of finding the minimum number of boolean n-vectors such that every k-projection of them yields all possible k-vectors. In our notation this is $\kappa(n, n-k)$. They determined $\kappa(n, n-2)$, gave a construction for sets in $\mathscr{S}(n, n-3)$ of non-optimal size, and used essentially the same probabilistic argument as in [26] to obtain an upper bound for $\kappa(n, n-k)$. Becker and Simon [3], apparently unaware of the work in [7], repeated many of these results for κ, and used the same methods to establish bounds on λ. They also gave a construction, based on the work of Friedman [12], which yields an upper bound for $\kappa(n, n-k)$ that has the correct growth behavior for fixed $n-k$. In [27], Levitin and Karpovsky considered the problem of exhaustive testing of combinatorial devides with n inputs, where each output is a boolean function of at most k binary input variables. They used MDS codes to construct sets in $\mathscr{S}(n, m)$, although the sets were not of optimal size.

Several persons have worked on a problem complementary to determining $\kappa(n, m)$. Some time ago, Erdős asked for the maximum size of any set of nodes of Q_n for which the induced subgraph contains no 4-cycle. Johnson and Entringer [24] found this maximum size and characterized the extremal graphs for this case. Let $f(n, m)$ denote the maximum size of any set of nodes of Q_n for which the induced subgraph contains no Q_m, and $g(n, m)$ denote the corresponding number for edges. Note that $f(n, m) = 2^n - \kappa(n, m)$ and $g(n, m) = n2^{n-1} - \lambda(n, m)$. Thus, the Johnson and Entringer result determines $\kappa(n, 2)$. In [20, 21, 23], Johnson has considered $f(n, m)$ and obtained bounds for the cases $m = 3, 4, 5$, and in

[22] has evaluated $g(5, 2)$. Responding to a related question of Erdős [9] (see Section 3.4), F. Chung (personal communication, July 1988) established an upper bound for $g(n, 2)$, thus providing a lower bound for $\lambda(n, 2)$.

1.2. *Organization*

The following sections contain our new results on κ and λ as well as an extensive survey of related work. In view of the fault-tolerance applications on the one hand, and the exhaustive testing and resource distribution applications on the other, we address both the problem of determining the values of κ and λ and the problem of the construction of small fault sets. In Section 2 we derive several bounds for κ. We establish new bounds for the maximum size of 3-independent families by using the non-constructive methods of Erdős *et al.* [10]. These bounds yield an improved upper bound for $\kappa(n, n-3)$.

We also give a construction for small sets in $\mathscr{S}(n, n-3)$ which yields a new recursive inequality for $\kappa(n, n-3)$, producing the best known upper bounds for it with n of any practical size. We make use of the results obtained by Kleitman and Spencer [26] for k-independent subsets to establish a new lower bound for $\kappa(n, m)$.

Many of the techniques of Section 2 are easily modified to give corresponding results for λ. These results are described in Section 3 and include an improved upper bound for $\lambda(n, m)$ for m small relative to n, a new lower bound for $\lambda(n, m)$ that is the best known for n large, and a new lower bound for $\lambda(n, m)$ that is the best when m is small. Here, as in Section 2, all but one of the bounds are established by constructive methods.

The asymptotic behavior of $\kappa(n, m)$ and $\lambda(n, m)$, discussed in Section 4, is not well understood for general n and m. However, the new bounds we establish in Sections 2 and 3 do give new information for the cases when m is small relative to n and when $n-m$ is small. In Section 5, we use a combination of the results of earlier sections together with computer programs to construct optimal or near optimal fault sets, thereby determining exact values or tight bounds for $\kappa(n, m)$, for $0 \leqslant m \leqslant n \leqslant 10$, and $\lambda(n, m)$, for $1 \leqslant m \leqslant n \leqslant 7$. In Section 6 we describe techniques for constructing fault sets when n is large. Section 7 contains a discussion of various related open problems and some generalizations.

Because of the large number of results and techniques in Section 2 and 3, the reader may prefer to initially skim these sections, proceeding to Sections 4, 5, and 6. These latter sections help to put the various inequalities into perspective. The reader may then return to the initial sections for a more careful reading.

Throughout, lg denotes $\log_2$ and ln denotes $\log_e$.

2. The Values of κ

The theorems in this section are organized according to the methods employed in their proofs. Theorem 1 and 2 are proved by quite elementary means. A labeling technique is used to prove Theorem 3, whereas Theorem 4 is proved by the use of level sets, yielding a good upper bound for $\kappa(n, m)$ for fixed m. The results in Theorem 8 through 11 rely on the connection between κ and independent sets mentioned in Section 1. A partitioning technique which can be viewed as an extension of the 2-independent set construction yields a recursive inequality for $\kappa(n, n-3)$. The final theorem of this section uses a construction somewhat related to the partitioning method to establish a second recursive upper bound for $\kappa(n, n-3)$. Whenever we establish a recursive inequality, the proof shows how to combine minimum fault sets for the larger side to get a fault set for the smaller side satisfying the inequality.

For a node q in Q_n, the *weight of q* will denote the number of 1's in its string. Extending our notation of n-bit strings for nodes, we will denote the subcubes of Q_n by strings from $\{0, 1, *\}^n$, where the number of *s in the string is the dimension of the subcube.

2.1. *Elementary Bounds*

The theorems in this section are proved by quite elementary and constructive means.

Theorem 1. *For $n \geqslant 1$,*

(i) $\kappa(n, n) = 1$

(ii) $\kappa(n, n-1) = 2$

(iii) $\kappa(n, 0) = 2^n$

(iv) $\kappa(n, 1) = 2^{n-1}$.

Proof. Parts (i) and (iii) follow directly from the definition of $\kappa(n, m)$.

For (ii), note that at least one node must be removed from each of two disjoint copies of Q_{n-1} in Q_n. Moreover, if we remove any pair of antipodal nodes of Q_n, the remaining graph contains no Q_{n-1}. Thus (ii) holds.

For (iv), let Q' and Q'' denote two disjoint copies of Q_{n-1} in Q_n, and consider those edges with one node in Q' and the other in Q''. Since at least one node of each of these edges must be removed in order to remove all the Q_1's from Q_n, we must have $\kappa(n, m) \geqslant 2^{n-1}$. On the other hand, if we remove from Q_n all nodes of even weight then no Q_1 can remain since every edge contains exactly one node of even weight. Part (iv) now follows. ∎

In the next theorem we give recursive upper and lower bounds for $\kappa(n, m)$.

THEOREM 2. *For* $n, m \geqslant 1$,

(i) $\kappa(n, m) \leqslant \kappa(n-1, m-1) + \kappa(n-1, m)$.

(ii) $\kappa(n, m) \geqslant \max\{2\kappa(n-1, m), \kappa(n-1, m-1)\}$.

Proof. Let Q' and Q'' be two node-disjoint copies of Q_{n-1} in Q_n.

For (i), let $S_1 \subseteq Q'$, $S_2 \subseteq Q''$ be sets of size $\kappa(n-1, m-1)$, $\kappa(n-1, m)$ in $\mathscr{S}(n-1, m-1)$, $\mathscr{S}(n-1, m)$. Clearly $S_1 \cup S_2$ is in $\mathscr{S}(n, m)$. For (ii), note that at least $\kappa(n-1, m)$ nodes must be removed from each of Q' and Q'' so that no Q_m remains in either $(n-1)$-cube. Thus $\kappa(n, m) \geqslant 2\kappa(n-1, m)$. To prove the second of the implied inequalities in (ii), let S be a set in $S(n, m)$ of size $\kappa(n, m)$ and let S', S'' be the nodes of S in Q', Q''. Denote by T' the set of nodes of Q' that are adjacent to the nodes of S''. If Q' contains an $(m-1)$-cube A' that is disjoint from $S' \cup T'$, then Q'' contains a corresponding $(m-1)$-cube A'' which combines with A' to form an m-cube disjoint from S. Since this contradicts the choice of S, we may conclude that $S' \cup T'$ must contain at least $\kappa(n-1, m-1)$ nodes and, therefore, $\kappa(n, m) \geqslant \kappa(n-1, m-1)$. ▮

Table 1 shows that sometimes the first term on the right side of the inequality in Theorem 2(ii) is the largest (for example, at $n=7$ and $m=2$) and sometimes the second term is the largest (for example, at $n=6$ and $m=4$). Part (ii) of Theorem 2 shows that $\kappa(n, m)$ is strictly increasing in n. Further, given any fault set S in $\mathscr{S}(n, m)$, removal of any single node of S gives a fault set S' in $\mathscr{S}(n, m+1)$, since any $(m+1)$-cube consists of two disjoint m-cubes, at least one of which is still faulty in S'. Therefore $\kappa(n, m)$ is strictly decreasing in m.

The next theorem generalizes part (i) of Theorem 2. Consider the $(n-1)$-dimensional subcubes $A = 0{*}\cdots{*}$ and $B = 1{*}\cdots{*}$ of Q_n. We may visualize Q_n as a 1-cube with "supernodes" A and B, where we label A with 0 and B with 1. Let S_1 be a subset of A whose removal from A leaves no m-cubes, and S_2 a subset of B whose removal from B leaves no $(m-1)$-cubes. Part (i) of Theorem 2 was proved by observing that $S_1 \cup S_2$ is in $\mathscr{S}(n, m)$. As a first step in generalizing this idea, visualize Q_n as a 2-cube with supernodes A_{00}, A_{01}, A_{10}, and A_{11}, where $A_{ij} = ij{*}\cdots{*}$ is an $(n-2)$-cube of Q_n for $i, j \in \{0, 1\}$. Assign label l_{ij} to supernode A_{ij} as follows: $l_{00} = l_{11} = 0$, $l_{10} = 1$, and $l_{01} = 2$. Next, for each $i, j \in \{0, 1\}$, choose a minimum set S_{ij} of nodes of A_{ij} whose removal from A_{ij} leaves no $(m - l_{ij})$-cube. We see that $\bigcup_{i,j \in \{0,1\}} S_{ij}$ is in $\mathscr{S}(n, m)$ and so

$$\kappa(n, m) \leqslant 2\kappa(n-2, m) + \kappa(n-2, m-1) + \kappa(n-2, m-2). \qquad (1)$$

This result is not a consequence of iterating the inequality in part (i) of Theorem 2, for one iteration yields

$$\kappa(n, m) \leqslant \kappa(n-2, m) + 2\kappa(n-2, m-1) + \kappa(n-2, m-2) \tag{2}$$

which is weaker than inequality (1). Figure 2 illustrates this labeling.

THEOREM 3. *Let r be a non-negative integer. Label the nodes of Q_r with integers in the interval $[0, r]$ such that for every j in $0 \cdots r$ each j-cube of Q_r has a node with label at least as large as j. If $l(q)$ is the label of node q in Q_r then*

$$\kappa(n, m) \leqslant \sum_{q \in Q_r} \kappa(n-r, m-l(q)) \tag{3}$$

for $n \geqslant m \geqslant r$.

