

On the Fault Tolerance of Fat-Trees^{*}

SOTIRIS NIKOLETSEAS¹ GRAMMATI PANTZIOU¹ PANAGIOTIS PSYCHARIS²
PAUL SPIRAKIS¹

¹ Computer Technology Institute, P.O. Box 1122, 26110 Patras, Greece
Phone: (+30)61-273496, Fax: (+30)61-222086
email: {nikole,pantziou,spirakis}@cti.gr
² Intrasoftware, Greece

Abstract. We examine the reliability properties of ideal fat-trees, a general model used to capture both distance and bandwidth constraints of various classes of fat-tree networks. We allow the edges and the vertices of the network to fail independently with probability f , and show that: (1) Any fat-tree G can always be partitioned into an upper (G_H) and a lower (G_L) part. After the faults, the remaining part of G_L guarantees that a *linear fraction* of the leaves of the fat-tree *still connect to the upper part*, with high probability. (2) G_H is robust, in the sense that, after the faults, at least half of the edge-disjoint paths between any set of “leaves” of G_H are *preserved* with probability tending to 1, even in the case of failure probabilities as high as $f < 0.25$. The robust properties of G_H hold for the case that fat-nodes do not have internal edges and also for the case that fat-nodes are random regular graphs. (3) For the special case of a pruned butterfly, there is a *critical* probability p_c for the existence of a linear sized component surviving the failures and including a large fraction of terminal nodes. We show that $p_c \geq 0.42$.

1 Introduction

Fat-trees play an important role in modern multiprocessor architectures because of the efficient (with respect to VLSI design and routing) universality properties they have ([17], [3], [11]) and, also, because they provide a natural basis for the development of efficient routing algorithms ([10], [18], [6]) and for the implementation of interconnection networks in parallel supercomputers ([16]).

Aiming at a general study of various existing fat-tree networks, we use the ideal fat-tree model introduced by Bilardi et al ([7]), which captures both distance and bandwidth constraints and provides a framework for designing algorithms portable among different fat-tree networks. An abstract definition of the ideal fat-tree network is the following: The N terminal nodes (processors) are placed at the leaves of a complete binary tree whose internal nodes are switching modules building a routing network among the processors. The nodes of the tree

^{*} This research was partially supported by the EU Long Term Research Projects GEP-PCOM (contract No. 9072) and ALCOM-IT (contract No. 20244).

are connected via channels of appropriate capacity, which is equal to the number of edges joining each node (except from the root) to its parent. The communication over the network is subject to certain distance and capacity constraints: 1) A message which is sent from a source s it arrives at its destination t after $2h$ steps, where $h \in \{1, \dots, \log N\}$ is the height of the lowest common ancestor of s and t , and 2) the number of messages traversing any channel should not exceed the channel's capacity. This model is ideal in the sense that the routing time is the minimum satisfying distance and capacity constraints, while most existing fat-tree networks would incur a polylogarithmic (with respect to the ideal fat-tree) slowdown in routing time. Also, a general capacity function $w(n)$ usually obeys the following natural conditions: (i) $w(2n) \geq w(n)$, that is, the outgoing bandwidth of a subtree does not decrease with its size, and (ii) $w(2n) \leq 2w(n)$, that is, the bandwidth toward the parent does not exceed the total bandwidth toward the children. Here $n \in \{1, \dots, N/2\}$ is the number of leaves of the subtree rooted at a given node and $w(n)$ denotes the capacity of the fat-edge (or the number of edges) to the parent. Usually, $w(n) = \Theta(n^{1/2})$ is assumed since it satisfies the above conditions, and also ensures *area universality* (see [5]) of the ideal fat-tree.

