Fault tolerant graphs, perfect hash functions and disjoint paths

M. Ajtai * N. Alon † J. Bruck * R. Cypher * C.T. Ho * M. Naor *
E. Szemerédi ‡

Abstract

Given a graph G on n nodes we say that a graph T on $n+k$ nodes is a k-fault tolerant version of G, if we can embed G in any n node induced subgraph of T. Thus T can sustain k faults and still emulate G without any performance degradation. We show that for a wide range of values of n, k and d, for any graph on n nodes with maximum degree d there is a k-fault tolerant graph with maximum degree $O(kd)$. We provide lower bounds as well: there are graphs G with maximum degree d such that any k-fault tolerant version of them has maximum degree at least $\Omega(d\sqrt{k})$.

*IBM Almaden Research Center, 650 Harry Road, San Jose, CA 95120, USA

†Department of Mathematics, Raymond and Beverly Sackler Faculty of Exact Sciences, Tel Aviv University, Tel Aviv, Israel. Research supported in part by a United States Israel BSF Grant.

‡Department of Mathematics, Rutgers University, New Brunswick, NJ, USA

1 Introduction

Given a graph G on n nodes we say that a graph T on $n+k$ nodes is a `k-fault tolerant version` of G, if we can embed G in any n node induced subgraph of T. The embedding should be a one-to-one function, such that any two adjacent nodes in G are mapped into adjacent nodes in the subgraph of T. An equivalent way to phrase this requirement is to say that any k nodes of T can be declared faulty and the remaining graph can embed G in the above manner. Thus, if G represents some desired interconnection, then T is a network that can sustain up to k faults and still emulate G without any performance degradation.

If we take T to be the complete graph on $n+k$ nodes, then we know that T is k-fault tolerant for any G on n nodes, Our goal, however, is to minimize the maximum degree in T. Our results show that this degree can be $O(kd)$ where d is the maximum degree in G.

We present a general method of constructing such fault tolerant versions of graphs. This method depends only on n and k and not on the specific graph. It requires the construction of a directed graph M and a collection of mappings from T to M. M has k distinguished nodes called *spares*. The property that the mappings and M should have is that for any set of k faults in T there should be a mapping for which there are k disjoint paths in M connecting the mapped faulty nodes and the k spares. This method is presented in Section 2. In general, it can construct graphs T with maximum degree $O(kd)$. Sections 3 and 4 are devoted to constructing M and the mappings.

In Section 3 we consider a graph M based on the grid. We show that a good collection of mappings exists based on the existence of a certain family of perfect hash functions. Such a family exists when $k < n^{1/4}$. We provide an *explicit* construction for n and k such that n is $\Omega(k^{\log k})$. For larger k's the explicit construction yields graphs T with maximum degree larger than $O(kd)$.

In Section 4 we consider expander based graphs. We show (non-constructively) that our method works (i.e. yields T's with maximum degree $O(kd)$) for all n and k such that k is $O(n^{1-\beta})$ for some $\beta > 0$.

In Section 5 we show our lower bounds: there are graphs on n nodes with degree d such that any k-fault tolerant version of them has maximum degree at least $\Omega(d\sqrt{k})$. The relationship between n, k and d that must hold in order for the lower bound proof to work are that d be larger than $\log^2 n/\log\log n$ and k be smaller than $c_1\sqrt{d/\log d}$ and $n/2d$, where $c_1 > 0$ is some absolute constant.

1.1 History, motivation and comparison with other work

As mentioned above, the main motivation for the problem of constructing k-fault tolerant graphs is finding fault tolerant architectures. The graph G represents a desired interconnection and the graph T the actual fault tolerant network that implements it. The graph T can sustain any k node (processor) faults and still simulate the graph G with no slowdown. Note that there is no assumption about the distribution of the k faults, and that the redundancy is the lowest possible, as the number of spares equals the number of faults. This approach to fault tolerance has been considered in a number of papers previously. The papers [14, 25, 21, 7, 8, 5, 6] considered the problem for some specific graphs such as cycles, arrays and trees. Rosenberg [23] has provided VLSI area lower bounds for such architectures. Dutt and Hayes [8] and Bruck, Cypher and Ho [5] have considered the problem of constructing k-fault tolerant versions for arbitrary graphs and have shown a simple $O(k^2d)$ construction outlined in Section 2.

There are different approaches for achieving fault tolerance. The approach represented by works such as [13, 15] is for the faulty network to attempt to simulate the original architecture with as little slowdown as possible. However, as there is no redundancy, some slowdown is necessary and when no assumptions are made regarding the distribution of the faults, this slowdown is often non-constant. Other approaches assume that there is some additional switching mechanism that is fault-free (see e.g. [22, 16]) or require that there be many more spare processors than faults (see e.g. [1]), or some combination of the above.

The advantages that k-fault tolerant graphs offer are clear (e.g. minimal redundancy, no slowdown, no assumptions on fault distribution, no requirement for a fault-free switching mechanism). The question is what price must be paid in terms of the degree of the network in order to obtained these advantages. Answering this question is the main goal of the paper.

Finally, we should note that some of the techniques developed in this paper might be applicable elsewhere. The results of Section 4 yield an efficient way to construct superconcentrators from expander graphs. The construction of families of perfect hash functions in Section 3 may have applications in data structures.

1.2 Notation

In the following the input graph will always be $G = (V, E)$ and the k-fault tolerant graph will be $T = (U, D)$. We will also use an intermediate directed graph $M = (W, C)$. Unless stated otherwise, $|V| = n$ and $|U| = |W| = n + k$. Given a mapping $f : U \mapsto V$ and a subset $S \subset U$, $f(S)$ is the set of all the elements of V that are the image of some element of S. For the directed graph $M = (W, C)$ and for a node $w \in W$ let $\Gamma_M(w)$ denote the set of all out-neighbors of w in M *including* w itself.