Proof. For each node $a = a_1 a_2 \cdots a_r$ in Q_r, let $Q_n(a)$ be the $(n-r)$-dimensional subcube of Q_n given by $a_1 \cdots a_r * \cdots *$. Let $S(a)$ be a set of $\kappa(n-r, m-l(a))$ nodes of $Q_n(a)$ whose removal from $Q_n(a)$ leaves no $(m-l(a))$-cube. We claim that the removal of the set

$$S = \bigcup_{a \in Q_r} S(a)$$

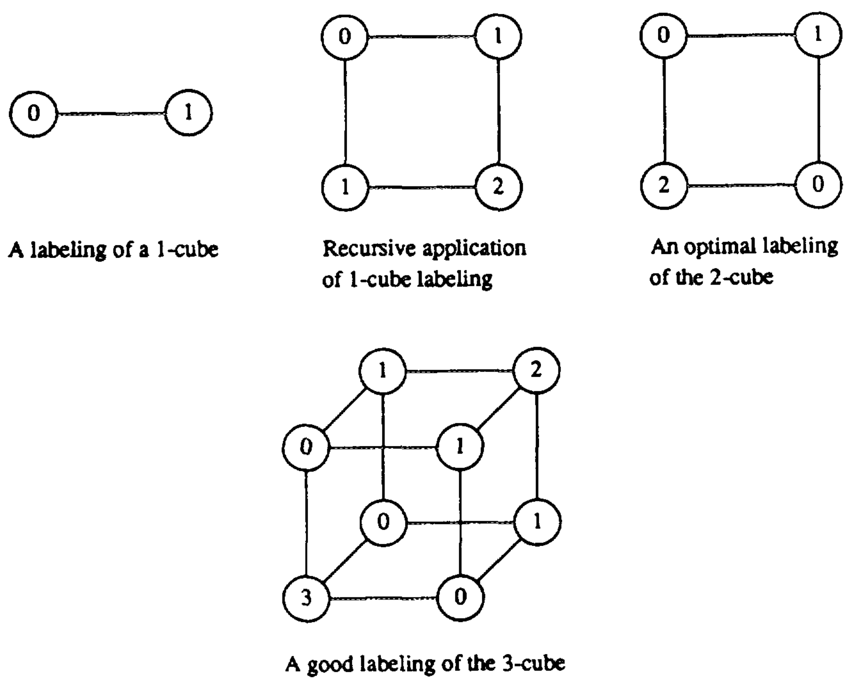

FIG. 2. Labelings of small hypercubes.

from Q_n leaves no m-cube. For suppose T is an m-cube of Q_n, say $T = w_1 \cdots w_n$, where $w_{i(j)} = *$ for $1 \leqslant i(1) < \cdots < i(m) \leqslant n$. Let $t = \max\{j \mid i(j) \leqslant r\}$ and consider the t-dimensional subcube T' of Q_r given by $T' = w_1 \cdots w_r$. (Conceptually, T can be thought of as a product of a t-dimensional subcube of Q_r and an $(m-t)$-dimensional subcube of Q_{n-r}.) By our assumption on the labeling of Q_r, there is some node $v \in T'$ whose label $l(v)$ is at least t. Thus, the $(n-r)$-dimensional cube $Q_n(v)$ has no $(m-l(v))$-dimensional subcube after the removal of $S(v)$. Since $T \cap Q_n(v) = v_1 \cdots v_r w_{r+1} \cdots w_r$ has dimension $m-t$, which is at least $m - l(v)$, this subcube must contain at least one element of S. ∎

In the theorem just proved, if we take $r = 1$ and choose the labels 0 and 1, then inequality (3) reduces to the statement in part (i) of Theorem 2. An iteration of this inequality corresponds to selecting the labels 0, 1 (from 0) and 1, 2(from 1) for Q_2 which yields inequality (2). However, with $r = 2$ and labels 0, 1, 1, 2 assigned to the appropiate nodes of Q_2, we obtain the stronger inequality (1). For each value of r, it is clear that there is a labeling of Q_r which gives an inequality for $\kappa(n, m)$ which is stronger than that supplied by using a labeling obtained by iteration corresponding to a smaller value of r.

The results expressed in Theorem 3 are most useful in the construction of near optimum sets in $\mathscr{S}(n, m)$ based on good constructions for near optimum sets in $\mathscr{S}(n, m-j)$ for $0 \leqslant j \leqslant r$ for some $r \leqslant m$. In applying Theorem 3, the actual choice of r will be determined by what is known about the optimum or near optimum sets in $\mathscr{S}(n, m-j)$ for $0 \leqslant j < m$. In addition, since determining optimum labelings for Q_r for large r is a challenging combinatorial problem in itself, usually only near optimum labelings would be available. Consider, for example, the following construction. For each j in 0, ..., r, pick a set S_j of $\kappa(r, j)$ nodes of Q_r that contains $0 \cdots 0$ and whose removal from Q_r leaves no j-cube. Define a labeling h as follows: for q a node of Q_r, let $h(q) = \max\{k \mid q \in S_k\}$. Clearly, for each collection of sets $S_1, S_2, ..., S_r$, the resulting labeling h satisfies the requirements set forth in Theorem 3, but to obtain near optimum labelings by this method, one would want to choose the sets so that S_i overlaps S_j, for $j < i$, as much as possible. Whatever the selection, $0 \cdots 0$ will have label r, and those nodes with label 0 will not be in S_1. In the worst case, we would construct by this method a set R in $\mathscr{S}(n, m)$, where

$$|R| \leqslant \kappa(r, r)\, \kappa(n-r, m-r) + [\kappa(r, 0) - \kappa(r, 1)]\, \kappa(n-r, m) + \sum_{j=1}^{r-1} [\kappa(r, j) - 1]\, \kappa(n-r, m-j).$$

2.2. *Level Sets*

Our next upper bound on κ is established by the simple device of removing all nodes at given distances from the origin $0 \cdots 0$ of Q_n. An example of such a fault set appears in Fig. 1.

THEOREM 4. *If $n, m \geqslant 1$ and a is any integer, then*

$$\kappa(n, m) \leqslant \sum_{k \equiv a \bmod m+1} \binom{n}{k}.$$

Further, this sum is minimized when $a = \lfloor (n-m-1)/2 \rfloor$.

Proof. The nodes of Q_n can be partitioned into *levels*, where level i consists of all nodes of weight i, $0 \leqslant w \leqslant n$. Any m-dimensional subcube of Q_n must include nodes from $m+1$ consecutive levels. Consequently, if all the nodes are removed from at least one level in every set of $m+1$ consecutive levels, then no Q_m will remain. This can be accomplished by removing all nodes whose weights are in a fixed congruence class a modulo $m+1$. Furthermore, we can minimize the number of nodes removed in this way by judicious choice of a. The level size is monotone decreasing away from the center level (or levels, for n odd). Selecting $a = \lfloor (n-1-m)/2 \rfloor$ results in the removal of levels as far from the center level(s) as possible. A straightforward term-by term comparison shows the optimality of this value of a. ▮

While many authors [3, 7, 20–22, 24] utilize the approach of the theorem just proved, most choose to express their result in the following simpler but weaker form.

COROLLARY 4.1. *For $n \geqslant m \geqslant 1$,*

$$\kappa(n, m) \leqslant \frac{2^n}{m+1}.$$

Proof. The desired result follows from the identity

$$\sum_{a=0}^{m} \sum_{k \equiv a \bmod m+1} \binom{n}{k} = 2^n. \quad ▮$$

The bound given by Theorem 4 in the case $k = 2$ is sharp according to the results of Johnson and Entringer [24], who used constructive methods to determine $f(n, 2)$, the complement of $\kappa(n, 2)$. We state their result in terms of κ.

THEOREM 5 [24]. *For $n \geqslant 2$,*

$$\kappa(n, 2) = \lfloor 2^n/3 \rfloor.$$

Before further discussion concerning the use of level sets, let us simplify notation by letting $C(n, m, a) = \sum_{k \equiv a \bmod m+1} \binom{n}{k}$, and setting $C^*(n, m) = \min\{C(n, m, a) : 0 \leqslant a \leqslant m\}$. In [21, 28] it was noted that, for fixed m, $C^*(n, m)$ satisfies a recursive equation, and this was later solved for $m = 3, 4, 5$ in [20, 21, 23]. These results yield upper bounds for $\kappa(n, m)$ for $m = 3, 4, 5$ which are improvements over those provided by Corollary 4.1. We summarize these in the following.

THEOREM 6 [20, 21, 23]. *For $n, m \geqslant 1$,*

(i) $$\kappa(n, 3) \leqslant 2^n/4 - 2^{\lfloor n/2 \rfloor}/2.$$

(ii) $$\kappa(n, 4) \leqslant \begin{cases} 2^n/5 - (2/5)\, L_n & n \text{ odd} \\ 2^n/5 - (1/5)\, L_{n+1} & n \text{ even,} \end{cases}$$

where L_n, the Lucas number, is $[(1 + \sqrt{5})^n + (1 - \sqrt{5})^n]/2^n$.

(iii) $$\kappa(n, 5) \leqslant \begin{cases} 2^n/6 - 3^{\lfloor n/2 \rfloor}/2 + 1/6 & n \text{ odd} \\ 2^n/6 - 3^{\lfloor n/2 \rfloor}/3 + 1/3 & n \text{ even} \end{cases}.$$

Johnson [21] suggested that the bound $C^*(n, m)$ given by Theorem 4 may be sharp, and formally conjectured equality in that case $m = 4$. However, for any fixed $m > 2$, equality between $\kappa(n, m)$ and $C^*(n, m)$ cannot hold for all $n \geqslant m$. For $m = 3$ this follows from the fact that $\kappa(7, 3) = 24$, from Table 1, Section 5, whereas $C^*(7, 3) = 28$. For $m > 3$, we see that equality fails between $\kappa(m + 2, m)$, whose value is given by Theorem 9, and $C^*(m + 2, m)$, whose value is $m + 3$.

For fixed m and large n, $C^*(n, m)$ is the best upper bound known for $\kappa(n, m)$, but it may still be far from optimal, for, as we shall see in Section 4, there is a large gap between $C^*(n, m)$ and the known lower bounds in these cases.

2.3. *Independent Sets*

We now turn to the theory of independent sets to help us in our study of κ. A family F of sets is *k-independent* if for every pair of disjoint subsets S_1 and S_2 of F such that $|S_1| + |S_2| = k$, there is at least one element common to all the sets in S_1 which is in none of the sets in S_2. The following lemma shows the close relationship between k-independent sets and sets in $\mathscr{S}(n, n-k)$. To state it, we first need some additional notation. Let $\mathscr{F}(r, k)$ denote all k-independent sets of subsets of $\{1, ..., r\}$. For any set T of i elements, by the *orderings* of T we mean the set of $i!$ i-tuples which, when viewed as unordered sets, are equal to T. Let $\hat{\mathscr{F}}(r, k)$ denote the set of all orderings of all elements of $\mathscr{F}(r, k)$, and let $\hat{\mathscr{S}}(n, n-k)$ denote the set of all orderings of all elements of $\mathscr{S}(n, n-k)$.