In this work, we analyze the tolerance of ideal fat-trees in the presence of faults, which may lead to unavailability of parts of fat-nodes or edges of the fat-tree (interconnection) network of the multiprocessor architecture. Let G be an undirected graph. A random graph G^* of type- G is obtained by selecting edges of G independently and with probability p . We can thus, represent a communication network in which edges fail independently and with probability $f = 1 - p$. In [20, 22] this model was defined and used in order to prove the preservation with high probability of multiconnectivity and expander properties in various graphs of type- G . Let G be the tree-structured graph representing the initial fat-tree before the faults. The corresponding random graph G^* of type- G represents the faulty fat-tree obtained by allowing edges to fail independently and with probability f . We show that various properties in type- G fat-tree networks are preserved for a wide range of failure probabilities, including f as high as $O(1)$. (Note that usually, the assumption $f = o(1)$ (e.g. $f < 1/n$) captures many realistic permanent failure patterns. A value of $f = \Theta(1)$ (independent of n) is considered to be a worst-case assumption.) In particular, we show that: (1) Any fat-tree, G , can always be partitioned into a lower part, G_L , and an upper part G_H . We show that G_H is robust in the sense that *it preserves most of its edge-disjoint paths* (between its "terminals") with high probability even when $f < 0.25$. (2) We also show that *a linear fraction of the leaves* (terminals) of G *stays connected* (through what remains out of G_L) to G_H , with high probability.

There is a substantial body of literature concerning the fault tolerance of interconnection networks which we will not review in detail here due to space limitations (see e.g. [24, 14, 15, 8, 23, 12]). In [24] a fault tolerant network which is area-universal over all network layouts with few bends is given, and impossibility results with respect to the area-universality of fault tolerant networks under arbitrary patterns of faults, are presented. In [14, 15, 8, 23] the fault tolerance

properties of butterflies, multibutterflies, and randomly-wired splitter networks are studied. Note that although, the above important results can be (in some cases) probably extendable to other networks including the fat-trees, our results concerning the robustness properties of ideal fat-trees are of interest since they can be applied to any tree-structured interconnection network including a variety of fat-tree networks.

In this work we also consider the pruned butterfly, which is an interesting special case of a fat-tree with fat-nodes having no internal edges. For this special case we study (Section 5) the critical probability p_c for the existence of a linear sized component surviving the failures and including a large fraction of terminal nodes, and we show that $p_c \geq 0.42$. The existence of a linear-sized component in the faulted version of a network is an important measure of its robustness. If such a component does not exist, then the computation in the faulted version of the network is bound to be more than a constant factor slower than the computation in the fault-free version. A number of results for special networks have been presented in the literature. We refer here to the works of Karlin et al ([12]) for the butterfly network, Kesten ([13]) for the d by d two-dimensional mesh, Erdős and Renyi ([9]) for the complete graph, and Ajtai et al ([1]) for the hypercube of dimension d . Our result for the case of the pruned butterfly is a judicious modification of the approach in [12, 1].

2 Faulty fat-trees remain “fat-enough” even for constant failure probabilities: The case where fat-nodes have no internal edges

Let G be the undirected graph representing the tree-structured fat-tree network and T_G be the corresponding tree taken by viewing each fat-node as a single (super)vertex and by collapsing the parallel edges into one, with the appropriate capacity. Bilardi and Bay ([5]) have suggested the following definition of “ γ -fat” tree-networks, which is based on the relation between the number of edge disjoint paths connecting two subsets A, B of terminal nodes and the maximum flow $F(T_G, A, B)$ that can be pushed from A to B over the network. The “ γ -channel-sufficient” definition of Bilardi and Bay ([5]) is as follows:

Definition [Bilardi and Bay, [5]] *A tree-structured graph G is “ γ -channel-sufficient” if for every disjoint sets of leaves A, B there exist at least $\gamma F(T_G, A, B)$ edge disjoint paths from A to B in G , where T_G is the tree corresponding to graph G and $\gamma \in (0, 1)$ is a constant.* \square

We consider here the fat-tree partitioned into two parts, a higher (closer to the root) G_H and a lower G_L part. We first show that G_H is robust, even for $f = O(1)$. For the G_L part we just show (in section 4) that at least a constant fraction of the terminal nodes remains connected to (what remains of) G_H despite the faults.

2.1 Preservation of minimum cuts in upper (G_H) fat-trees with edge failures

Let H be such that all channels in the upper H levels of the graph G have capacity at least $\alpha\sqrt{n}\sqrt{\log n}$, where $\alpha \geq \sqrt{2}$ a constant. Denote by G_H this upper part of the graph G . Remark that a cut in G_H with respect to any two subsets A, B of “terminal” nodes of G_H is a set of edges intersecting any path between nodes of A and B . Note that the minimum cut between any subset of the “terminal” nodes of G_H (i.e. the fat-nodes of G_H to which G_L connects) is at least $\alpha\sqrt{n}\sqrt{\log n}$ before the faults. Now let $C_{A,B}(G_H)$ be the size of minimum (in cardinality) cut in G_H , for A, B fixed, and $C_{A,B}(G_H^*)$ be the corresponding minimum cut value in the graph of type- G_H .

