

Measuring the Vulnerability of Interconnection Networks in Embedded Systems

V. Lakamraju, Z. Koren, I. Koren, and C. M. Krishna

Department of Electrical and Computer Engineering
University of Massachusetts, Amherst, MA 01003

Abstract. Studies of the fault-tolerance of graphs have tended to largely concentrate on classical graph connectivity. This measure is very basic, and conveys very little information for designers to use in selecting a suitable topology for the interconnection network in embedded systems. In this paper, we study the vulnerability of interconnection networks to the failure of individual links, using a set of four measures which, taken together, provide a much fuller characterization of the network. Moreover, while traditional studies typically limit themselves to uncorrelated link failures, our model deals with both uncorrelated and correlated failure modes. This is of practical significance, since quite often, failures in networks are correlated due to physical considerations.

1 Introduction

The interconnection network is an integral part of most embedded systems. It has often as considerable an impact on the system's performance as the nodes themselves. The choice of an appropriate interconnection network is therefore key to determining the performance of the embedded system. Performance measures for interconnection networks are essential to guide the designer in choosing an appropriate topology. In large systems – especially those which must operate for long durations without any possibility of repair – the probability is significant that one or more nodes and/or links are down at any time and this can affect the performance of the system considerably.

Studies of the fault-tolerance of networks have tended to largely concentrate on measures such as classical node (link) connectivity. They measure the extent to which the network can withstand the failure of individual links and nodes while still remaining functional. Such measures are very basic and limited in what they can express of reliability (see [4] for a survey of measures of network vulnerability). They are worst-case measures and convey very little information for designers to use in selecting a suitable topology for the interconnection network in embedded systems.

In this paper we study the vulnerability of an interconnection network to the failure of individual links, using a set of four measures which, taken together, provide a much fuller characterization of the network. Moreover, while traditional studies typically limit themselves to independent link failures, our studies deal with correlated failure modes, as well.

We start in Section 2 by defining four measures of network vulnerability. We follow this in Section 3 with some numerical results. A brief discussion in Section 4 concludes the paper.

2 The Performance Measures

The four performance metrics used to assess network vulnerability can be grouped into two pairs. The first pair assesses the tendency of the topology under study to become disconnected due to link failures. The two measures under this category are:

- 1. The probability that the network becomes disconnected, π_d .
- 2. The size of the biggest connected component, χ_{max} .

The probability that the network becomes disconnected gives us guidance as to the chance that all the processors remain usable (assuming the processors themselves do not fail) by being reachable from every other processor. If the network does get disconnected, we are interested in what happens to the splinters that are left. In particular, we are concerned with whether the graph breaks up into a large number of small components, or whether there is one large component which contains most of the nodes. The latter is obviously preferable. All other things being equal, therefore, we would prefer a network which would disconnect in such a way that the biggest component left after disconnection contains a large fraction of the nodes.

The second pair of measures focuses on node-pair distances. They are:

- 3. The diameter of the network, Δ .
- 4. The average distance between node pairs, \overline{D} .

Node pair distances play a role in determining the time it takes for messages to be sent from one node to another. A graph whose diameter is relatively stable is obviously superior to another whose diameter exhibits rapid variations upon link failure.

The notion of diameter stability is not new: the previously-defined measure of edge *persistence* [3] is the minimum number of edges that must be removed to increase the graph diameter. Persistence, however, being a worst-case measure, conveys much less information about graph vulnerability than does the diameter, as a function of the component failure probability.

Inter-node distances play a large role in determining the communication delays between nodes. Algorithms that assign tasks to nodes (processors) have to account for inter-node communication delays when dealing with tasks which communicate with one another. The smaller the delays between the nodes, the greater are the options available to the task assignment algorithm. This is especially true when the original task assignment (on a computer without any failures) is sought to be done in such a way that any task reassignment required upon failure is reduced. For hard real-time systems, it becomes important that the system state on the failed node be transferred to another node with very little delay. This parameter gives a good estimate of the amount of delay that would be involved in the movement of data that would be required to re-establish the state. A close estimation of such delays can help in the efficient calculation of fault-recovery policies[2]. It also gives an indication of how closely the nodes are connected to each other and this can help in the scheduling of tasks.

3 Simulation Models and Results

We consider two link failure models: uniform and clustered. In the uniform model, link failures follow an IID (independent and identically distributed) stochastic process. Each link fails with probability p_f , and link failures are independent of one another. In the clustered model, a probability of either $p - \delta$ or $p + \delta$ (for some given p , δ) is randomly selected for each *node*. Each link incident on a node fails with the failure probability drawn for that node. This failure mechanism results in adjacent links being correlated

with regard to faults, and consequently, in bigger clusters of faulty links and of fault-free links than those generated by the IID link failures. δ is the *clustering parameter*. The greater the value of δ , the more clustered the failing links will be. Note that since the failure probability is applied twice to the same link, the actual probability of a random link failure in the correlated model is $p_f = 1 - (1 - p)^2$.

Three different classes of topologies have been used for the simulation runs, namely, the mesh, the hypercube and the generalization of a chordal ring proposed by Arden and Lee[1]. This is a chordal ring in which extra links are added (apart from the 2 links connected to each of its neighbors) among the nodes in some regular fashion. The exact placement of these extra links has an impact on both the traditional measures as well as the ones proposed here.