LEMMA 7. *Given positive integers k, r, n, there is a natural bijection between the n-tuples of $\hat{\mathscr{F}}(r, k)$ and the r-tuples of $\hat{\mathscr{S}}(n, n-k)$.*

Proof. Let $F=(F_1, ..., F_n)$ be an n-tuple of $\hat{\mathscr{F}}(r, k)$. F can be used to construct an r-tuple in $\hat{\mathscr{S}}(n, n-k)$ as follows. Let $M=(m_{ij})$ be the $r \times n$ matrix defined by

$$m_{ij} = \begin{cases} 1 & \text{if } \ i \in F_j \\ 0 & \text{otherwise} \end{cases}.$$

Then, for each $1 \leqslant j \leqslant n$, the jth column of M represents the characteristic function of the set F_j. Moreover, for each $1 \leqslant i \leqslant r$, the ith row of M can be associated with the element i of $\{1, ..., r\}$, and also represents a node of Q_n, where the jth entry of the row is the jth bit of the node's label. We denote by S_M the r-tuple of nodes represented by the rows of M, and claim that S_M is in $\hat{\mathscr{S}}(n, n-k)$. To see why this is the case, let $A = a_1 a_2 \cdots a_n$ be an $(n-k)$-cube in Q_n and define $S_1 = \{F_i : a_i = 1\}$ and $S_2 = \{F_i : a_i = 0\}$. Since A is an $(n-k)$-cube, $|S_1| + |S_2| = k$, and since F is k-independent there is at least one element, say x, that is in each set in S_1 and is in none of the sets in S_2. Thus the node represented by row x is in both A and S_M, proving that S_M is in $\hat{\mathscr{S}}(n, n-k)$.

It is clear that the above mapping from n-tuples of $\hat{\mathscr{F}}(r, k)$ to r-tuples of $\hat{\mathscr{S}}(n, n-k)$ is 1–1. To see that it is onto, let S be an r-tuple in $\hat{\mathscr{S}}(n, n-k)$. Create the $r \times n$ matrix M by setting m_{ij} equal to the jth bit of the ith element of S, and construct an n-tuple $F=(F_1, F_2, ..., F_n)$ of subsets of $\{1, ..., r\}$ by interpreting the jth column of M as the characteristic function of the set F_j. We claim that F is k-independent. To prove this, suppose S_1 and S_2 are disjoint subsets of F, where $|S_1| + |S_2| = k$, and J_1, J_2 are their index sets defined by $J_p = \{i : F_i \in S_p\}$ for $p = 1, 2$. Let $B = b_1 b_2 \cdots b_n$ be the $(n-k)$ – dimensional subcube described by

$$b_i = \begin{cases} 1 & \text{if } \ i \in J_1 \\ 0 & \text{if } \ i \in J_2 \\ * & \text{otherwise.} \end{cases}$$

Since S is in $\hat{\mathscr{S}}(n, n-k)$, B must contain at least one element, say the yth element, of S. This means that $y \in F_i$ for each $i \in J_1$ and $y \notin F_i$ for $i \in J_2$, which allows us to conclude that F is k-independent. ∎

The correspondence established in the lemma, used in [3, 7], gives the following result.

THEOREM 8 [3, 7]. *Let $F(r, k)$ denote the maximum size of a k-independent family of subsets of a set of r elements. Then*

$$\kappa(n, m) = \min\{r \mid F(r, n-m) \geq n\}.$$

Schönheim [34], Brace and Daykin [6], and Kleitman and Spencer [26] determined the maximum size of a family of 2-independent sets. Kleitman and Spencer proved that $F(r, 2) = \binom{r-1}{\lfloor r/2 \rfloor - 1}$, observing that this maximum is attained by taking all subsets of size $\lfloor r/2 \rfloor$ that contain a fixed element of X. Using this result and the above theorem, one immediately obtains the following.

THEOREM 9. *$\kappa(n, n-2)$ is the minimum positive integer r such that $\binom{r-1}{\lfloor r/2 \rfloor - 1} \geq n$.*

Chandra *et al.* [7] rediscovered this result and the following corollary, as did Becker and Simon [3].

COROLLARY 9.1 [3, 7]. $\kappa(n, n-2) = \lg n + \frac{1}{2} \lg \lg n + O(1)$, *where the $O(1)$ term is non-negative.*

Kleitman and Spencer also obtained bounds for $F(r, k)$. They proved an upper bound for $F(r, k)$ for $k \geq 3$ [26, inequality (17)], from which we deduce the more convenient but slightly weaker form

$$F(r, k) \leq \frac{1}{2} \left\{ (k-2)!\,2 \binom{r}{p} \Big/ \binom{x}{p} \right\}^{1/(k-2)} + (k-3), \tag{4}$$

where $x = \lfloor r/2^{k-2} \rfloor + 1$ and $p = \lfloor x/2 \rfloor + 1$. When we combine this result with Theorem 8, we obtain the following.

THEOREM 10. *For $n \geq k \geq 3$,*

$$\kappa(n, n-k) \geq \frac{k-2}{H(1/2^{k-1}) - 1/2^{k-2}} \lg(n-k+3) - k \lg k - 2 \lg \lg n,$$

where $H(x) = -[x \lg x + (1-x) \lg(1-x)]$.

At present, the lower bound just obtained is the best known for $\kappa(n, n-k)$ for k fixed and large n, k. When it is rewritten in the slightly weaker form

$$\kappa(n, n-k) \geq 2^{k-1} \left(\frac{k-2}{k-3} + \lg e \right) \lg(n-k+3) - k \lg k - 2 \lg \lg n, \tag{5}$$

it can easily be compared with the improvement gained over the bound from [3]

$$\kappa(n, n-k) \geq 2^{k-2}[\lg(n-k+2) + 0.125 \lg \lg(n-k+2)],$$

which is the result of applying Theorems 2 and 9.

Now, in the other direction, Kleitman and Spencer [26] used a non-constructive probabilistic argument to prove that

$$F(r, k) \geq (1/2)(k!)^{1/k} (2^k/(2^k-1))^{r/k}. \tag{6}$$

When this inequality is combined with Theorem 8, it is straightforward to show that

$$\kappa(n, n-k) \leq -\frac{k}{\lg(1-2^{-k})} \lg n. \tag{7}$$

This inequality, first established in [7] and later in [3], provides the best known upper bound for fixed k and large n, k. It will be discussed further in Section 4.

Using the non-constructive methods of Erdős *et al.* [10], we next derive a new upper bound for $\kappa(n, n-3)$ that, for n large, is superior to any other known bounds. The best upper bound known previously, given by inequality (7) with $k=3$, is $\kappa(n, n-3) \leq 15.571 \lg n$.

THEOREM 11. *For n sufficiently large, $\kappa(n, n-3) < 7.57 \lg n$.*

Proof. Let r be an even positive integer and let X be a set of r elements. We will prove that there is a 3-independent family of subsets of X that contains at least $(1.0959)^r$ elements when r is sufficiently large. From this we will be able to conclude that $\kappa(n, n-3) < \lg n/\lg 1.0959$ for n sufficiently large, which well complete the proof of the theorem.

Let X' be the set of all subsets of X of size $r/2$, and let p be a real number, $0 < p < 1$, whose value will be determined later. Denote by S a random collection of subsets obtained by choosing independently and with probability p each of the subsets in X'. Using S, we form a 3-independent family by successively deleting any set A from S for which there are sets B and C in S that satisfy either

(1) $A \subseteq B \cup C$, or

(2) $B \cap C \subseteq A$.

For a fixed $A \in X'$, let $b_1(A, r)$ denote the number of pairs $(B, C) \in X' \times X'$ for which (1) holds and let $b_2(A, r)$ be defined analogously for (2). Setting $b(r) = \sum_{A \in X'} [b_1(A, r) + b_2(A, r)]$, we see that the expected number of members deleted from S is at most $p^3\binom{r}{r/2} b(r)$. By choosing $p = (2b(r))^{1/2}$, the existence of a 3-independent set with at least $(1/2)(2b(r))^{-1/2}\binom{r}{r/2}$ members can be guaranteed.

We need an upper bound for $b(r)$, but it will suffice to determine an upper bound for $b_1(A, r)$ because $b_1(A, r) = b_2(X' - A, r)$ for $A \in X'$ and $b(r) = 2\sum_{A \in X'} b_1(A, r)$. To this end, suppose A is a given set in X'. The pairs $(B, C) \in X' \times X'$ for which condition (1) holds can be put in one-to-one correspondence with the four-tuples of sets (U_1, U_2, V_1, V_2) which satisfy the set of restricitions

$$\begin{aligned} R\colon\quad & V_1 \subseteq U_1 \subseteq A, \\ & U_2,\ V_2 \subseteq X - A, \\ & |U_1| + |U_2| = r/2, \\ & |V_1| + |V_2| = |U_1|. \end{aligned}$$

To illustrate the intended correspondence, if we are given the pair (B, C) for which condition (1) holds, take $U_1 = A \cap B$, $U_2 = B - A$, $V_1 = U_1 \cap C$, and $V_2 = C - A$. It is straightforward to check that (U_1, U_2, V_1, V_2) does satisfy each condition of R. Conversely, if the four-tuple (U_1, U_2, V_1, V_2) satisfies all of the conditions listed in R, then with $B = U_1 \cup U_2$ and $C = V_1 \cup (A - U_1) \cup V_2$, the pair (B, C) satisfies condition (1). It follows that $b_1(A, r)$ is the number of such four-tuples satisfying the conditions in R. Hence,

$$b_1(A, r) = \sum_{0 \leqslant x \leqslant r/2} \binom{r/2}{x}\binom{r/2}{r/2 - x} \sum_{0 \leqslant y \leqslant x} \binom{x}{y}\binom{r/2}{x - y}.$$

Since the ratio of consecutive terms in the sum for $b_1(A, r)$ is monotone decreasing, the maximum term occurs where this ratio is approximately 1, namely for $x \sim 0.309r$. Using Stirling's approximation, $n! \sim (n/e)^n (2\pi n)^{1/2}$, we find that $b(r) < (3.3302)^r$ and so $(2b(r))^{-1/2}\binom{r}{r/2} > 2(1.0959)^r$ for r sufficiently large. ∎

2.4. *Partitions*

We now introduce another technique for obtaining upper bounds for $\kappa(n, m)$. While the results of this section are asymptotically weaker than those of the previous section, they provide good recursive constructions for small values of n which are not available from the probabilistic arguments employed.

Consider a collection $P_1, P_2, \ldots, P_r$ of partitions of $\{1, 2, \ldots, n\}$ with the following property.

Property $\mathscr{P}(n, k)$: for every pair of disjoint subsets U and V of $\{1, 2, \ldots, n\}$ for which $|U \cup V| = k$, there is some partition P_i

such that no cell of P_i contains both an element of U and an element of V.

A collection of r partitions satisfying property $\mathscr{P}(n, k)$ can be used to construct a k-independent family of size n as well as a set in $\mathscr{S}(n, n-k)$. In view of the correspondence described at the beginning of Section 2.3 between k-independent sets and sets in $\mathscr{S}(n, m)$, it suffices to show how a collection of partitions satisfying property $\mathscr{P}(n, k)$ can be used to construct a set in $\mathscr{S}(n, n-k)$.