Theorem 1. $\forall A, B: E(C_{A,B}(G_H^*)) = (1 - f) C_{A,B}(G_H)$

Proof (sketch): Consider a random indicator variable X_e for every edge e in a minimum $A - B$ cut of G_H , taken value 1 (0) when e remains in G_H^* (or not, respectively). The theorem follows easily by linearity of expectation. \square

Now sort arbitrarily the edges of G_H and let I_j be a random indicator variable showing whether edge e_j remains in G_H^* or not, that is

$$I_j = \begin{cases} 1 & \text{if } e_j \in G_H^* \\ 0 & \text{otherwise} \end{cases}$$

Consider the sequence X_0, X_1, \dots, X_t of random variables such that

$$\begin{aligned} X_k &= E(C_{A,B}(G_H^*) | I_1, \dots, I_k) \\ X_0 &= E(C_{A,B}(G_H^*)) \\ X_t &= \text{the value of minimum } A - B \text{ cut in } G_H^* \end{aligned}$$

Lemma 2. X_0, X_1, \dots, X_t is a (Doob) martingale (in fact, an edge exposure martingale).

Proof: Just define the filter F_k, F_{k-1}, \dots, F_0 where F_k is the σ -field generated by the events corresponding to I_1, \dots, I_k . Clearly, $E(X_{k+1} | F_k) = X_k$, and the martingale criterion is met. \square

Lemma 3. $|X_k - X_{k-1}| \leq 1$

Proof: Remark that taking into account the indicator variable I_k , corresponding to whether $e_j \in G_H^*$ is holding or not, can only decrease the expectation of the minimum $A - B$ cut of the “so-far exposed” graph by at most 1 (this may happen when $e_j \notin G_H^*$, otherwise $X_k = X_{k-1}$). \square

Thus, the Lipschitz condition holds, so we can employ the powerful method of bounded differences and get, by Azuma’s inequality (see for example [19]), that, $\forall \lambda > 0$

$$\Pr\{|X_t - X_0| \geq \lambda\sqrt{t}\} \leq 2e^{-\lambda^2/2} \quad (1)$$

We are now ready to prove the basic theorem of this section of the paper:

Theorem 4. For any $\alpha \geq \sqrt{2}$,

$$\forall A, B : \Pr\{|C_{A,B}(G_H^*) - E(C_{A,B}(G_H^*))| > \alpha\sqrt{n}\sqrt{\log n}\} \leq \frac{2}{n}$$

Proof: By the definition of X_0, X_1, \dots, X_t we get that $\Pr\{|C_{A,B}(G_H^*) - E(C_{A,B}(G_H^*))| > \alpha\sqrt{n}\sqrt{\log n}\} = \Pr\{|X_t - X_0| > \lambda\sqrt{n}\}$, where now $\lambda = \alpha\sqrt{\log n}$, $\alpha \geq \sqrt{2}$. Then by equation (1) we get

$$\Pr\{|X_t - X_0| > \lambda\sqrt{n}\} \leq 2e^{-\frac{\alpha^2 \log n}{2}} = 2n^{-\frac{\alpha^2}{2}} \leq \frac{2}{n}$$

□

Thus, the minimum $A-B$ cut of G_H^* is around $(1 \pm \beta)E(\min_{A,B}\text{-cut of } G_H^*)$, i.e., (by theorem 1) around $(1 \pm \beta)(1 - f)(\min_{A,B}\text{-cut of } G_H)$ with probability $\geq 1 - \frac{2}{n}$ for some $\beta \in (0, 1)$. By choosing a β close to zero, we get

Corollary 5. For any A, B the maximum flow of G_H between A, B is only fractionally reduced (by $(1 - f)$) almost certainly.

Note that this Corollary gives us information about the flow that can be pushed in the whole graph only if the remaining part of G_L can push such flow into G_H . Note also that in [4], the authors have shown a related result, namely, that the value of every cut in a *compressed* version G^* of any particular graph G is very close to the corresponding cut in G . They have shown this to hold also for $s-t$ cuts. However, the way the graph is compressed (in [4]) does not capture our fault model because their model includes each edge e in G^* with probability p_e (the compression probability) but also gives it a weight of $1/p_e$ if it is included. Thus, one cannot just apply the results in [4].