2 The General Scheme

In this section we describe our general scheme. It involves coming up with an intermediate graph M and a collection of mappings H with certain properties. These properties are a function of n and k only, and do not involve G. Given a graph M with maximum in-degree δ and a collection of mappings H of size ℓ, our scheme constructs for any input graph G with maximum degree d a k-fault tolerant graph T with maximum degree $\ell(\delta+1)^2 d$.

Consider first the method suggested by Bruck, Cypher and Ho [5]. Let $V = \{1,\ldots,n\}$ and $U = \{1,\ldots,n+k\}$. For every edge $(i,j) \in E$ create the following edges in T: completely connect in T the nodes $\{i,\ldots,i+k\}$ to the nodes $\{j,\ldots,j+k\}$. Given the set of k faults the embedding strategy is greedy: scan the nodes of V in increasing order and embed node i in the first available (not occupied and not faulty) node in U. For any set of k or fewer faults, we know that $i \in V$ is embedded in one of $\{i,\ldots,i+k\} \subset U$. Therefore for any edge $(i,j) \in E$ we know that there is an edge between the node where i is embedded and the node where j is embedded. Every edge in E induces $(k+1)^2$ edges in D and the maximum degree in T is $(k+1)^2 d$.

We now discuss the limitations of such a method. Given $G = (V,E)$ and a k-fault tolerant version of it $T = (U,D)$, then for any $v \in V$ let

$$S(v) = \{u \in U | \exists \text{ set of } k \text{ faults such that } v \text{ is embedded in } u\}.$$

In the method above $S(i) = \{i,..,i+k\}$. It is clear that for all $v \in V$ we must have $|S(v)| \geq k+1$, otherwise all of $S(v)$ might be faulty and there will be no place to embed v. If for every $(v_1,v_2) \in E$ we must connect $S(v_1)$ to all of $S(v_2)$, we get that every edge in G induces $(k+1)^2$ edges in T.

Our approach in reducing the degree is to partition the sets $S(v)$ in such a way that if $v_1 \in V$ is mapped to part j in $S(v_1)$, then $v_2 \in V$ is mapped to part j in $S(v_2)$. In order to implement it we choose $M = (W,C)$, a directed graph on $n+k$ nodes, and a one-to-one mapping $m : V \mapsto W$. We call the nodes $w \in W$ with no $m^{-1}(w)$ the *spare nodes*. We also choose a collection H of one-to-one mappings from U to W.

The directed graph M and the collection H now define for any input graph $G = (V,E)$ a graph $T = (V,E)$ that should be a k-fault tolerant version of G. For each edge $(v_1,v_2) \in E$ and each $h \in H$ create the following edges in T: completely connect $h^{-1}(\Gamma_M(m(v_1))$ and $h^{-1}(\Gamma_M(m(v_2))$. Thus

$$D = \bigcup_{h \in H} \bigcup_{(v_1,v_2) \in E} \{(u_1,u_2\} | h(u_1) \in \Gamma_M(m(v_1)) \text{ and } h(u_2) \in \Gamma_M(m(v_2))\}. \quad (1)$$

We say that a set S of k faults is *good* for a mapping $h : U \mapsto W$ if there are k node disjoint paths in M connecting $h(S)$ to the k spare nodes. Given a set S of k faults such that h is good for S we claim that the following $\varphi : V \mapsto U$ embeds G in the non faulty nodes: fix the k node disjoint paths; for all $v \in V$, if $m(v)$ is not on any of the k paths, then $\varphi(v) = h^{-1}(m(v))$. Otherwise, let $w \in W$ be the node next to $m(v)$ on the path (connecting a faulty node with a spare one) in the direction towards the spare and let $\varphi(v) = h^{-1}(w)$. Clearly:

Claim 2.1 *φ as defined above is one-to-one.*

Also:

Claim 2.2 *For all $(v_1, v_2) \in E$ there is an edge $(\varphi(v_1), \varphi(v_2)) \in D$.*

Proof: Since v_1 is embedded in $h^{-1}(\Gamma_M(m(v_1)))$ and v_2 is embedded in $h^{-1}(\Gamma_M(m(v_2)))$ by the definition (1) of D there is an edge in D between $\varphi(v_1)$ and $\varphi(v_2)$. $\Box$

Every $h \in H$ and every edge $(v_1, v_2) \in E$ induces $(deg_M(m(v_1) + 1)(deg_M(m(v_2) + 1)$ edges in T, where here $deg_M(w)$ denotes the out-degree of w in M. Thus we would like to choose as small H as possible so that for any set S of k faults, at least one $h \in H$ is good for S.

If the in-degree of M is bounded by δ and the size of H is ℓ, then we get that the degrees in T are bounded by $\ell(\delta + 1)^2 d$. Therefore we have the following:

Theorem 2.1 *Suppose that there exists a graph $M = (W, C)$ with maximum in-degree δ and H a family of mappings $U \mapsto W$ of size ℓ such that for any set of k faulty nodes in U there is a mapping $h \in H$ such that there are k node disjoint paths in M connecting $h(S)$ to the k spare nodes. Then for any graph $G = (V, E)$ with maximum degree d there is a k-fault tolerant graph $T = (U, D)$ with maximum degree $(\delta + 1)^2 \ell d$.*

Once we have M and H as in Theorem 2.1 the embedding strategy is straightforward: for a given set of faults S, find an $h \in H$ such that h is good for S and construct φ as above. Using a simple observation we can see the limitation of our general scheme:

Theorem 2.2 *Let M, H, δ and ℓ be as in Theorem 2.1, then $(\delta + 1)\ell > k$.*

Proof: Construct a set S as follows: pick a non-spare node $w \in W$ and make

$$S = \bigcup_{h_i \in H} h_i^{-1}(\Gamma_M(w)).$$

The size of S is at most $\ell|\Gamma_M(w)|$. No $h_i \in H$ can have node disjoint paths connecting $h_i(S)$ to the spare nodes, since w will have no non-faulty out-neighbors under any h_i. Therefore if $k \geq |\Gamma_M(w)|\ell$ then H and M do not satisfy the requirements of Theorem 2.1. Because w was an arbitrary non-spare node, it follows that $(\delta + 1)\ell > k$. $\Box$

By Theorem 2.2 it is clear that the strategy in constructing M and H should be to minimize δ as possible, since the actual degree of T will be at least $\delta k d$. In the next section we will see a construction where $\delta = 2$ and ℓ is $O(k)$ that works for $k < n^{1/4}$. A construction with $\delta = O(1/\beta)$

and $\ell k = O(k)$ that works for $k < n^{1-\beta}$ is given in section 4.