Let P be a partition of $\{1, 2, ..., n\}$ with non-empty cells $A_1, A_2, ..., A_p$, and for each i, $1 \leqslant i \leqslant n$, let $\phi(i)$ be the unique integer j for which $i \in A_j$. Further, if $J \subseteq \{1, 2, ..., n\}$, let $\phi(J) = \{\phi(i): i \in J\}$. We use P, and therefore ϕ, to construct a function τ from $\{0, 1\}^p$ to $\{0, 1\}^n$ as follows:

$$\tau(a_1, a_2, ..., a_p) = (a_{\phi(1)}, a_{\phi(2)}, ..., a_{\phi(n)}).$$

That is, for each subset W of $\{1, 2, ..., p\}$, τ maps the characteristic function of W to the characteristic function of $\bigcup_{i \in W} A_i$ as a subset of $\{1, 2, ..., n\}$. Now, suppose $P_1, P_2, ..., P_r$ is a collection of partitions of $\{1, 2, ..., n\}$ satisfying property $\mathscr{P}(n, k)$, where P_i has c_i non-empty cells $A_{i1}, A_{i2}, ..., A_{ic_i}$. Further, let τ_i and ϕ_i be the functions obtained from P_i as described above. For each i, $1 \leqslant i \leqslant r$, choose a minimum size set $S_i \in \mathscr{S}(c_i, c_i - k)$ and let $X_i = \{\tau_i(s): s \in S_i\}$. Since we can always choose a minimum set in $\mathscr{S}(n, m)$ that contains (0, ..., 0), we will do so, and then modify each X_i for $i > 1$ by removing the n-tuple (0, ..., 0). We claim that the resulting set $X = \bigcup_{i=1}^{r} X_i$ is in $\mathscr{S}(n, n-k)$. To prove this, suppose $U = u_1 u_2 \cdots u_n$ is a subcube of Q_n of dimension $n-k$, and let J_1 and J_2 be the index sets determined by

$$u_i = \begin{cases} 1 & \text{if } i \in J_1 \\ 0 & \text{if } i \in J_2 \\ * & \text{otherwise.} \end{cases}$$

We want to show that there is some element of X that is in U. Since property $\mathscr{P}(n, k)$ is satisfied by the collection $P_1, P_2, ..., P_r$, at least one of these partitions, say P_u, is such that none of its cells contains both an element of J_1 and an element of J_2. Thus, we can define the subcube $V = v_1 v_2 \cdots v_{c_u}$ of Q_{c_u} by

$$v_i = \begin{cases} 1 & \text{if } i \in \phi_u(J_1) \\ 0 & \text{if } i \in \phi_u(J_2) \\ * & \text{otherwise} \end{cases}$$

and conclude that V is of dimension $c_u - k$. Furthermore, since $S_u \in \mathscr{S}(c_u, c_u - k)$, there is an element $x = x_1 x_2 \cdots x_{c_u} \in (V \cap S_u)$, which shows that $\tau_u(x) \in (X \cap U)$. We formalize this result in terms of κ in the following.

THEOREM 12. *Let $n \geqslant k \geqslant 1$. Suppose $P_1, P_2, ..., P_r$ is a collection of partitions of $\{1, 2, ..., n\}$ satisfying property $\mathscr{P}(n, k)$, where P_i has c_i non-empty cells for $1 \leqslant i \leqslant r$. Then*

$$\kappa(n, n-k) \leqslant \sum_{i=1}^{r} \kappa(c_i, c_i - k) - r + 1.$$

This inequality is very useful when a small collection of partitions satisfying property $\mathscr{P}(n, k)$ can be found, as illustrated in the following for $k = 3$.

COROLLARY 12.1. *For all integers s and t such that $st \geqslant n \geqslant s \geqslant t \geqslant 3$,*

$$\kappa(n, n-3) \leqslant 2\kappa(s, s-3) + \kappa(t, t-3) - 2.$$

Proof. Choose the integers s and t in the given range. Let P_1 denote the partitions of $\{1, 2, ..., n\}$ with cells

$$A_{1j} = \{m : 1 \leqslant m \leqslant n, m \equiv j \bmod s\},$$

for $0 \leqslant j \leqslant s-1$; let P_2 denote the partition with cells

$$A_{2j} = \left\{m : 1 \leqslant m \leqslant n, \left\lfloor \frac{m}{s} \right\rfloor = j\right\},$$

for $0 \leqslant j \leqslant t$; and let P_3 denote the partition with cells

$$A_{3j} = \left\{m : 1 \leqslant m \leqslant n, \left\lfloor \frac{m}{s} \right\rfloor + m \equiv j \bmod s\right\},$$

for $0 \leqslant j \leqslant s-1$. It is relatively easy to check that this collection of partitions P_1, P_2, P_3 does indeed have property $P(n, 3)$. ∎

Although Corollary 12.1 gives an upper bound for $\kappa((n, n-3)$ which is $O((\lg n)^{\lg 3})$, its results, when combined with Theorem 13 and the exact values of κ in Table 1, actually give a better upper bound than that provided by Theorem 11 for $n \leqslant 1600$.

Friedman [12] showed how to construct, for any fixed k and n, a collection of $O(\lg n)$ partitions of $\{1, 2, ..., n\}$ such that for any subset T of $\{1, 2, ..., n\}$ of size k, there is a partition in which each of its cells contain at most one element of T. Becker and Simon [3] used Friedman's result to construct sets in $\mathscr{S}(n, n-k)$ of size at most $\lg n(k^4/\lg k)\, 2^{2k \lg k + 3k}$. While this construction yields sets of the right order of magnitude, namely $O(\lg n)$, for small k they are impractically large. For example, if $k=3$ and $n=40$, they are of size $2^{20} \lg 40$, whereas the above corollary with $s=9$ and $t=5$ yields a set of size 32.

The next theorem contains an upper bound for $\kappa(n, n-3)$ which is also established by constructive means. Although its methods yield sets of size $O((\lg n)^2)$, as pointed out earlier, when it is combined with Corollary 12.1 and Table 1, it yields superior bounds for $\kappa(n, n-3)$ for $n \leqslant 1600$. A similar construction was used by Chandra *et al.* [7] to construct sets in $\mathscr{S}(n, n-k)$ of size $O((\lg n)^{k-1})$, but for specific n and k, their sets are somewhat larger than ours because they could not utilize the results in Table 1.

THEOREM 13. *For* $n \geqslant 5$,

$$\kappa(n, n-3) \leqslant \kappa(\lceil n/2 \rceil, \lceil n/2 \rceil - 3) + \kappa(\lceil n/2 \rceil, \lceil n/2 \rceil - 2).$$

Proof. If n is odd then the bound given for $\kappa(n, n-3)$ is the same as the one for $\kappa(n+1, n-2)$. Since part (ii) of Theorem 2 shows that $\kappa(n, n-3) \leqslant \kappa(n+1, n-2)$, it suffices to consider only the case when n is even.

For an arbitrary positive integer m, let x and y be two binary m-bit strings and denote by xy the $2m$-bit string formed by the concatenation of x and y. The complement of x is denoted by x', that is, x' is the string $x'_1 \cdots x'_m$ where $x'_i = 0$ if $x_i = 1$ and $x'_i = 1$ otherwise, $1 \leqslant i \leqslant m$.

Now, let S_1, S_2 denote sets in $\mathscr{S}(n/2, n/2-3)$, $\mathscr{S}(n/2, n/2-2)$, respectively, of minimum size. Define the set S by

$$S = \{xx \mid x \in S_1\} \cup \{yy' \mid y \in S_2\}.$$

To complete the proof, we show that S is in $\mathscr{S}(n, n-3)$.

Suppose i, j, and k are integers in $\{1, ..., r\}$ and $a, b, c \in \{0, 1\}$. Let $T = T(r; i, j, k: a, b, c)$ denote the $(r-3)$-cube of Q_r given by $u_1 \cdots u_r$, where $u_i = a$, $u_j = b$, $u_k = c$, and $u_h = *$ for all $h \neq i, j, k$, $1 \leqslant h \leqslant r$. Analogously, we denote by $W(r; i, j: a, b)$ the $(r-2)$-cube of Q_r given by $u_1 \cdots u_r$, where $u_i = a$, $u_j = b$, and $u_h = *$ for all $h \neq i, j, k$, $1 \leqslant h \leqslant r$.

Let T be an $(n-3)$-cube of Q_n, say $T = T(n; i, j, k: a, b, c)$, where we may suppose, without loss of generality, that $i < j < k$. We divide the proof into cases and show, in each case, that T contains an element of S. The details of the proof in the case where both $i \leqslant n/2$ and $n/2 < j$ are omitted since they are handled almost identically to those listed below.

• $k \leqslant n/2$ or $n/2 < i$. If $k \leqslant n/2$, then $T(n/2; i, j, k: a, b, c) \cap S_1 \neq \phi$, whereas if $n/2 < i$, then $T(n/2; i-n/2, j-n/2, k-n/2: a, b, c) \cap S_1 \neq \phi$. In either case, if x is an element of this intersection, then $xx \in T \cap S$.

• $j \leqslant n/2, n/2 < k$, and $k - n/2 \neq i, j$. Since $T(n/2; i, j, k-n/2: a, b, c) \cap S_1 \neq \phi$, it follows, as in the previous case, that $T \cap S \neq \phi$.

• $j \leqslant n/2, k-n/2 = i$, and $c = a$ or if $j \leqslant n/2, k-n/2 = j$, and $c = b$. Choose an integer k_1 in $1 \cdots n/2$ other than i, j. Since $T(n/2; i, j, k_1: a, b, 0) \cap S_1 \neq \phi$, we have $T \cap S \neq \phi$.

• $j \leqslant n/2, k-n/2 = i$, and $c \neq a$ or if $j \leqslant n/2, k-n/2 = j$, and $c \neq b$. Since $W(n/2; i, j: a, b) \cap S_2 \neq \phi$, if y is an element of this intersection, then $yy' \in T \cap S$. ▮

The proof of the above theorem can be extended to show that for all $n \geqslant 2h \geqslant 2$,

$$\kappa(n, n-h) \leqslant \kappa\left(\left\lceil \frac{n}{2} \right\rceil, \left\lceil \frac{n}{2} \right\rceil - h\right) + \sum_{i=2}^{h-2} \kappa\left(\left\lceil \frac{n}{2} \right\rceil, \left\lceil \frac{n}{2} \right\rceil - h + i\right)\left(\kappa\left(\left\lceil \frac{n}{2} \right\rceil, \left\lceil \frac{n}{2} \right\rceil - i\right) - 1\right).$$

A slightly weaker result appears in [7], where the factor $\kappa(\lceil \frac{n}{2} \rceil, \lceil \frac{n}{2} \rceil - i) - 1$ in the above summation is replaced by $\kappa(\lceil \frac{n}{2} \rceil, \lceil \frac{n}{2} \rceil - i)$.

3. The Values of λ

Turning to the corresponding questions involving edge faults instead of node faults, we find that many of the results and proof techniques for κ have their analogs for λ. Once again, proofs of recursive bounds show how to construct small fault sets. By a slight abuse of notation, we will use $a_1 \cdots a_i {*} a_{i+2} \cdots a_n$ to denote the edge of the 1-cube as well as the 1-cube itself. When we speak of removing the edge $a_1 \cdots a_i {*} a_{i+2} \cdots a_n$ from Q_n, we remove the edge but not the nodes to which it is incident.

3.1. *Elementary Bounds*

Theorem 14. *For $n \geqslant 1$,*

(i) $\lambda(n, n) = 1$

(ii) $\lambda(n, n-1) = 3$ *for* $n \geqslant 3$

(iii) $\lambda(n, 1) = n2^{n-1}$.

Proof. Parts (i) and (iii) are immediate. To establish (ii), let Q' and Q'' denote two disjoint copies of an $(n-1)$-cube in Q_n. Clearly, at least one edge must be removed from each of Q' and Q''. Moreover, at least one edge with one endpoint in Q' and the other in Q'' must be removed to prevent an $(n-1)$-cube made up of corresponding $(n-2)$-cubes in Q' and Q''. Thus, $\lambda(n, n-1) \geqslant 3$. To realize this bound for $n \geqslant 3$, take the set of edges of Q_n given by $T = (*0 \cdots 0, 11*1 \cdots 1, 0*10 \cdots 0\}$. It is easy to check that T is in $\mathscr{T}(n, n-1)$. ∎

Recursive upper and lower bounds corresponding to Theorem 2 are now established for λ.