2.2 Preservation of edge disjoint paths under constant vertex and edge failures

Recall that in this first part of the paper, we consider the (worst, as far as connectivity issues are concerned) case where the fat-nodes of the tree-structured graph G have no internal edges at all. Now let H' be such that each fat-node of the part $G_{H'}$ consisting of the upper H' levels of graph G , has at least $k \log n$ (internal) vertices ($k \geq 2$ a constant). The next theorem shows that, for f constant, a constant fraction of the number of vertices in each fat-node survive the vertex faults with high probability.

Definition 6. For a particular fat-node V in $G_{H'}$ let E_V be the event “at least $(1 - \beta)(1 - f)|V|$ vertices of V remain in $G_{H'}^*$ ”, where $\beta \in (0, 1)$ a constant. Also, let $E = \bigcap_V E_V$, for all $V \in G_{H'}$.

Theorem 7. $\Pr\{E\} \geq 1 - n^{-(k'-1)}$, where $k' \geq 2$ a constant.

Proof: See full paper [21]. \square

The above theorem implies the survival of a constant fraction of vertices after the vertex faults within each fat-node. We will now prove that the edges connecting the fat-nodes tolerate the edge faults well-enough to guarantee the preservation of at least half of the edge disjoint paths when $f < 0.25$, almost certainly. We show, in particular, that at least half of the edge-disjoint paths in $G_{H'}$ (connecting any sets of “terminal” vertices of $G_{H'}$) are preserved in $G_{H'}^*$, with high probability. Note that this implies that for any two sets of terminal nodes of the fat tree that remain connected to $G_{H'}^*$ after the faults, the edge-disjoint paths that join them are fractionally preserved, at least as far as the $G_{H'}^*$ part of the graph is concerned. We assume here that at least a constant fraction of terminal nodes connect to the upper part G_H .

Theorem 8. *At least half of the edge disjoint paths in the graph $G_{H'}$ connecting any sets A, B of terminal nodes remain in $G_{H'}^*$, with probability at least $1 - \Theta\left(\frac{1}{\log n}\right)$, for all failure probabilities $f < 0.25$.*

Proof: We will show this holding (on the average) in the following section. From Theorem 7, at least $(1 - \beta)(1 - f)|V|$ vertices within each fat-node survive the vertex faults. Remark that when $(1 - \beta)(1 - f) > 3/4$, then for any two children V_1, V_2 of the same parent V in $G_{H'}^*$, more than half of the nodes of V get edges from both V_1, V_2 . Now, for a fat-node V of $G_{H'}$ at level $k + 1$, let $p_\epsilon(k + 1)$ be the (bad-event) probability that in $G_{H'}^*$, the remaining edge disjoint paths connecting any sets A, B of terminal nodes that arrive at the node V are less than half of the corresponding paths in $G_{H'}$. Since, by our previous remark, more than half of the vertices of V get edges from both V_1, V_2 this bad event holds if the remaining edge disjoint paths arriving from either V_1 or V_2 are less than half of the ones before the failures, thus:

$$p_\epsilon(k + 1) = 1 - (1 - 0.5 p_\epsilon(k))^2 = p_\epsilon(k) - \frac{p_\epsilon^2(k)}{4} \quad (2)$$

By choosing $q(k)$ such that $p_\epsilon(k) = \frac{1}{q(k)+1}$ we get from equation 2 that

$$\begin{aligned} \frac{1}{q(k+1)+1} &= \frac{1}{q(k)+1} - \frac{1}{4(q(k)+1)^2} = \frac{4q(k)+3}{4(q(k)+1)^2} \Leftrightarrow \\ q(k+1) &\sim q(k) + 2 + \frac{1}{q(k)} \end{aligned}$$

By induction, it follows that $k < q(k) < k + H_{k-1} + 3$, where H_k is the k -th harmonic number, implying that $q(k) = \Theta(\log n) \Leftrightarrow p_\epsilon(k) = \Theta\left(\frac{1}{\log n}\right)$. To complete the proof, remark that we get the best possible f satisfying $(1 - \beta)(1 - f) > 3/4$ by setting $\beta \simeq 0$, i.e. $1 - f > 3/4 \Leftrightarrow f < 1/4 = 0.25$. \square

3 The case where the fat-nodes are random regular graphs

In this section, we study the reliability properties of fat-trees whose fat-nodes, instead of having no edges at all (as in the previous section), now have quite a lot of “built-in” connectivity. To be more specific, consider the case where the fat-nodes of the tree-structured graph G are random regular graphs of degree d (hence they are also expanders with high probability).