3 Grid Based Constructions

We now describe the application of the general scheme that works for k's which are smaller than $n^{\frac{1}{4}}$. It is fully constructive for k such that n is $\Omega(k^{\log k})$.

Let n and k be such that $k \mid n$ and let $M = (W, C)$ be the following graph: $W = \{(a,b) | a \in \{0, \ldots, n/k\}, b \in \{0, \ldots, k-1\}\}$. Let

$$\begin{aligned} C = \quad & \{((a,b),(a-1,b)) | a \in \{1,\ldots,n/k\}, b \in \{0,\ldots,k-1\}\} \\ & \bigcup \{((a,b),(a,b-1 \bmod k)) | a \in \{1,\ldots,n/k\}, b \in \{0,\ldots,k-1\}\}. \end{aligned}$$

We refer to M as a grid and we consider all the nodes with the same a to be a row and all those with the same b to be a column. The nodes V are mapped by m arbitrarily to $\{(a,b) | a \in \{1,\ldots,n/k\}, b \in \{0,\ldots,k-1\}\}$. Therefore the bottom '0' row constitutes the set of spares.

The following claim provides a sufficient condition for the existence of k disjoint paths connecting the faults to the spares.

Claim 3.1 *Given a set $S = \{w_1 = (a_1,b_1), w_2 = (a_2,b_2) \ldots w_k = (a_k,b_k)\}$ of k nodes in W such that all the members of S are in different rows, i.e all the a_i's are different, then there exists a set of k disjoint paths connecting S to the bottom row.*

Proof: The path can be found in a greedy manner, processing the members of S bottom up. Assume that $a_1 < a_2, \ldots < a_k$. Suppose that we have arranged paths connecting the first $i-1$ nodes to the bottom row and we wish to connect w_i. We will maintain the property that the paths for the first $i-1$ nodes use at most $i-1$ different *columns*, i.e. no node from a different column is on any of the paths. Let c_i be the first column to the left of b_i that has not been used (that is, c_i is the vacant row that minimizes $b_i - c_i \bmod k$). The path connecting (a_i,b_i) to the bottom row consists of the segment in row a_i connecting (a_i,b_i) to (a_i,c_i) and the segment connecting in column c_i (a_i,c_i) to $(0,c_i)$. Since a_i is the highest node processed so far, it is clear that all the nodes on $(a_i,b_i)-(a_i,c_i)$ are vacant, and since column c_i has not been used so far, the segment $(a_i,c_i)-(0,c_i)$ is vacant as well. □

Obviously, the condition that all faults are mapped to different rows is not necessary. However, in all subsequent constructions in this section we will try to fulfill it. Therefore from Claim 3.1 we have:

Theorem 3.1 *Let $M = (W,C)$ be as defined above and suppose that there exists a collection H*

of mappings $U \mapsto W$ of size ℓ with the property that for all $S \subset U$ where $|S| = k$ there is an $h \in H$ such that all the nodes $h(S)$ are in different rows, then M and H satisfy the properties of Theorem 2.1.

Therefore we are left with the problem of finding the collection H. This is essentially a perfect hashing problem with the added restriction that the size of the preimages is fixed at k. For more information on perfect hashing, see [20, 11, 10, 3]. Unfortunately, no existing construction from the literature seems to fit our desired parameters.

By (simple) results in perfect hashing it follows that the minimum size of H we can hope for is $\Omega(k)$. We now present the first construction. This is not an explicit construction, but the function does have a succinct and efficient representation. It will serve as the basis for the next construction as well.

3.1 A shift based construction

Let $U = \{(x, y) | x \in \{0, \ldots, n/k\}, y \in \{0, \ldots, k-1\}\}$ and assume that $p = n/k + 1$ is a prime. The functions $h \in H$ are of the form $h((x, y)) = (ay + x \bmod p, y)$ for some $a \in \{1, \ldots, n/k\}$. In order to construct H, choose ℓ elements $a_1, a_2, \ldots a_\ell$ uniformly at random from $\{0, \ldots, n/k\}$ and let $h_i((x, y)) = (a_i y + x \bmod p, y)$. For every pair of distinct elements (x_1, y_1) and (x_2, y_2) if a_i is random, $Prob[h_i((x_1, y_1)) = h_i((x_2, y_2))] = \frac{1}{p} < \frac{k}{n}$ in case $y_1 \neq y_2$ and 0 otherwise. Therefore, we get that for any set $S = \{(x_1, y_1), (x_2, y_2), \ldots, (x_k, y_k)\}$ of k nodes in U the probability that h_i is *not* good for S is less than $\frac{k^3}{n}$. By summing up for all subset $S \subset U$ we get that the probability that there exists a set S such that none of the $h_i \in H$ is good for S is at most

$$\binom{n+k}{k}\left(\frac{k^3}{n}\right)^{\ell}. \tag{2}$$

If $k < n^{1/4}$ and $\ell = 4k$, then (2) is smaller than 1 and we can conclude that a good H exists. Therefore, by applying Theorem 3.1 we have:

Theorem 3.2 *For any k and n such that $k < n^{1/4}$ and any graph G with maximum degree d, there is a k-fault tolerant graph T with maximum degree which is $O(kd)$.*

Remark: In the construction above we assumed that k divides n and that $\frac{n}{k} + 1$ is a prime. However, the theorem itself can be proved without these assumptions by taking totally random functions h_i and by allowing at most one non-full row in our grid.