THEOREM 15. *For* $n \geqslant m$,

(i) $\lambda(n, m) \geqslant \max\{(n/m)\, 2^{n-m}, \lceil 2\lambda(n-1, m)\, n/(n-1) \rceil, \kappa(n, m)\}$

(ii) $$\lambda(n, m) \leqslant \min \begin{cases} 2\lambda(n-1, m) + \kappa(n-1, m-1), \\ \lambda(n-1, m-1) + \lambda(n-1, m), \\ (n-m+1)\, \kappa(n, m) \end{cases}$$

(iii) *If* $\lambda(n+1, m+1) < n+1$ *then* $\lambda(n, m) \leqslant \lambda(n+1, m+1)$.

Proof. For (i), first note that there are $\binom{n}{m} 2^{n-m}$ copies of Q_m in Q_n and each edge is contained in $\binom{n-1}{m-1}$ of the Q_m's. Thus $\lambda(n, m) \geqslant (n/m)\, 2^{n-m}$.

To show that $\lambda(n, m) \geqslant 2\lambda(n-1, m) + \lceil \lambda(n, m)/n \rceil$, from which the second implied inequality of part (i) follows, let T denote a set of minimum size in $\mathscr{T}(n, m)$. There exist at least $\lceil \lambda(n, m)/n \rceil$ parallel edges in T, and without loss of generality, we may suppose these are parallel to $*0 \cdots 0$. The desired inequality now follows from the observation that the two node-disjoint cubes of dimension $(n-1)$ given by $0* \cdots *$ and $1* \cdots *$ must each contain at least $\lambda(n-1, m)$ edges of T.

To see that $\lambda(n, m) \geqslant \kappa(n, m)$, observe that if T is a set of size $\lambda(n, m)$ in $\mathscr{T}(n, m)$ then the set $S = \{v \mid \{v, w\} \in T$ and $\text{weight}(v) < \text{weight}(w)\}$ is in $\mathscr{S}(n, m)$.

In order to show the first of the implied inequalities in (ii), we construct a set T in $\mathscr{T}(n, m)$ as follows. Let Q', Q'' be node-disjoint $(n-1)$-cubes of Q_n. Choose sets T_1, T_2 each of $\lambda(n-1, m)$ edges from Q', Q'', respectively, whose removal from Q', Q'' leaves no Q_m. Further, choose a set of S of $\kappa(n-1, m-1)$ nodes of Q' whose removal from Q' leaves no Q_{m-1}, and let T_3 be the set of edges of Q_n with one endpoint in S and the other in Q''. It is straightforward to verify that $T = T_1 \cup T_2 \cup T_3$ is in $\mathscr{T}(n, m)$. Thus, $\lambda(n, m) \leqslant 2\lambda(n-1, m) + \kappa(n-1, m-1)$.

The inequality $\lambda(n, m) \leqslant \lambda(n-1, m-1) + \lambda(n-1, m)$ can be proved in the same way as the corresponding inequality for κ in Theorem 2.

To prove the last of the implied inequalities in (ii), choose a set of nodes S of size $\kappa(n, m)$ in $\mathscr{S}(n, m)$ and let

$$T = \{\{u, v\} \mid \{u, v\} \text{ is an edge of } Q_n,\ u \in S,$$
$$v \text{ has the same first } m-1 \text{ components as } u\}.$$

Since any m-cube of Q_n contains at least one node in S, it will contain at least one edge of T. Thus, $T \in \mathscr{T}(n, m)$ and $|T| \leqslant (n-m+1)\,\kappa(n, m)$.

For the proof of (iii), suppose for some n that $\lambda(n+1, m+1) < n+1$, and let T be a set of minimum size in $\mathscr{T}(n+1, m+1)$. By our assumption on n, there is some j, $1 \leqslant j \leqslant n+1$, for which no element of T has a $*$ for its jth component. If we project T on this component we see that the resulting set is in $\mathscr{T}(n, m)$. ∎

In part (i) of Theorem 15, each of the first two terms providing a lower bounds for $\lambda(n, m)$ is larger than the remaining two terms for certain values of n and m. For $m = 1$, $(n/m)\,2^{n-m} = \lambda(n, m)$, and for $n = 7$ and $m = 4$, the term $\lceil 2\lambda(n-1, m)\,n/(n-1) \rceil$ gives the best bound. We have not found an example for which the third term, $\kappa(n, m)$, exceeds the other two, but neither have we been able to prove that it is always at most the maximum of the other two. In the inequality occurring in part (ii) of the above theorem, we find that for $m = 1$, the first and third terms are equal to $\lambda(n, 1)$, whereas for $n = 7$ and $m = 5$, the second term is less than the other two. We have not found an example for which the third term, $(n-m+1)\,\kappa(n, m)$, is less than the other two, nor have we been able to show that it is always at least as the minimum of the other two. The example $8 = \lambda(4, 2) < \lambda(3, 1) = 12$ shows that, unlike the corresponding inequality for κ, the conclusion of part (iii) of our theorem does not hold for all $n \geqslant m$. Figure 1 illustrates $\lambda(4, 2)$.

We state two straightforward consequences of Theorem 15 which were also observed in [3].

COROLLARY 15.1. *For* $n \geqslant 3$,

(i) $\kappa(n, n-2) \leqslant \lambda(n, n-2) \leqslant \kappa(n-1, n-3) + 6$

(ii) $\lambda(n, n-2) = \lg n + \frac{1}{2} \lg \lg n + O(1)$.

The labeling technique used in Theorem 3 has an analog for $\lambda(n, m)$.

THEOREM 16. *Let r be a fixed integer, $0 \leqslant r \leqslant m$, and let $r' = \min\{r, m-1\}$. Label the nodes of Q_r with integers in $[0, r]$ and label some subset E_r of the edges of Q_r with integers in $[0, r']$ in such a way that every l-cube of Q_r has either a node or an edge whose label is at least l,*

$0 \leqslant l \leqslant r$. If $l(q)$ is the label of node q in Q_r and $l'(e)$ is the label of edge e in E_r then

$$\lambda(n, m) \leqslant \sum_{q \in Q_r} \lambda(n-r, m-l(q)) + \sum_{e \in E_r} \kappa(n-r, m-l'(e)).$$

Proof. We construct a set T of edges of Q_n as follows. For each $q = q_1 \cdots q_r \in Q_r$, choose a set T_q of $\lambda(n-r, m-l(q))$ edges of the $(n-r)$-cube $Q_n(q) = q_1 \cdots q_r * \cdots *$ whose removal from $Q_n(q)$ leaves no $(m-l(q))$-cube. Further, if $e = u_1 \cdots u_i * u_{i+2} \cdots u_r$ is an edge in E_r, choose a set S_e of $\kappa(n-r, m-l'(e))$ nodes of the $(n-r)$-cube $Q_n(u_1 \cdots u_i 0 u_{i+2} \cdots u_r) = u_1 \cdots u_i 0 u_{i+2} \cdots u_r * \cdots *$ whose removal from $Q_n(u_1 \cdots u_i 0 u_{i+2} \cdots u_r)$ leaves no cube of dimension $(m-l'(e))$. Now, let T_e be the set of edges with one endpoint in S_e and the other in $Q_n(u_1 \cdots u_i 1 u_{i+2} \cdots u_r) = u_1 \cdots u_i 1 u_{i+2} \cdots u_r * \cdots *$. It is straightforward to verify that the set $T = (\bigcup_{q \in Q_r} T_q) \cup (\bigcup_{e \in E_r} T_e)$ is in $\mathscr{T}(n, m)$. ∎

3.2. *Level Sets*

We now construct sets in $\mathscr{T}(n, m)$ by removing edges from Q_n whose nodes are at a specified distance from the origin. The size of the sets constructed by this technique are, for fixed m and large n the smallest yielded by any of the known constructions.

THEOREM 17. *If $n \geqslant m \geqslant 1$ and a is any integer, then*

$$\lambda(n, m) \leqslant (n-m+1)\left[\sum_{\substack{k \equiv a \bmod m \\ k < n/2}} \binom{n-1}{k} + \sum_{\substack{k \equiv a+1 \bmod m \\ k > (n+1)/2}} \binom{n-1}{k-1}\right].$$

Further, this sum is minimized when $a = \lfloor (n-m)/2 \rfloor$.

Proof. Consider, as in the proof of Theorem 4, the nodes of Q_n partitioned into levels in which all nodes of weight i compose level i, $0 \leqslant i \leqslant n$. Suppose that in every set of $m+1$ consecutive levels there are two consecutive levels, say level i and $i+1$, in the set from which we have removed each edge that joins a node in level i to a node in level $i+1$. Clearly, no Q_m can remain. We can improve upon this, for if we fix some $n-m+1$ dimensions, we need only remove those edges that join a node in level i to a node in level $i+1$ along these dimensions. Equivalently, we could have removed the edges that join nodes in level i and $i-1$ along these dimensions. To be more explicit, let $N_0(i, j)$, $N_1(i, j)$ denote the set of nodes of Q_n at level i with jth component equal to 0, 1, respectively.

If $0 \leqslant i < n/2$ and $0 \leqslant j \leqslant n-m+1$, let T_{ij} denote the set of edges of Q_n with one endpoint in $N_0(i, j)$ and the other in $N_1(i+1, j)$; if $(n+1)/2 < i \leqslant n$ and $0 \leqslant j \leqslant n-m+1$, let T_{ij} be the set of edges of Q_n with one endpoint in $N_1(i, j)$ and the other in $N_0(i-1, j)$. Further, setting $T_i = \bigcup_{0 \leqslant j \leqslant n-m+1} T_{ij}$, for $0 \leqslant i \leqslant n$, we see that $|T_i| = (n-m+1)\binom{n-1}{i}$. Now, fix some integer a. If we remove the edges in T_i for $i < n/2$ and $i \equiv a \bmod m$ together with the edges in T_k for $k > (n+1)/2$ and $k \equiv a+1 \bmod m$ then no set of $m+1$ consecutive levels can contain an m-dimensional subcube of Q_n. Again, as in Theorem 4, we choose a in order to minimize the number of edges removed in this manner. The value $a = \lfloor (n-m)/2 \rfloor$ ensures that, where possible, we avoid removing edges incident with level $n/2$ when n is even and levels $(n-1)/2$ and $(n+1)/2$ when n is odd. ∎

COROLLARY 17.1. *For $n \geqslant m \geqslant 1$,*

$$\lambda(n, m) \leqslant (n-m+1)\, 2^{n-1}/m.$$

Proof. The desired result is a consequence of the identity

$$\sum_{a=0}^{m-1} \left[\sum_{\substack{k \equiv a \bmod m \\ k < n/2}} \binom{n-1}{k} + \sum_{\substack{k \equiv a+1 \bmod m \\ k > (n+1)/2}} \binom{n-1}{k-1} \right] \leqslant 2^{n-1}.$$ ∎

3.3. *Partitions*

Using techniques similar to those in Section 2.4, a collection of r partitions satisfying property $\mathscr{P}(n, k)$ can be used to construct sets in $\mathscr{T}(n, n-k)$. We merely need to modify the method used to construct sets in $S(n, m)$ by changing τ as follows. If P is a partition of $\{1, 2, ..., n\}$ with non-empty cells $A_1, A_2, ..., A_c$, and if $t = (t_1, t_2, ..., t_c)$ is an edge of Q_c, where $t_e = *$, say, then let $\tau(t_1, t_2, ..., t_c) = (d_1, d_2, ..., d_n)$, where $d_i = t_{\phi(i)}$ if i is not in $\phi^{-1}(e)$, $d_i = *$ if i is the least element in $\phi^{-1}(e)$, and $d_i = 0$ otherwise. The same methods as those used to establish Theorem 12 and Corollary 12.1 can be used to prove the following.