Let H' be such that all channels in the upper H' levels of the graph G have capacity at least $k \log n$, where $k > 0$ a constant. Denote by $G_{H'}$ this upper part of the graph (including the random regular fat-nodes) and by $T_{G_{H'}}$ the corresponding tree-structured graph (where we view each fat-node as a single supervertex). Let C be a capacity in $T_{G_{H'}}$ and C^* the corresponding capacity in $T_{G_{H'}}^*$ (the type- $T_{G_{H'}}$ random graph representing what remains from $T_{G_{H'}}$ after the edge faults). Note that the capacity of a channel connecting two fat-nodes is equal to the number of edges joining these nodes.

Definition 9. Let E_C be the event “ C^* is within $(1 \pm \beta)(1 - f)C$ ”, where $\beta \in (0, 1)$ is a constant and f the edge failure probability. \square

The following theorem shows that at least a constant fraction of the capacity of any channel of this part of the fat-tree (and thus the number of edges connecting the corresponding endvertices of the channel) is *preserved* with high probability.

Theorem 10. *There is a constant $k' \geq 2$ such that $\Pr\{\cup_{C \in T_{G_{H'}}} \overline{E_C}\} \leq n^{-(k'-1)}$*

Proof (sketch): As mentioned earlier, a channel capacity of size C in $G_{H'}$ implies the existence of C edges connecting the endpoints of the channel. But the edges in $G_{H'}$ fail independently and with probability f . Thus, $G_{H'}^*$ is the result of performing on each channel C Bernoulli trials with success probability $p = 1 - f$. By Chernoff bounds, we thus have:

$$\begin{aligned} \Pr\{E_C\} &= \Pr\{C^* \in (1 \pm \beta)(1 - f)C\} \geq 1 - e^{-\frac{\beta^2}{2}(1-f)C} \\ &\geq 1 - e^{-\frac{\beta^2}{2}(1-f)k \log n} = 1 - n^{-k'} \end{aligned}$$

where $k' = \frac{\beta^2}{2}(1-f)k$ can be made at least 2 by choosing appropriate values for β, k . Thus, $\Pr\{\exists C \text{ in } T_{G_{H'}} : \overline{E_C} \text{ in } T_{G_{H'}}^*\} \leq n n^{-k'} = n^{-(k'-1)}$ \square Note that the above theorem implies the preservation of a constant fraction of the edges joining any vertex pair, by choosing a small (but even constant) failure probability f and β close to 0. In order to investigate what happens in the interior of the fat-nodes we remark that each fat-node is a member of $G_{n,p}^d$, the class of all random regular graphs of degree d , whose edges fail independently and with probability f . In [20], [22] the following facts have been shown:

Fact 1 $G_{n,p}^d$ is highly disconnected when $f = 1 - p$ is constant and $d < \frac{1}{2}\sqrt{\log n}$, almost certainly.

Fact 2 When $G_{n,p}^d$ is disconnected, it still has a giant connected component for any $f < 1 - \frac{32}{d}$ with high probability, for any $d \geq 64$.

Fact 3 The giant connected component of a random member of $G_{n,p}^d$ remains, with high probability, a certifiable efficient expander, despite the edge faults, provided that $f < 1 - \frac{256}{d}$.

From Fact 3, we conclude that a constant degree $d > d_0$, $d_0 = \frac{256}{1-f}$ suffices to guarantee that each fat-node remains, with high probability, an efficient expander despite the constant edge failure probability f . By theorem 10 and fact 3 we get:

Corollary 11. Let E be the number of edge disjoint paths in fat-trees whose fat-nodes are random regular graphs of degree d . At least $(1-f)E$ edge disjoint paths survive the constant failure probability edge faults almost certainly, provided that $d > \frac{256}{1-f}$. \square

The following special case is indicative of the power of Corollary 11:

Corollary 12. Let E be the number of edge disjoint paths in fat-trees whose fat-nodes are random regular graphs of degree d . At least $(3/4)E$ edge disjoint paths survive the failures almost certainly, provided that $f < 0.25$ and $d > d_0$, where $d_0 > (4/3)256$ i.e. $d_0 > 340$. \square