3.2 Explicit Constructions

Let $P \geq k$ be a prime and suppose r is an integer and $P^r = (n+k)/k$. We label the nodes W (and U) with the vectors of length r over $GF[P]$ such that all the nodes on the same row receive the same label. Thus, a node $w \in W$ is $(\vec{x}, y)$ where $\vec{x} = (x_1, x_2, \ldots x_r)$ and $x_j \in GF[P]\}$ and $0 \leq y \leq k-1$.

The collection H is defined by ℓ vectors $\vec{a}_1, \vec{a}_2, \ldots \vec{a}_r$ of length r over $GF[P]$. For each $1 \leq i \leq \ell$ the function $h_i \in H$ is defined as $h_i((\vec{x}, y)) = (\vec{x} + y\vec{a_i}, y)$ where the addition is vector addition over $GF[P]$ and the multiplication is that of a vector by a scalar, and where y is treated as an element in $GF[P]$.

For such a collection to be bad means that there is a set $S = \{(\vec{x}, y_1), \vec{(x)}_2, y_2), \ldots (\vec{x}_k, y_k)\}$ such that no member of H is good for S. Therefore, for every element $h_i \in H$ there is a pair $w_{i_1} = (\vec{x}_{i_1}, y_{i_1})$ and $w_{i_2} = (\vec{x}_{i_2}, y_{i_2})$ such that $h_i(w_{i_1}) = h_i(w_{i_2})$, and

$$\vec{x}_{i_1} + y_{i_1}\vec{a}_i = \vec{x}_{i_2} + y_{i_2}\vec{a}_i.$$

Pick *one* such *colliding pair* for each $h_i \in H$ and consider the following graph: there are k nodes, each corresponding to an element of S. Connect two nodes if they were picked as a colliding pair for some h_i. Thus the number of edges is ℓ. The following fact is well-known (see [4] p. 126):

Proposition 3.1 *A graph on k nodes with ℓ edges where $\ell \geq 2k$ has a cycle of length at most*

$$g = 2\left\lceil \frac{\log k}{log\lfloor \ell/k \rfloor} \right\rceil + 1.$$

Find the cycle of length at most g whose existence the proposition ensures, and let its nodes (after relabeling) be $w_1, w_2, \ldots w_g$. Therefore we have that

$$\begin{aligned}
\vec{x}_1 + y_1\vec{a}_1 &= \vec{x}_2 + y_2\vec{a}_1 \\
\vec{x}_2 + y_2\vec{a}_2 &= \vec{x}_3 + y_3\vec{a}_2 \\
&\ldots \\
\vec{x}_g + y_g\vec{a}_g &= \vec{x}_1 + y_1\vec{a}_g
\end{aligned}$$

by summing up we get that $\sum_{i=1}^{g}(y_{i+1} - y_i)\vec{a}_i = 0$. Since, for example, $y_1 \neq y_2$, we get that there must exist g vectors in $\vec{a}_1, \vec{a}_2, \ldots \vec{a}_\ell$ that are *linearly dependent* over $GF[P]$.

Our strategy in choosing $\vec{a}_1, \vec{a}_2, \ldots \vec{a}_\ell$ is, therefore, to make them such that no g of them are linearly dependent. This is exactly the condition required from the columns of a parity check matrix of a linear error correcting code over GF[P]. See Macwilliams and Sloane [18] for information on error correcting codes. Take, for instance, a Reed Solomon code of length ℓ, minimum distance $g+1$ and dimension $\ell - r$. Such a code exists if $g+1 \leq r$ and $\ell \leq P-1$. The columns of the parity check matrix are indexed by the non-zero elements of $GF[P]$. For $i \in GF[P]$ we have $\vec{a}_i = (1, i, i^2, \ldots, i^{r-1})$ where all the computation is over $GF[P]$.

Suppose, for example, that $k \leq n^{1/20}$, $\ell = k\lceil 2^{\frac{2\log^2 k}{log n}} \rceil$, P is a prime between ℓ and 2ℓ and $r = \frac{\log((n+k)/k)}{\log P}$ is an integer. By Proposition 3.1 we can take g to be $2\lceil \frac{\log k}{log \ell/k} \rceil + 1$ and we can verify that $g+1 \leq r$, which implies that the appropriate code exists. Although this example required that $(n+k)/k$ be a power of P, the approach can be extended to remove this requirement, yielding:

Theorem 3.3 *For any n and k where $k \leq n^{1/20}$ and for any graph G with maximum degree d, there is an* explicit *construction for a k-fault tolerant graph T with maximum degree which is $O(\lceil 2^{\frac{2\log^2 k}{\log n}} \rceil kd)$. Note that when n is $\Omega(k^{\log k})$ this bound is $O(kd)$.*

4 Expander Based Constructions

In this section we replace the grid that appears in the constructions in Section 3 by an appropriate expander. This enables us to extend the $O(kd)$ upper bound for a wider range of the parameters n, k and d. The crucial property of an expander which is used here is the fact that some of its expansion properties remain even after deleting many of its edges. This is related to the main idea in [1], but here we need several additional ideas. Let us call an undirected graph R an (n, δ, λ)-graph if it is a δ-regular graph on n nodes, and the absolute value of each of the eigenvalues of its adjacency matrix, besides the largest, is at most λ. It is known ([12]) that random δ-regular graphs on n-nodes are almost surely (n, δ, λ)-graphs for $\lambda = O(\sqrt{\delta})$ and that for every $\delta = p+1$ where p is a prime congruent to 1 modulo 4 there is an explicit infinite family of $(n, \delta, 2\sqrt{\delta - 1})$-graphs ([17], [19]). For more details on such graphs see [2].