THEOREM 18. *Let $n \geqslant k \geqslant 1$. Suppose $P_1, P_2, ..., P_r$ is a collection of partitions of $\{1, 2, ..., n\}$ satisfying property $\mathscr{P}(n, k)$, where P_i has c_i non-empty cells for $1 \leqslant i \leqslant r$. Then*

$$\lambda(n, n-k) \leqslant \sum_{i=1}^{r} \lambda(c_i, c_i - k) - r + 1.$$

This theorem is very useful when a small collection of partitions satisfying property $\mathscr{P}(n, k)$ can be found, as illustrated in the following for $k = 3$.

COROLLARY 18.1. *For all integers s and t such that $st \geqslant n \geqslant s \geqslant t \geqslant 3$,*

$$\lambda(n, n-3) \leqslant 2\lambda(s, s-3) + \lambda(t, t-3) - 2.$$

The next result is an analog of Theorem 13.

THEOREM 19. *For* $n \geq 6$,

$$\lambda(n, n-3) \leq \lambda(\lceil n/2 \rceil, \lceil n/2 \rceil - 3) + \lambda(\lceil n/2 \rceil, \lceil n/2 \rceil - 2).$$

Proof. We use an argument similar to that in the proof of Theorem 13 except that special consideration is needed for the case n odd.

First we introduce some notation to show how two edges of Q_k will be used to form an edge in Q_{2k}. If $x = x_1 \cdots x_i * x_{i+2} \cdots x_k$ is a 1-cube in Q_k, let

$$x\tilde{x} = x_1 \cdots x_k x_1 \cdots x_i 0 x_{i+2} \cdots x_k$$

and

$$x\tilde{x}' = x_1 \cdots x_k x_1' \cdots x_i' 0 x_{i+2}' \cdots x_k'$$

Thus, $x\tilde{x}$ and $x\tilde{x}'$ are 1-cubes in Q_{2k}.

Suppose $n = 2p$. Choose sets T_1, T_2 of minimum size in $\mathscr{T}(p, p-3)$, $\mathscr{T}(p, p-2)$, repectively and define

$$T = \{x\tilde{x} \mid x \in T_1\} \cup \{y\tilde{y}' \mid y \in T_2\}.$$

The proof that $T \in \mathscr{T}(n, n-3)$ is almost identical to that used for Theorem 13 and so we suppress the details.

Now, suppose $n = 2p - 1$. Let W_1, W_2 be sets in $\mathscr{T}(p, p-3)$, $\mathscr{T}(p, p-2)$, respectively, each of minimum size. We form a set W in $\mathscr{T}(n, n-3)$ in the same way as we formed T in the n even case, except that we project the $(n+1)$-tuples on the last component. That is, if $x = x_1 \cdots x_{n+1} \in Q_{n+1}$ and $P_{n+1}(x) = x_1 \cdots x_n$ we take $W = \{P_{n+1}(x\tilde{x}) \mid x \in W_1\} \cup \{P_{n+1}(y\tilde{y}') \mid y \in W_2\}$. We suppress the details of the proof that W is in $\mathscr{T}(n, n-3)$ as they are straightforward. ∎

3.4. *Lower Bounds for* λ

Since $\lambda(n, m) \geq \kappa(n, m)$, from Theorem 15, various lower bounds for $\lambda(n, m)$ can be derived from the lower bounds for $\kappa(n, m)$. In particular, the new lower bound proved in Theorem 10 gives us the improved lower bound for $\lambda(n, n-k)$ for fixed k and large n which we state below.

THEOREM 20. *For* $n \geq k \geq 3$,

$$\lambda(n, n-k) \geq \frac{k-2}{H(1/2^{k-1}) - 1/2^{k-2}} \lg(n-k+3) - k \lg k - 2 \lg \lg n,$$

where $H(x) = -[x \lg x + (1-x) \lg(1-x)]$.

At present, the lower bound just obtained is the best known for $\lambda(n, n-k)$ for k fixed and large n, k. We write it in the following slightly weaker form to make it easier to see the size of the bound:

$$\lambda(n, n-k) \geqslant 2^{k-1}\left(\frac{k-2}{k-3} + \lg e\right) \lg(n-k+3) - k \lg k - 2 \lg \lg n. \quad (8)$$

The next theorem gives lower bounds for $\lambda(n, m)$ which, for small m, are better than those available from the inequality in Theorem 20. Its proof is an extension and generalization of an argument used by Johnson [22], who proved that $g(5, 2) \leqslant 56$. Our extension establishes that $g(n, 2) \leqslant k2^{n-1} + b/2$, where k is the integer such that $4\binom{k}{3} \leqslant \binom{n}{3} \leqslant 4\binom{k+1}{3}$, and b is the largest integer such that $(2^n - b)\binom{k}{3} + b\binom{k+1}{3} \leqslant \binom{n}{3} 2^{n-2}$, and further generalizes this to arbitrary $g(n, m)$. F. Chung (personal communication, July 1988) independently proved a result which is essentially the same as our result for $g(n, 2)$, namely that the edge density $x = g(n, 2)/(n2^{n-1})$ must satisfy $(n-1)(n-2) \geqslant 4x(xn-1)(xn-2)$. Thus, for large n, the edge density is bounded above by $(1/4)^{1/3}$. In terms of $\lambda(n, 2)$, we see that, for n large, at least 0.37 of the edges must be faulty in order that every Q_2 is faulty. By Theorems 1 and 15(ii), at most $\frac{1}{2}$ of the edges need be faulty to ensure that every Q_2 is faulty. Some time ago, Erdős [9] conjectured that, for every $\varepsilon < 0$, there is an n_ε such that, for all $n > n_\varepsilon$, $g(n, 2) < (\frac{1}{2} + \varepsilon) n2^{n-1}$, i.e., that the edge density becomes arbitrarily close to $\frac{1}{2}$. He also conjectured that $g(n, k) < cn^{a_k}2^n$, where $a_k < 1$ and $a_k \to 0$ as $k \to \infty$.

THEOREM 21. *For $n \geqslant m \geqslant 1$, let $g(n, m)$ be the largest number of edges in a subgraph of Q_n that contains no Q_m. Then*

$$g(n, m) \leqslant k2^{n-1} + b/2,$$

where k is the integer determined by

$$2^{m+1}\binom{k}{m+1} \leqslant (2^{m+1} - 6)\binom{n}{m+1} < 2^{m+1}\binom{k+1}{m+1}$$

and

$$b = \left\lfloor \left((2^{m+1} - 6)\binom{n}{m+1} 2^{n-m-1} - \binom{k}{m+1} 2^n\right) \Big/ \binom{k}{m} \right\rfloor.$$

Proof. We call a node in Q_n together with its n incident edges an n-star, and refer to the node as its center. We first note that any induced subgraph of Q_{m+1} with at least $2^{m+1} - 5$ of the $(m+1)$-stars must contain a Q_m. For suppose H is an induced subgraph of Q_{m+1} that is lacking only five

$(m+1)$-stars. Split H into two node-disjoint subgraphs H_0 and H_1, where the nodes of H_i have first coordinate i, $i=0, 1$. If either H_0 or H_1 lacks only one $(m+1)$-star, then it must be a Q_m. Without loss of generality, suppose H_0 lacks only two $(m+1)$-stars, centered at nodes p_0 and q_0. Then p_0 and q_0 must be adjacent, for otherwise H_0 would be a Q_m. Let A_0 denote an $(m-1)$-cube of H_0 that does not contain nodes p_0 and q_0 and let A_1 be its neighbor in H_1. Since A_1 must be missing at least one edge, it must contain at least two of the centers, say r and s, of the missing $(m+1)$-stars of H_1 and these must be adjacent. Now, there are at least two node-disjoint $(m-1)$-cubes of H_0, say B_0 and C_0, with p_0 a node of B_0 and q_0 a node of C_0. At least one of their neighbors in H_1, B_1 or C_1, contains all of its possible $(m+1)$-stars and so will form an m-cube with its neighbor in H_0.

Let E denote a set of edges of Q_n, let G denote the subgraph of Q_n induced by E, and suppose G contains no Q_{m+1}. Since there are $\binom{n}{m+1} 2^{n-m-1}$ cubes of dimension $m+1$ in Q_n, G can contain at most $(2^{m+1}-6)\binom{n}{m+1} 2^{n-m-1}$ of the $(m+1)$-stars. Let x_i denote the number of nodes of G of degree i, for $1 \leqslant i \leqslant n$, then $|E| = 1/2 \sum_{i=1}^{n} i x_i$. We must have

$$\sum_{k=m+1}^{n} \binom{k}{m+1} x_k \leqslant (2^{m+1}-6) \binom{n}{m+1} 2^{n-m-1}.$$

Now, let $M(z) = M(z_1, ..., z_n) = \frac{1}{2} \sum_{i=2}^{n} i z_i$ and consider the problem of maximizing M subject to the following three constraints:

C1: z_i is an integer in $[0, 2^n]$ for $1 \leqslant i \leqslant n$,

C2: $\sum_{i=1}^{n} z_i \leqslant 2^n$,

C3: $\sum_{i=m+1}^{n} \binom{n}{m+1} z_i \leqslant (2^{m+1}-6)\binom{n}{m+1} 2^{n-m-1}$.

Note that if $y = (y_1, ..., y_n)$ satisfies these constraints and if, say, y_i, y_{i+1}, and y_{i+2} are all non-zero, then the n-tuple $y' = (y'_1, ..., y'_n)$ also satisfies these constraints, where $y'_i = y_i - 1$, $y'_{i+1} = y_{i+1} + 2$, $y'_{i+2} = y_{i+2} - 1$, and $y'_j = y_j$ otherwise. Moreover, $M(y') = M(y)$. In view of this property, if $x = (x_1, ..., x_n)$ yields a maximum value for M subject to these constraints, then we may assume without loss of generality that all but at most two of the $x_i's$ are 0, and that these two are consecutive. We see that the integer k given in the statement of the theorem is the integer for which $x_k \neq 0$, and $x_i = 0$ for all $i \neq k$, $k+1$. The theorem now follows from the simpler problem of maximizing $M = (k x_k + (k+1) x_{k+1})/2$ subject to the modified constraints:

C1′: x_k and x_{k+1} are non-negative integers,

C2′: $x_k + x_{k+1} = 2^n$,

C3′: $\binom{k}{m+1} x_k + \binom{k+1}{m+1} x_{k+1} \leqslant (2^{m+1}-6)\binom{n}{m+1} 2^{n-m-1}$. ∎

COROLLARY 21.1. *For $n \geqslant m \geqslant 1$,*

$$\lambda(n, m) \geqslant (n-k)\, 2^{n-1} - b/2,$$

where k is the integer determined by

$$2^{m+1}\binom{k}{m+1} \leqslant (2^{m+1}-6)\binom{n}{m+1} < 2^{m+1}\binom{k+1}{m+1}$$

and

$$b = \left\lfloor \left((2^{m+1}-6)\binom{n}{m+1} 2^{n-m-1} - \binom{k}{m+1} 2^n \right) \Big/ \binom{k}{m} \right\rfloor.$$

4. Asymptotics

What is known about the behavior of κ and λ for large values of n and m falls roughly into two categories: results for $n-m$ fixed and those for m fixed. The successful techniques for studying the case $n-m$ fixed are quite different from those that succeed in the case of m fixed. Moreover, the bounds obtained for fixed $n-m$ are not useful for fixed m, and conversely. In this section we describe the best known bounds for each of these cases and mention several open problems concerning the relative sizes of κ and λ.