4 The connectivity of terminal nodes to the upper part

Let T be the set of terminal nodes of a fat tree R . Let G_L, G_H be the lower (respectively, higher) part of the “tree” graph as previously. Let $F(G_L, T, G_H)$ be the maximum flow from T to the nodes of the G_H through G_L . Assuming that R is γ -channel-sufficient we know that there exist at least $\gamma F(G_L, T, G_H)$ edge disjoint paths from T to G_H in R ([5]). If $|T| = N$, then the minimum cut between T and the nodes of G_H is clearly at least θN ($0 < \theta < 1$ a constant) due to the channel capacities. Thus, the nodes in T are connected to the upper part via $N' = \gamma\theta N$ disjoint paths. The survival probability of each such path is $q = (1-f)^{\Theta(\log \log N)}$. For f constant, q becomes at least $\frac{1}{\log^c n}$, where $c > 1$ a constant. From the Bernoulli of N' trials and success probability q we get (by Chernoff bounds) that:

Lemma 13. At least $(1 \pm \beta)qN$ terminal nodes are connected to G_H via edge disjoint paths with probability at least $1 - e^{-\frac{\beta^2}{2}qN}$, for any $\beta \in (0, 1)$.

For example, when f is constant then the expected number of terminal nodes that connect to G_H via edge disjoint paths is $\Theta\left(\frac{n}{\log^c N}\right)$. Let X_0, X_1, \dots be the sequence where X_j is the expected number of terminal nodes in T that connect (possibly through non-edge-disjoint paths) to level j of the lower part (level 0 is T). Assume that the fat-tree is 1-channel sufficient (e.g., a pruned butterfly). Then, because $\gamma = 1$ and $f < 0.5$, it is $E(X_{j+1}) = X_j$ because each vertex sends at least two edges upwards, thus we get (for a proof see full paper) that:

Lemma 14. $\{X_j\}$ is a martingale, for any $f < 0.5$. □

By using Azuma's inequality and sums of Poisson trials we then get:

Theorem 15. For any $f < 0.5$ and $\gamma = 1$ the expected number of terminal nodes that connect to G_H (possibly via overlapping paths) is $\Theta(N)$, and the actual number is concentrated around $\Theta(N)$. □

Thus, we can suggest that robust fat-trees can be built in the following way: (i) The lower part should consist of fat-nodes guaranteeing $\gamma = 1$ (e.g., a pruned butterfly or cliques or tree of meshes). (ii) The upper part should be built via fat-nodes being concentrators (e.g., regular random graphs) because of our results in the previous sections.

5 The special case of the pruned butterfly

The pruned butterfly is a particular case of a fat-tree with $\gamma = 1$ (see e.g. [6] for a definition of the pruned butterfly). Given a pruned butterfly P_d , let P_d/p denote the random pruned subbutterfly of type- P_d obtained by considering each edge independently and including (excluding) it in the pruned subbutterfly with probability p (respectively, $f = 1 - p$). In this section we show that there is a critical probability p_c such that for $p > p_c$ the pruned butterfly P_d/p has a connected component of linear size which includes at least $h \cdot 2^d$ terminal nodes, where $0 < h < 1$ a constant, with high probability, while for $p < p_c$ such a component does not exist almost certainly.

We follow closely, but judiciously modify, the approach of Karlin, Nelson and Tamaki in [12] for the case of the butterfly network, and the approach of Ajtai, Komlos and Szemerédi in [1] for the hypercube network. We first choose $p_1 > p_c$ and $p_2 > 0$ such that $(1 - p_1)(1 - p_2) = 1 - p$ and therefore $P_d/p_1 \cup P_d/p_2 = P_d/p$. Then, we prove that a constant fraction of the nodes of P_d/p_1 are in connected components of an appropriate size, which are called *atoms*. Finally, we show that the additional edges in P_d/p_2 connect the atoms into a linear-sized component. Due to space limitations, details are given in the full paper ([21]).

The characterization of the critical probability given above provides an algorithm to compute it. In the full version of this paper we give the details of the computation of the critical probability. We also give a lower bound on the value of the critical probability p_c ($p_c \geq 0.42$).

Acknowledgement: We wish to thank G. Bilardi for his encouragement and fruitful discussions.