Let $R = (V(R), E(R))$ denote, throughout this section, an (n, δ, λ)-graph. For a node v of R let $N(v)$ denote the set of all its neighbors. We say that two sets of nodes of equal cardinality S and S' in R can be *linked* if there are $|S|$ node disjoint paths in R connecting the nodes of S with these of S'. The main combinatorial property we need is the following result, which seems to be of independent interest. Proofs of lemmas and theorems in this section appear in the appendix.

Theorem 4.1 *Let R be as above and suppose*

$$8\frac{\lambda}{\delta} + 8(\frac{\lambda}{\delta})^2 < 1. \tag{3}$$

Then, every two sets of nodes S and S' in R, where $|S| = |S'|$ and where for every node v of R $|N(v) \cap S| \leq \delta/2$ and $|N(v) \cap S'| \leq \delta/2$, can be linked.

To prove this theorem we first need the following simple lemma, whose proof can be found in

[2], page 122.

Lemma 4.1 *If $R = (V(R), E(R))$ is as above and $X \subset V(R)$ is a set of xn nodes then*

$$\sum_{v \in V(R)} (|N(v) \cap X| - x\delta)^2 \leq \lambda^2 x(1-x)n.$$

This lemma implies the following two assertions.

Lemma 4.2 *For any two sets B and B' in R each of cardinality at least $\frac{\lambda}{\delta}n$ there is at least one edge between B and B'.*

Lemma 4.3 *Let B and X be two sets of nodes of R and suppose that each node of B has at least $\delta/2$ neighbors in X. If $|X| = xn$ and $|B| = bn$ then*

$$x \geq \frac{b\delta^2}{4(b\delta^2 + \lambda^2)}. \tag{4}$$

We can now combine Theorem 4.1 and Theorem 3.1 to prove the following.

Theorem 4.2 *For any n and k and any graph G with n nodes and maximum degree d there exists a k-fault tolerant version T of G with maximum degree*

$$O\left(\frac{\log n}{\log(n/k)}kd\right).$$

Note that for $k < n^{1-\beta}$ for any fixed $\beta > 0$ this bound is $O(kd)$.

5 Lower bounds

For a graph $G = (V, E)$ let $g(G, k)$ denote the minimum possible value of the maximum degree of a k-fault tolerant version of G. As shown in the previous section, if d is the maximum degree of G and k is not too large as a function of n (e.g., $k \leq n^{1/4}$) then $g(G, k) \leq O(dk)$. In this section we prove that for a wide range of values of k, d and n there exists a graph G on n nodes with maximum degree d so that $g(G, k) \geq \Omega(d\sqrt{k})$. It would be interesting to close the gap between these two estimates. The precise statement of our lower bound theorem is the following. Proofs of lemmas and theorems in this section appear in the appendix.

Theorem 5.1 *There is an absolute positive constants c_1 such that if d, k and n satisfy*

$$n > d \geq \frac{(\log n)^2}{\log\log n},$$

$$k \leq c_1 \frac{d^{1/2}}{(\log d)^{1/2}}$$

and

$$k < \frac{n}{2d},$$

then there exists a graph G on n nodes with maximum degree at most d so that $g(G,k) \geq \frac{1}{16}d\sqrt{k}$.

The proof employs a probabilistic construction of a suitable random graph G. Very roughly, it is based on the fact that any k-fault tolerant version T of G must contain many copies of G as subgraphs, and for an appropriate random G these copies cannot share too many common edges, due to the asymmetry of random graphs. This forces the maximum degree of any such T to be large. In what follows we omit all floor and ceiling signs whenever these are not crucial, to simplify the notation. In addition, we make no attempt to optimize the constants in the various estimates. For a graph G and for a subset W of its set of nodes let $G[W]$ denote the induced subgraph of G on the set of nodes W.

Lemma 5.1 *Let m and q be integers and let $1 > \epsilon > 1$ be a real satisfying*

$$2mqe^{-0.01\epsilon m} + 2q\binom{m}{2}e^{-0.01\epsilon^2 m} + q^2 2^{2m}(2m)!e^{-0.001\epsilon^3 m^2} < 1. \tag{5}$$

Then there exists a graph $G = (V,E)$ on $n = 2mq$ nodes with the following four properties.

(i) G is the node disjoint union of q bipartite graphs $G_1, \ldots, G_q$, where the two node classes of G_i are A_i and B_i and $|A_i| = |B_i| = m$ for $1 \leq i \leq q$.

(ii) The degree of every node of G is at least $\frac{1}{2}\epsilon m$ and at most $2\epsilon m$.

(iii) No two nodes of G have more than $2\epsilon^2 m$ common neighbors.

(iv) For any $1 \leq i \leq j \leq q$ and for any two disjoint subsets $B \subset B_i$ and $B' \subset B_j$ satisfying $|B| = |B'| = \epsilon m/4$, the number of edges in any graph which is isomorphic to a subgraph of $G[B \cup A_i]$ and to a subgraph of $G[B' \cup A_j]$ does not exceed $\epsilon^3 m^2/2$.

Lemma 5.2 *Suppose m, q and ϵ satisfy (5), define $n = 2mq$ and let k, $n \geq k \geq 1$ be an integer, $\epsilon k < 1/2$. Define $d = 2\epsilon m$ and let $G = (V,E)$ be a graph satisfying the conditions (i)-(iv) in Lemma 5.1. Then $g(G,k) \geq \frac{1}{16}d\sqrt{k}$.*

Remark A close inspection of the proof of Lemma 5.2 shows that the lower bound holds even if we allow the fault tolerant graph T to have $2n \geq n + k$ nodes. Thus, if the parameters n, d and k satisfy the assumptions in Theorem 5.1 then there is a graph G with n nodes and maximum degree d such that the following holds. If T is a graph with at most $2n$ nodes and any subgraph of T obtained by deleting k nodes of T contains a copy of G, then the maximum degree of a node of T is at least $\Omega(d\sqrt{k})$.