Kleitman and Spencer [26] used probabilistic methods to determine bounds for the maximum size of families of k-independent sets. Chandra *et al.* [7] used a probabilistic argument equivalent to that in [26] to prove the following bound on κ. Becker and Simon [3] rediscovered this result, and used similar arguments to establish an upper bound for λ. These bounds are stated in the following.

THEOREM 22. [3, 7]. *For all $n \geqslant m \geqslant 1$,*

$$\kappa(n, m) \leqslant (\ln 2)(n-m)\, 2^{n-m} \lg n$$
$$\lambda(n, m) \leqslant (\ln 2)(n-m)\, 2^{n-m}(n/m) \lg n. \quad \blacksquare$$

Combining these bounds with those given by Theorems 10 and 20, we see that both $\kappa(n, m)$ and $\lambda(n, m)$ are $\Theta(\log n)$ for $n-m$ fixed, but there are significant gaps between these bounds.

Question. For fixed $n-m$, does the limit $\lim_{n \to \infty} \kappa(n, n-m)/\lg n$ exist?

This limit exists for $n-m=2$ by the result on 2-independent sets as does the corresponding limit with λ in place of κ. Another question suggested by the slowly increasing nature of κ along the diagonals is:

Question. For fixed $n-m$, is it true that $\kappa(n+1, m+1)-\kappa(n, m) \leqslant 1$ for all sufficiently large n?

The same question could, of course, be asked for λ as well.

When m is fixed, and m and n are large, Theorem 10 and Corollary 4.1 combine to show that $\kappa(n, m)=\Theta(2^n)$. Analogously, Theorem 20 and Corollary 17.1 show that $\lambda(n, m)=\Theta(n2^n)$, but here, too, there are significant gaps between the upper and lower bounds for both κ and λ. Let

$$\alpha_m = \lim_{n \to \infty} \kappa(n, m)/2^n.$$

We see that α_m exists for all $m \geqslant 0$, for Theorem 2 implies that $\kappa(n, m)/2^n$ is non-decreasing for fixed m. Moreover, the inequalities $2\alpha_{m+1} \geqslant \alpha_m \geqslant \alpha_{m+1}$ follow from this same theorem. The definition of α_m and the results of Theorems 1 and 5 show that

$$\alpha_0 = 1, \quad \alpha_1 = \tfrac{1}{2}, \quad \alpha_2 = \tfrac{1}{3},$$

and that α_m satisfies the inequality

$$\frac{\kappa(n, m)}{2^n} \leqslant \alpha_m \leqslant \frac{1}{m+1}$$

for every n. It would be interesting to know the exact values of α_m for $m \geqslant 3$.

In the case of edges, there are many analogies to the above. Let

$$\beta_m = \lim_{n \to \infty} \lambda(n, m)/(n2^{n-1}).$$

The fact that β_m exists for all $m \geqslant 1$ is a consequence of part (i) of Theorem 15 which shows that the sequence $\lambda(n, m)/n2^{n-1}$ is nondecreasing for fixed m. Theorem 14 shows that $\beta_1 = 1$, and Corollary 17.1 together with the definition of β_m shows that

$$\frac{\lambda(n, m)}{n2^{n-1}} \leqslant \beta_m \leqslant \frac{1}{m}$$

for $m \geqslant 1$. Using Theorem 21, we find $\beta_2 \geqslant 0.37$ and $\beta_3 \geqslant 0.112$. Table 2 gives $\lambda(7, 4) \geqslant 19$, which yields $\beta_4 \geqslant 0.042$. The bound $\beta_5 \geqslant 0.016$ is from Theorem 21. Stronger lower bound results and extensions of the tables of values of κ and λ would be of considerable interest as they can improve our asymptotic estimates as well as yield more information about α_m and β_m.

Considering the relative sizes of κ and λ, we see that $\lambda(n, m)/\kappa(n, m)$ is $\Theta(n)$ when m is fixed. Another question concerning the behavior of these functions along the diagonal is:

Question. If $n-m$ is constant, is it true that $\lambda(n, m)/\kappa(n, m)$ is $\Theta(1)$?

Along the same lines, we have seen that, for $n-m \leqslant 2$, the difference $\lambda(n, m)-\kappa(n, m)$ is bounded, which prompts us to ask the following:

Question. For $n-m$ fixed, is it true that $\lambda(n, m)-\kappa(n, m)=O(1)$?

5. Exact Values

In application to hypercubes, the behavior of κ and λ for relatively small values of n and m is more important than their asymptotic values. We must keep in mind that n represents the dimension of the hypercube and so values of $n \geqslant 50$, say, represent a hypercube with more than a quadrillon processors! Consequently, in most applications, the exact values and constructive bounds that yield good approximations for $n<50$ are most useful.

Values of $\kappa(n, m)$ for $0 \leqslant m \leqslant n \leqslant 10$ are represented in Table 1, where exact values are given if known, and otherwise lower and upper bounds are given in the form lower-upper. The exact values for $m=0, 1, n-1$, and n follow from Theorem 1, the values of $\kappa(n, n-2)$ are from Theorem 9, and the $\kappa(n, 2)$ values are obtained from Theorem 5. A computer program using a greedy heuristic was developed to construct small sets S in $\mathscr{S}(n, m)$. To a partially constructed set S, the program randomly adds a node to S that is in the largest number of remaining fault-free m-cubes. This program found sets that resulted in the upper bounds for $\kappa(9, 4)$, $\kappa(10, 5)$, $\kappa(10, 6)$, and $\kappa(10, 7)$ shown in the table. In the remaining cases, the upper and

TABLE 1

Values of $\kappa(n, m)$

	m										
n	0	1	2	3	4	5	6	7	8	9	10
0	1										
1	2	1									
2	4	2	1								
3	8	4	2	1							
4	16	8	5	2	1						
5	32	16	10	6	2	1					
6	64	32	21	12	6	2	1				
7	128	64	42	24	12	6	2	1			
8	256	128	85	48-56	24	12	6	2	1		
9	512	256	170	96-120	48-64	24	12	6	2	1	
10	1024	512	341	192-240	96-165	48–68	24-25	12-13	6	2	1

TABLE 2

Values of $\lambda(n, m)$

	m						
n	1	2	3	4	5	6	7
1	1						
2	4	1					
3	12	3	1				
4	32	8	3	1			
5	80	24	8	3	1		
6	192	59-64	20-22	8	3	1	
7	448	142–160	47-62	19-20	7	3	1

lower bounds for $\kappa(n, m)$ are determined from part (ii) of Theorem 2 and Theorem 4.

Much less has been determined about the exact values of $\lambda(n, m)$. Table 2 displays values of $\lambda(n, m)$ for $1 \leqslant m \leqslant n \leqslant 7$, showing lower and upper bounds when exact values are not known. The values for $m = 1$, $n-1$, and n follow from Theorem 14. All remaining lower bounds can be obtained from Theorems 21.1 and 15. Upper bounds for $\lambda(n, 2)$ are from part (ii) of Theorem 15, while the remaining upper bounds in the table were found by construction. A computer program analogous to the one for κ was designed to construct small sets in $\mathscr{T}(n, m)$. A separate program was developed to determine that $\lambda(7, 5) = 7$.

(A copy of the sets constructed for κ and λ may be obtained by writing to Quentin F. Stout.)

It would be very useful to extend the table of values of κ and λ both for practical instances and because it would, in turn, yield improvements in known bounds for $\kappa(n, m)$ and $\lambda(n, m)$ not included in the table. However, finding small sets in $\mathscr{S}(n, m)$ and $\mathscr{T}(n, m)$ is computationally very difficult.

6. Constructions

The construction of fault sets that are of nearly minimum size is of interest to saboteurs, to computer architects solving resource allocation problems such as those described in [28], and to persons needing to construct k-independent sets for testing purposes [27]. Unfortunately, finding such sets is a very difficult problem in general. Arguments in [26, 7] show that nondeterministic methods have a high probability of success for n large and $n-m$ fixed. Probabilistic arguments similar to those in [26] were used in [3, 7] to prove that, with high probability, a randomly chosen set of $(\ln 2)(n-m)\,2^{n-m}\lg n$ nodes of Q_n is in $\mathscr{S}(n, m)$. An analogous argument shows that, with high probability, a randomly chosen set of $(\ln 2)(n-m)\,2^{n-m}\,(n/m)\lg n$ edges of Q_n is in $\mathscr{T}(n, m)$ [3].

Levithin and Karpovsky [27] developed constructive methods for a problem equivalent to the study of $\kappa(n, m)$. The problem involves the exhaustive testing of devices with n inputs where each output is a Boolean function of at most k binary input variables. Using MDS codes, they constructed an $r \times n$ binary matrix such that all 2^k possible binary k-vectors appear in each of the k columns, where $r = O(\log^w n)$ and w can be chosen arbitrarily close to 1. Their results give a construction of a set in $\mathscr{S}(n, m)$ of size $O(\log^w n)$, where w can arbitrarily close to 1. Alon [1] has given a construction of a family of k-independent subsets of a set of size r. Through the correspondence between independent sets and elements of $\mathscr{S}(n, m)$ shown in the proof of Theorem 8, this yields a construction of a set in $\mathscr{S}(n, m)$ of size $O((n-m)^{c(n-m)^2}\log n)$ for some constant c about 24. For fixed $n-m$, this size is the same order of magnitude as a minimum set in $\mathscr{S}(n, m)$, but even for $n-m=3$, say, it is more than $3^{216}\lg n$.

As discussed in Section 2.4, Becker and Simon [3] used results of Friedman [12] to construct sets in $\mathscr{S}(n, n-k)$ of size at most $\lg n(k^4/\lg k)\,2^{2k\lg k+3k}$. When $n-m=k$ is fixed, this has the right order of growth, but for small values such as $n-m=3$ and $n=20$, say, this bound is more than $2^{20}\lg 20$. On the other hand, the construction in Theorem 13 yields a set in $\mathscr{S}(20, 17)$ of size 19. Even the construction using level sets, Theorem 4, yields a set of size 40 in this case.

When m is fixed, the constructions in [12] and [3] give sets whose sizes are far from the same order of magnitude as $\kappa(n, m)$. In this case the best constructions for near minimum fault sets are given by the level sets in Theorems 4 and 17. It would certainly be of interest to find constructions of sets in $\mathscr{S}(n, m)$ of size $\sim\kappa(n, m)$ and corresponding sets in $\mathscr{T}(n, m)$ of size $\sim\lambda(n, m)$ for m fixed.

To construct small fault sets for practical sizes of n and m, the best strategy is to employ the constructive methods that led to the recursive inequalities of Sections 2 and 3 coupled with the computational results that led to Tables 1 and 2.

7. Concluding Remarks

Our analysis of subcube fault-tolerance assumes that it is sufficient to find an arbitrary fault-free m-dimensional subcube. However, the problem of determining a fault-free subcube of a given dimension is computationally intensive, so in practice the allocation routines examine the availability of only a certain subset of the subcubes of a given dimension. Most allocation schemes use some variant of the "buddy system" allocating only m-cubes of the form $a_1 \cdots a_{n-m} * \cdots *$ [32].