References

1. M. Ajtai, J. Komlos, and E. Szemerédi, "Largest random component of a k-cube", *Combinatorica*, 2(1):1-7, 1982.
2. N. Alon and J. Spencer, "The Probabilistic Method", John Wiley, 1992.

3. P. Bay and G. Bilardi, "An area-universal VLSI circuit", in Proc. of the Symp. on Integrated Systems, pp. 53-67, 1993.
4. A. Benczur, D. Karger, "Approximating $s - t$ Minimum Cuts in $\tilde{O}(n^2)$ time", in the Proc. of the Symp. on Theory of Computing (STOC'96), pp. 47-55, 1996.
5. G. Bilardi and P. Bay, "An area lower bound for a class of fat-trees", in Proceedings of the European Symposium on Algorithms, pp. 413-423, 1994.
6. P. Bay and G. Bilardi, "Deterministic on-line routing on area-universal networks", Journal of the ACM, 42(3): 614-640, 1995.
7. G. Bilardi, B. Codenotti, G. Del Corso, G. Pinotti and G. Resta, "Broadcast and Associative Operators on Fat-Trees", to appear in the Proc. of EUROPAR'97.
8. R. Cole, B. Maggs and R. Sitaraman, "Routing on butterfly networks with random faults", in the Proc. of FOCS'95, pp. 558-570, 1995.
9. P. Erdős and A. Renyi, "On the evolution of random graphs", Publ. Math. Inst. Hungar. Acad. Sci., 5:17-61, 1960.
10. R. Greenberg and C. Leiserson, "Randomized routing on fat-trees", in Randomness and Computation, pp. 345-374, JAI Press, 1989.
11. R. Greenberg, "The fat-pyramid and universal parallel computation independent of wire delay", IEEE Transactions on Computers, C-43 (12): 1358-1364, 1994.
12. A.R. Karlin, G. Nelson, and H. Tamaki, "On the Fault Tolerance of the Butterfly", in the Proc. of the Symp. on Theory of Computing (STOC'94), pp. 125-133, 1994.
13. H. Kesten, "The critical probability of bond percolation on the square lattice equals $1/2$ ", Communication in Mathematical Physics, 74:41-59, 1981.
14. , T. Leighton, B. Maggs and R. Sitaraman, "On the Fault Tolerance of Some Popular Bounded-Degree Networks", In the Proc. of FOCS'92, pp. 542-552.
15. F. T. Leighton and B. M. Maggs, "Fast Algorithms for Routing Around Faults in Multibutterflies and Randomly-Wired Splitter Networks", In IEEE Transactions on Computers, Vol. 41:5, pp. 578-587, 1992.
16. C. Leiserson , Z. Abuhamedh, D. Douglas, C. Feynman, M. Ganmukhi, J. Hill, W. Hillis, B. Kuszmaul, M. Pierre, D. Wells, M. Wong, S. Yang and R. Zak, "The network architecture of the Connection Machine CM-5", in Proc. of the 4th ACM Symp. on Parallel Algorithms and Architectures, pp. 272-285, 1992.
17. C. Leiserson, "Fat-trees: Universal networks for hardware-efficient supercomputing", IEEE Transactions on Computers, C-34 (10):892-900, 1985.
18. T. Leighton, B. Maggs and S. Rao, "Universal packet routing algorithms", in Proceedings of the 29th Symp. on the Foundations of Computer Science, 1988.
19. R. Motwani, P. Raghavan, "Randomized Algorithms", Cambridge Univ. Press, 1995.
20. S. Nikolettseas, K. Palem, P. Spirakis and M. Yung, "Short Vertex Disjoint Paths and Multiconnectivity in Random Graphs: Reliable Network Computing" , 21st Col. on Automata, Languages and Programming (ICALP), pp. 508 - 515, 1994.
21. S. Nikolettseas, G. Pantziou, P. Psycharis, P. Spirakis, "On the Fault Tolerance of Fat Trees", Computer Technology Institute Tech. Report, TR.97.05.24, 1996.
22. S. Nikolettseas and P. Spirakis, "Expander Properties in Random Regular Graphs with Edge Faults", in the Proc. of STACS'95, pp. 421-432, 1995.
23. H. Tamaki, "Efficient self-embedding of butterfly networks with random faults", in the Proc. of IEEE FOCS, 1992, pp. 533-541.
24. S. Toledo, "Competitive Fault-Tolerance in Area-Universal Networks", n the Proc. of SPAA'92, 1992, pp. 236-246