A Appendix: Proofs

Proof of Lemma 4.2: Suppose this is false and let B and B' be a counter-example. Put $|B| = bn$ and $|B'| = b'n$. By Lemma 4.1

$$\sum_{v \in B'} (|N(v) \cap B| - b\delta)^2 \leq \lambda^2 b(1-b)n.$$

However, by assumption there are no edges between B and B' and hence the left hand side of the last inequality is precisely $b'n(b\delta)^2$. It follows that

$$bb'\delta^2 \leq \lambda^2(1-b) < \lambda^2,$$

which is impossible, since both b and b' are at least $\frac{\lambda}{\delta}$. □

Proof of Lemma 4.3: By Lemma 4.1

$$\sum_{v \in B} (|N(v) \cap X| - x\delta)^2 \leq \lambda^2 x(1-x)n.$$

If $x \geq 1/2$ then inequality (4) follows trivially. Otherwise, the last inequality implies

$$bn(\frac{1}{2} - x)^2\delta^2 \leq \lambda^2 x(1-x)n,$$

i.e.,

$$\frac{b\delta^2}{4} \leq x(1-x)(\lambda^2 + b\delta^2) \leq x(\lambda^2 + b\delta^2).$$

This implies (4). □

Proof of Theorem 4.1 Suppose the theorem is false and let S and S' be a counter-example. By Menger's Theorem this implies that there exists a set C of cardinality $|C| < |S|$ $(= |S'|)$ such that in $R - C$ there is no path between a node in S and a node in S'. Let B be the set of all nodes in the connected components of $R - C$ that contain a node of S and let B' be the set of all nodes in the connected components of $R - C$ that contain a node of S'. Without loss of generality we may assume that $|B| \leq |B'|$. Put $|B| = bn$. Since obviously there is no edge between B and B', Lemma 4.2 implies that $b \leq \lambda/\delta$. Define $X = B \cup (C \setminus S)$, $|X| = xn$. We claim that $x < 2b$. Indeed, if this is false then $|X| \geq 2|B|$ and hence $|C \setminus S| \geq |B|$. Since $S \subset B \cup C$ this implies that

$$|C| = |C \cap S| + |C \setminus S| \geq |C \cap S| + |B| \geq |S|,$$

contradicting the fact that $|C| < |S|$. Hence $x < 2b$ as claimed.

Observe now that all the neighbors of every node v of B lie in $B \cup C \subset X \cup S$. Moreover, since by assumption v has at most $\delta/2$ neighbors in S it follows that it has at least $\delta/2$ neighbors in X. Therefore, by Lemma 4.3 and since $x < 2b$

$$2b > x \geq \frac{b\delta^2}{4(b\delta^2 + \lambda^2)}.$$

But since $b < \lambda/\delta$ this implies that

$$8(\lambda\delta + \lambda^2) \geq 8(b\delta^2 + \lambda^2) > \delta^2,$$

contradicting the assumption (3). This completes the proof. $\Box$

Proof Outline of Theorem 4.2: If n is large and k exceeds, say, $n/10$, the claimed bound is more than $n+k$ and we can simply take T to be complete. We can thus assume that $k \leq n/10$. Let R be an $(n+k, \delta, \lambda)$-graph, where $\delta = c \log n / \log(n/k)$, c is an absolute constant and inequality (3) holds. (Such graphs obviously exist, as mentioned in the beginning of the section, and can be given even constructively for some infinitely many values of n). Let M be the directed graph obtained from R by replacing each edge of R by two antiparallel directed edges. Let S' be a fixed set of k nodes in R (or M), so that no node of R has more than $\delta/2$ neighbors in S'. (A random set will have this property almost surely). This will be the set of spare nodes of M. The set of mappings H required in Theorem 3.1 is now simply chosen as a set of ℓ random bijections. For every fixed set S of k nodes, Theorem 4.1 implies that if a mapping h maps this set to a set of nodes of R that does not contain more than $\delta/2$ neighbors of any node of R, then h is good for S in the sense discussed in the previous section. It is easy to see that the probability that a random h is not good for a given fixed set S is at most $O(n2^{\delta}(k/n)^{\delta/2}) < 1/n$, where the last inequality holds if the constant c is chosen appropriately. Thus, the probability that there exists some set S for which no $h \in H$ is good is at most

$$O(\binom{n+k}{k}(n(4k/n)^{\delta/2})^{\ell}),$$

which is less than 1 for $\ell = c'k \log(n/k)/\log n$ for an appropriate c'. Combining this with the general procedure summarized in Theorem 3.1 and Theorem 2.2 the desired result follows. $\Box$

Proof of Lemma 5.1: Let $A_1, \ldots, A_q, B_1, \ldots, B_q$ be $2q$ pairwise disjoint sets of nodes, each of cardinality m. The graph G on the set of nodes

$$V = A_1 \cup \ldots \cup A_q \cup B_1 \cup \ldots \cup B_q$$

is the random graph obtained by choosing, for each i, $1 \leq i \leq q$ and for each $a \in A_i$ and $b \in B_i$, randomly and independently, the pair ab to be an edge with probability ϵ. The obtained graph obviously satisfies (i). To complete the proof we show that with positive probability it satisfies (ii), (iii) and (iv) as well.