Under a given allocation scheme $\mathscr{A}$, let $\mathscr{A}Q_n$ denote the set of all subcubes of Q_n that are recognized by $\mathscr{A}$. A natural extension of $\kappa(n, m)$ is to $\kappa(\mathscr{A}; n, m)$, which we define as the least number of nodes that need to be removed from Q_n so that the resulting graph contains no m-cube in $\mathscr{A}Q_n$. We define $\lambda(\mathscr{A}; n, m)$ in an analogous way. As an example, if $\mathscr{B}$ denotes the buddy allocation scheme then $\mathscr{B}Q_n = \{a_1 \cdots a_r * \cdots * \mid r = 0, ..., n\}$, and it is easy to check that $\kappa(\mathscr{B}; n, m) = \lambda(\mathscr{B}; n, m) = 2^{n-m}$ for $n \geqslant m \geqslant 1$. While the buddy system is the only allocation scheme used on hypercube computers thus far, we see it is not particularly fault-tolerant. For some specific allocation schemes of interest, Livingston and Stout [29] determined $\kappa(\mathscr{A}; n, m)$. For arbitrary allocation scheme $\mathscr{A}$, Becker and Simon [3] showed that the problem of determining $\kappa(\mathscr{A}; n, n-2)$ is equivalent to a graph-coloring problem. The general problem of determining $\kappa(\mathscr{A}; n, m)$ and $\lambda(\mathscr{A}; n, m)$ is open.

The fault-intolerance questions considered here can be generalized to arbitrary architectures and arbitrary graph properties. That is, given a graph G which represents the connectivity of the processors, how tolerant is G to retaining some specific graph property P under the removal of successive copies of a subgraph H? Here, we define the quantity $\kappa(P, H; G)$ as the minimum number of copies of H whose deletion from G leaves the resulting graph without property P. For example, suppose G is Q_n, P is the property of being connected, and H is a single edge; then $\kappa(P, H; Q_n) = n$. If G is Q_n, P is the property of containing an m-cube, and H is a single m-cube, then $\kappa(Q_m, Q_m; Q_n)$ is the mispacking number $\mathrm{mispac}_0(Q_m \subset Q_n)$ discussed in [13]. As a final example along these lines, consider the problem, described in [25], due to Yuzvinski: How many nodes of the n-cube must be removed in order that no connected component of the rest contains an antipodal pair of nodes? Kleitman [25] solved this problem by establishing the more general result that at least $\binom{n}{\lfloor n/2 \rfloor}$ nodes must be removed from Q_n if no connected component of the remaining graph is to contain more than 2^{n-1} nodes.

In the generalized problem considered above, asking for the minimum number of copies of H whose removal from G destroys P is appropiate in an adversarial situation, in certain resource allocation problems [28], in designing efficient test [27], or in constructing k-independent sets [26]. However, suppose each copy of H to be removed is selected uniformly and at random from the set of all copies of H in G. A natural question that arises is What is the expected number of copies of H that must be removed from G so that the resulting graph fails to have property P? Consider, for example, the case in which G is Q_n, H is a single node, and P is the property of containing an $(n-1)$-cube. In contrast to $\kappa(n, n-1) = 2$ we find that its expected value, denoted $\kappa_E(n, n-1)$, is $\Theta(\log n)$. Some of the properties of $\kappa_E(n, m)$ and and $\lambda_E(n, m)$ for arbitrary n and m, and of $\kappa_E(\mathscr{A}; n, m)$ and $\lambda_E(\mathscr{A}; n, m)$ for certain allocation schemes $\mathscr{A}$, are studied in [29, 30]. A related but somewhat different situation arises if we are only concerned that, with high probability, G fails to have property P. What is the expected number of copies of H that must be removed in this case? Becker and Simon [3] considered an instance of this question in which G is Q_n, H is a single node, and P denotes the property of containing an m-cube. They showed that if at least $(n-m)\,2^{n-m}\lg n$ nodes are removed from Q_n, the probability that there are no remaining m-cubes approaches 1 as n tends to infinity.

A variation of these questions appears in the work of Burton [2], and Erdős and Spencer [11]. Using a probabilistic model of Q_n in which each edge is deleted independently and with fixed probability p, they showed that if $P_1(Q_n, p)$ denotes the probability that the resulting subgraph of Q_n is connected then

$$\lim_{n \to \infty} P_1(Q_n, p) = \begin{cases} 1 & \text{if } p < 1/2; \\ 1/e & \text{if } p = 1/2; \\ 0 & \text{otherwise.} \end{cases}$$

When p is allowed to vary with n, Bollobás [4, 5] proved that if $\mu > 0$ and

$$p = p(n) = 1 - \tfrac{1}{2}\{\mu + o(1)\}^{1/n}$$

then

$$\lim_{n \to \infty} P_1(Q_n, p) = e^{-\mu}.$$

Suppose that instead of deleting edges from Q_n we delete nodes, together with their incident edges, with fixed probability p and define $P_0(Q_n, p)$ as the probability that the resulting subgraph of Q_n is connected. Najjar and Gaudiot [31], investing the reliability of the hypercube network in the presence of node faults, used Monte Carlo simulation to estimate $P_0(Q_n, p)$ for small n. In [30], and analog of the above results for $P_1(Q_n, p)$ is proved for $P_0(Q_n, p)$, namely

$$\lim_{n \to \infty} P_0(Q_n, p) = \begin{cases} 1 & \text{if } p < 1/2; \\ 1/2 & \text{if } p = 1/2; \\ 0 & \text{otherwise.} \end{cases}$$

ACKNOWLEDGMENTS

We thank the referees for their useful and incisive comments, and Karen Johnson and Fan Chung for providing comments and access to their unpublished work.

RECEIVED August 1, 1988; FINAL MANUSCRIPT RECEIVED February 28, 1991

REFERENCES

1. ALON, N. (1986), Explicit construction of exponential sized families of k-independent sets, *Discrete Math.* **58**, 191–193.
2. BURTON, Y. (1977), On the probability of connectedness of a random subgraph of the n-cube, *Problemy Peredachi Informatsii* **13**. [Russian–English summary]
3. BECKER, B., AND SIMON, H. (1988), How robust is the n-cube? *Inform. and Comput.* **77**, 162–178.
4. BOLLOBÁS, B. (1983), The evolution of the cube, *in* "Combinatorial Mathematics" (C. Berge, *et al.*, Eds)., pp. 91–97, North-Holland, Amsterdam.
5. BOLLOBÁS, B. (1985), "Random Graphs," pp. 336–346, Academic Press, London.
6. BRACE, A., DAYKIN, D.E. (1972), Sperner type theorems for finite sets, *in* "Proceedings, British Combinatorial Conference, Oxford," pp. 18–37.
7. CHANDRA, A., KOU, L., MARKOWSKY, G., AND ZAKS, S. (1983), On sets of boolean n-vectors with all k-projections surjective, *Acta Inform.* **20**, 103–111.
8. DUTT, S. (1990), "Designing and Reconfiguring Fault-Tolerant Multiprocessor Systems," Ph.D. Thesis, University of Michigan.
9. ERDŐS, P. (1984), Some problems in graph theory, combinatorial analysis and combinatorial number theory, *in* "Graph Theory and Combinatorics" (B. Bollobás, Ed.), 1–17, Academic Press, New York/London.
10. ERDŐS, P., FRANKL, P., AND FÜREDI, Z. (1982) Families of finite sets in which no set is covered by the union of two others, *J. Combin. Theory A* **33**, 158–166.
11. ERDŐS, P., AND SPENCER, J. (1979), Evolution of the n-cube, *Comput. Math. Appl.* **5**, 33–39.
12. FRIEDMAN, J. (1984), Constructing $O(n \log n)$ size monotone formulae for the kth elementary symmetric polynomial of n Boolean variables, *in* "25th Symposium on Foundations of Computer Science," pp. 506–515.
13. HARARY, F., HAYES, J. P., AND WU, H. (1988), A survey of the theory of hypercube graphs, *Comput. Math. Appl.* **15**, 277–289.
14. HASTAD, J., LEIGHTON, T., AND NEWMAN, M. (1987), Reconfiguring a hypercube in the presence of faults, *in* "Proceedings, 19th ACM Symposium, Theory of Computing," pp. 274–284.
15. "Hypercube Multiprocessors 1986," SIAM, Philadelphia.
16. "Hypercube Multiprocessors 1987," SIAM, Philadelphia.
17. "Proceedings, 3rd Conference on Hypercube Concurrent Computers and Applications" (1988), SIAM, Philadelphia.
18. "Proceedings, 4th Conference on Hypercube Concurrent Computers and Applications" (1989), SIAM, Philadelphia.
19. "Proceedings, 5th Distributed Memory Computer Conference" (1990).
20. JOHNSON, K. (1986), Largest Induced Subgraphs of the n-Cube That Contain No $2k$-Cycle, Ph.D. Thesis, University of New Mexico.

21. Johnson, K. (1988) An appearance of the Fibonacci and Lucas numbers in a hypercube problem, preprint, University of New Mexico.
22. Johnson, K. (1988), Buds, stems and fault-tolerance in hypercubes, preprint, University of New Mexico.
23. Johnson, K., McCanna, J., and Grassl, R. (1989), Pascalian rectangles modulo *m*, preprint, University of New Mexico.
24. Johnson, K., and Entringer, R. (1989), Largest induced subgraphs of the *n*-cube that contain no 4-cycles, *J. Combin. Theory Ser. B* **46**, 346–355.
25. Kleitman, D. (1986), On a problem of Yuzvinsky on separating the *n*-cube, *Discrete Math.* **60**, 207–213.
26. Kleitman, D. and Spencer, J. (1973), Families of *k*-independent sets, *Discrete Math.* **6**, 255–262.
27. Levitin, L., and Karpovsky, M. (1986), "Efficient Exhaustive Tests Based on MDS Codes," Technical Report, Boston University.
28. Livingston, M., and Stout, Q. F. (1988), Distributing resources in hypercube computers, *in* "Proceedings, Third Conference on Hypercube Concurrent Computers and Applications," pp. 222–231.
29. Livingston, M., and Stout, Q. F. (1988), Fault tolerance of allocation schemes in massively parallel computers, *in* "Proceedings, 2nd Symposium, Massively Parallel Computation," pp. 491–494.
30. Livingston, M., and Stout, Q. F. (in preparation), Expected fault-tolerance of hypercubes.
31. Najjar, W., and Gaudiot, J.-L. (1987), Reliability and performance modeling of hypercube-based multiprocessors, "2nd International Workshop on Applied Mathematics and Performance/Reliability Models of Computer/Communication Systems, University of Rome II," pp. 305–319.
32. "NCUBE Handbook, Version 1.0" (1986), NCUBE Corp., Beaverton, OR.
33. Rennels, D. A. (1986), On implementing fault-tolerance in binary hypercubes, *in* "Proceedings, 16th Symposium, Fault-Tolerant Computation," pp. 344–349.
34. Schönheim, J. (1971), A generalization of results of P. Erdős, G. Katona, and D. J. Kleitman concerning Sperner's theorem, *J. Combin. Theory Ser. A* **11**, 111–117.
35. Squire, J. S., and Palais, S. M. (1963), Programming and design considerations for a highly parallel computer, *AFIPS Proc.* **23**, 395–400.
36. Yanney, R. M. (1982), "Fault Recovery in Multiprocessor Networks," Ph.D. Thesis, University of Southern California.