The conditions (ii) and (iii) are very simple. The degree of every fixed node of G is a binomial random variable with expectation ϵm and variance $\epsilon(1-\epsilon)m$. By the well known tail estimates for binomial distributions (see, e.g., [2]), the probability that the degree of a fixed node is not between $\frac{1}{2}\epsilon m$ and $2\epsilon m$ is at most $e^{-0.01\epsilon m}$ (where the constant 0.01 can be easily improved). As there are $2mq$ nodes the probability that (ii) fails is at most

$$2mqe^{-0.01\epsilon m}. \tag{6}$$

The condition (iii) is similar. Two nodes can have common neighbors only if they both lie in the same set A_i or in the same set B_i. There are precisely $2q\binom{m}{2}$ such pairs of nodes. For a fixed pair, the number of their common neighbors is a binomial random variable with expectation $\epsilon^2 m$ and variance $\epsilon^2(1-\epsilon^2)m$. Therefore, by the known estimates for binomial distributions we conclude that the probability that (iii) fails is at most

$$2q\binom{m}{2}e^{-0.01\epsilon^2 m}. \tag{7}$$

It remains to estimate the probability that condition (iv) does not hold. Fix i and j, $1 \leq i \leq j \leq q$ and fix two disjoint subsets $B \subset B_i$ and $B' \subset B_j$. (Since B_i and B_j are disjoint for $i < j$ the assumption that B and B' are disjoint is informative only for $i = j$.) If a copy of some graph appears in $G[B \cup A_i]$ and in $G[B' \cup A_j]$ then there is a bijection $f : B \cup A_i \mapsto B' \cup A_j$ which maps the copy in $G[B \cup A_i]$ to that in $G[B' \cup A_j]$. Fix a bijection f as above. The maximum number of edges of G in a graph that is mapped into a copy of itself by this mapping is precisely the number of edges uv in $G[B \cup A_i]$ for which $f(u)f(v)$ is an edge in $G[B' \cup A_j]$. For a fixed pair uv the probability that both uv and $f(u)f(v)$ are edges is ϵ^2, as these two pairs are distinct since $B \cap B' = \emptyset$. Moreover, the events corresponding to distinct pairs uv are mutually independent, as no edge of G is examined twice in considering all pairs

$$\{uv : u \in B, v \in A_i\} \cup \{f(u)f(v) : u \in B, v \in A_i\}.$$

It follows that the number of edges in a maximum common subgraph mapped to a copy of itself by f is a binomial random variable with expectation $\epsilon^2|B|m = \epsilon^3 m^2/4$ and variance $\epsilon^2(1-\epsilon^2)|B|m$. Therefore, the probability it exceeds $\epsilon^3 m^2/2$ is at most $e^{-0.001\epsilon^3 m^2}$. The number of choices of B and B' is less than $q^2 2^{2m}$, and the number of choices of the bijection f is less than $(2m)!$ Therefore, the probability that (iv) does not hold is at most

$$q^2 2^{2m}(2m)!e^{-0.001\epsilon^3 m^2}.$$

Combining the last inequality with inequalities (6) and (7) and with the assumption (5) we conclude that with positive probability all the conditions (i)-(iv) hold, completing the proof. □

Proof of Lemma 5.2: Let $T = (U, D)$ be a k-fault tolerant version of G. For every node $v \in V$ let $S(v)$ denote the set of all nodes u of T so that there is a copy of G in T in which u plays the role of v. Clearly $|S(v)| > k$ for all v (since $T - S(v)$ contains no copy of G). Define $A = A_1 \cup \ldots \cup A_q$, $B = B_1 \cup \ldots \cup B_q$ and consider the set of all ordered pairs $\{(v, u) : v \in A, u \in S(v)\}$. The cardinality of this set is clearly more than $|A|k$ and hence there is some $u \in U$ so that

$$|\{v \in A : u \in S(v)\}| \geq \frac{|A|k}{n+k} \geq \frac{k}{4}.$$

Put $l = k/4$ and let A' be a set of l distinct nodes in A so that $u \in S(a)$ for $a \in A'$. Since G satisfies conditions (ii) and (iii), each $a \in A'$ has at least

$$\frac{1}{2}\epsilon m - (l-1)2m\epsilon^2 > \frac{1}{4}\epsilon m$$

neighbors which are not adjacent to any other member of A'. Fix, for each $a \in A'$, a set $B(a) \subset B$ of $\epsilon m/4$ neighbors of a which are not adjacent to any other member of A'. The sets $\{B(a) : a \in A'\}$ are clearly pairwise disjoint. Let, also $A(a)$ denote the set A_i where i is the unique index for which $a \in A_i$.

Let $C \subset U$ be the set of all neighbors of u in T, and let $|C| = x$ be the degree of u in T. For each $a \in A'$, let $G(a)$ be a copy of G in T in which u plays the role of a. Also, let $G'(a)$ be the copy of $G[B(a) \cup A(a)]$ in $G(a)$. Observe that by condition (ii), the number of edges of each $G'(a)$ is at least $\frac{1}{2}\epsilon m|B(a)|$. The crucial point now is that condition (iv) implies that there are no two distinct members of the set $\{G'(a) : a \in A'\}$ that share more than $2\epsilon^2|B(a)|m$ common edges. Therefore, each $G'(a)$ has at least

$$\frac{1}{2}\epsilon m|B(a)| - (l-1)2\epsilon^2 m|B(a)| \geq \frac{\epsilon m}{4}|B(a)| = \frac{\epsilon^2 m^2}{16}$$

edges that do not belong to any other $G'(a')$. Hence, the total number of edges in the union of all the subgraphs $\{G'(a) : a \in A'\}$ of T is at least

$$l\frac{\epsilon^2 m^2}{16} = \frac{1}{64}\epsilon^2 m^2 k.$$

However, all these edges are incident with the set C of the x neighbors of u in T. It follows that the maximum degree of a node of T is at least the maximum between x and $\frac{\epsilon^2 m^2 k}{64x}$, which is at least the geometric mean of these two quantities. This shows that the maximum degree is at least

$$\frac{1}{8}\epsilon m\sqrt{k} \geq \frac{1}{16}d\sqrt{k},$$

where the last inequality follows from the fact that G satisfies condition (ii). Since T was an arbitrary k-fault tolerant version of G this shows that $g(G,k) \geq \frac{1}{8}d\sqrt{k}$ and completes the proof of the lemma. $\square$

Proof of Theorem 5.1 Let b be a large absolute constant. If $d \geq bn^{2/3}(\log n)^{1/3}$ define $m = n/2$, $\epsilon = d/2m$ and $q = 1$, and check that these satisfy the assumptions of Lemma 5.2 (for appropriate c_1 and b). Otherwise, let m be the largest integer so that $d \geq bm^{2/3}(\log m)^{1/3}$ and define $\epsilon = d/2m$ and $q = n/2m$. Again, one can check that these satisfy the assumptions of Lemma 5.2. We omit the detailed computation. $\square$.

References

[1] N. Alon and F. R. K. Chung, Explicit Construction of Linear Sized Tolerant Networks, Discrete Math. 72(1988), 15-19.

[2] N. Alon and J. Spencer, **The Probabilistic Method**, John Wiley and Sons Inc., New York, 1992.

[3] F. Berman, M.E. Bock, E. Dittert, M.J. O'Donnell and D. Plank, Collections of Functions for Perfect Hashing, SIAM Journal on Computing, Vol 15, No. 2, May 1986, pp. 604 – 618.

[4] B. Bollobas, **Extremal Graph Theory**, Academic Press, London, 1978.

[5] J. Bruck, R. Cypher and C.T. Ho, Fault-tolerant parallel architectures with minimal numbers of spares, IBM Research Report, RJ 8029, March 1991.

[6] J. Bruck, R. Cypher and C.T. Ho, *Fault-Tolerant Mesh and Hypercube architectures*, IBM Research Report, RJ 8566, January 1992.

[7] S. Dutt and J. P. Hayes, *On Designing and Reconfiguring k-Fault-Tolerant Tree Architectures*, IEEE Trans. on Computers, vol. C-39, no. 4, pp. 490–503, April 1990.

[8] S. Dutt and J. P. Hayes, *Designing Fault-Tolerant Systems Using Automorphisms*, Journal of Parallel and Distributed Computing, vol. 12, pp. 249–268, 1991.

[9] S. Dutt and J. P. Hayes, *Some Practical Issues in the Design of Fault-Tolerant Multiprocessors*, Proceedings of the 21st International Symposium on Fault-Tolerant Computing, pp. 292–299, June 1991.

[10] M.L. Fredman and J. Komlós, On the Size of Separating Systems and Families of Perfect Hash Functions, SIAM Journal of Algebraic and Discrete Methods, Vol 5, No. 1, March 1984, pp. 61–68.

[11] M.L. Fredman, J. Komlós and E. Szemerédi, Storing a Sparse Table with $O(1)$ Worst Case Access Time, Journal of the Association for Computing Machinery, Vol 31, No. 3, July 1984, pp. 538–544.

[12] J. Friedman, J. Kahn and E. Szemerédi, On the second eigenvalue in random regular graphs, Proc. 21st ACM STOC, ACM Press (1989), 587–598.

[13] J. Hastad, F. T. Leighton and M. Newman, *Fast Computations using Faulty Hypercubes*, Proceedings of 21st Annual ACM Symposium on Theory of Computing, pp. 251–284, 1989.

[14] J. P. Hayes, *A Graph Model for Fault-Tolerant Computing Systems*, IEEE Trans. on Computers, vol. C-25, no. 9, pp. 875–884, September 1976.

[15] C. Kaklamanis, A. R. Karlin, F. T. Leighton, V. Milenkovic, P. Raghavan, S. Rao, C. Thomborson and A. Tsantilas, *Asymptotically Tight Bounds for Computing with Faulty Arrays of Processors*, Proceedings of 31st Annual IEEE Symposium on Foundations of Computer Science, pp. 285–296, October 1990.

[16] S. Y. Kung, S. N. Jean and C. W. Chang, *Fault-Tolerant Array Processors Using Single-Track Switches*, IEEE Trans. on Computers, vol. C-38, no. 4, pp. 501–514, April 1989.

[17] A. Lubotzky, R. Phillips and P. Sarnak, Explicit expanders and the Ramanujan conjectures, Proc. 18^{th} ACM STOC (1986), 240-246. See also: A. Lubotzky, R. Phillips and P. Sarnak, Ramanujan graphs, Combinatorica 8 (1988), 261-277.

[18] F. J. Macwilliams and N. J. A. Sloane, **The Theory of Error-Correcting Codes**, North Holland, Amsterdam, 1977.

[19] G. A. Margulis, Explicit group-theoretical constructions of combinatorial schemes and their application to the design of expanders and superconcentrators, Problemy Peredachi Informatsii 24(1988), 51-60 (in Russian). English translation in Problems of Information Transmission 24(1988), 39-46.

[20] K. Mehlhorn, **Data Structures and Algorithms 1: Sorting and Searching**, Springer-Verlag, Berlin Heidelberg, 1984.

[21] M. Paoli, W. W. Wong and C. K. Wong, *Minimum k-Hamiltonian Graphs, II,* J. of Graph Theory, Vol. 10, pp. 79–95, 1986.

[22] A. L. Rosenberg, *The Diogenes Approach to Testable Fault-Tolerant VLSI Processor Arrays*, IEEE Trans. on Computers, Vol. C-32, no. 10, pp. 902–910, October 1983.

[23] A. L. Rosenberg, *On Designing Fault-Tolerant Arrays of Processors*, Technical Report CS-1982-14, Dept. of Computer Science, Duke University, 1982.

[24] J. Spencer, **Ten Lectures on the Probabilistic Method**, SIAM (Philadelphia), 1987.

[25] W. W. Wong and C. K. Wong, *Minimum k-Hamiltonian Graphs*, J. of Graph Theory, Vol. 8, pp. 155-165, 1984.