All the simulation runs were on networks of 64 nodes. Some of the networks had degree 4 and the rest degree 6. A simple rectangular mesh as well as its counterpart, the mesh torus (a mesh with an end-around connection) and both 2-D and 3-D meshes were tested. Simulation runs were performed to measure the effect of the link failure probability, p_f , as well as the effect of the clustering parameter, δ , on the different performance measures, for the mentioned graph families. A number of interesting results can be concluded from the plots.

Figures 1, 2, 3 and 4 depict the dependence on the link failure probability p_f (in the IID link failure model) of the probability of network disconnection π_d , the maximum component size χ_{max} , the diameter Δ , and the average node-pair distance \overline{D} , respectively, for the different topologies.

The conclusions we can derive from these figures are as follows. Though the rectangular mesh is the topology of choice when scalability is concerned, it is certainly not the best topology when considering resistance to link failures. The probability that the network becomes disconnected increases rapidly as the probability of link failure increases. The other topologies in its class do better in all the other parameters as well. Similarly, among the degree-6 networks, the 3-D mesh performs very badly compared to the other topologies in its class.

The chordal ring of degree 4 has better diameter stability compared to the mesh torus. One word of caution though: The diameter of the chordal ring depends on the placement of the extra links (i.e. not those connected to immediate neighbors). For the simulations, an extensive search was performed to find a placement of links which would result in the minimum diameter.

The chordal ring of degree 6 performs only marginally better than the hypercube and the 3-D mesh torus in the diameter and average distance measures.

Figures 5 and 6 show the dependence of the probability of network disconnection, π_d , and that of the maximum component size, χ_{max} , respectively, on the fault clustering parameter, δ , for several graph topologies. The incidence of disconnected graphs increases with the failure clustering (even though the link failure probability remains the same). Again, the meshes without the end-around connection perform badly compared to the other networks. Each family of graphs has a distinctive sensitivity to the level of failure clustering.

The size of the largest connected component decreases as the degree of clustering increases. Also, the maximum component size is dependent on the clustering of links in the topology. This is illustrated in Figure 6 with the two types of the chordal ring. The *good placement* refers to an optimal placement of the links whereas *bad placement* refers to a sub-optimal placement. The dependence of the extra link placement on the component size becomes negligible as the degree of the network increases.

Our experiments also showed that the two other measures, namely, the diameter and the

average node-pair distance (in graphs that remain connected) are not very sensitive to the failure clustering, with the diameter being slightly more sensitive than the average distance. This result holds across graph families.

4 Discussion

In this paper, we have studied the vulnerability of various topologies to link failure. These results – and others like them – can be used by designers in choosing the appropriate topology. We have confined ourselves to a set of symmetric networks: we plan to extend our studies to irregular topologies.

There is also room for modeling correlated failures in other ways. One of them would be to use a “wave-propagation” model, in which the effect of the correlated failure at a node ripples through the network so that all the links which are at the same distance from the failed node has the same probability of failure and this probability decreases as the distance increases. It would also be interesting to look at the combined effect of both node and link failures.

Acknowledgment

This effort was supported in part by the Defense Advanced Research Projects Agency and the Air Force Research Laboratory, Air Force Materiel Command, USAF, under agreement number F30602-96-1-0341, order E349. The government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon.

The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, or the Defense Advanced Projects Agency, Air Force Research Laboratory, or the U. S. Government.

References

1. B.W.Arden and H.Lee. “Analysis of Chordal Ring Networks”, *IEEE Transactions on Computers* C-30, 1981.
2. M.Berg and I. Koren, “On Switching Policies for Modular Fault-Tolerant Computing Systems”, *IEEE Transactions on Computers* Vol. C-36, 1987.
3. F. T. Boesch, F. Harary, J. A. Kabell, “Graphs as Models of Communication Network Vulnerability: Connectivity and Persistence,” *Networks*, Vol. 11, 1981.
4. M.Choi and C.M.Krishna, “On Measures of Vulnerability of Interconnection Networks”, *Microelectronics and Reliability* Vol 29, No. 6, 1989.
5. T.Cormen, C. Leiserson and R.Rivest, *Introduction to Algorithms*, Cambridge: MIT Press, 1990.
6. C. M. Krishna and K. G. Shin, *Real-Time Systems*, New York: McGraw-Hill, 1997.

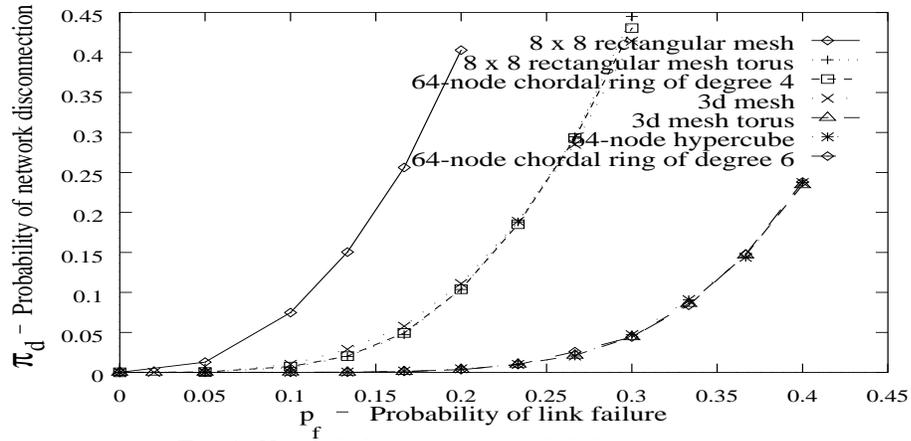


Fig. 1. Network disconnection probability versus p_f

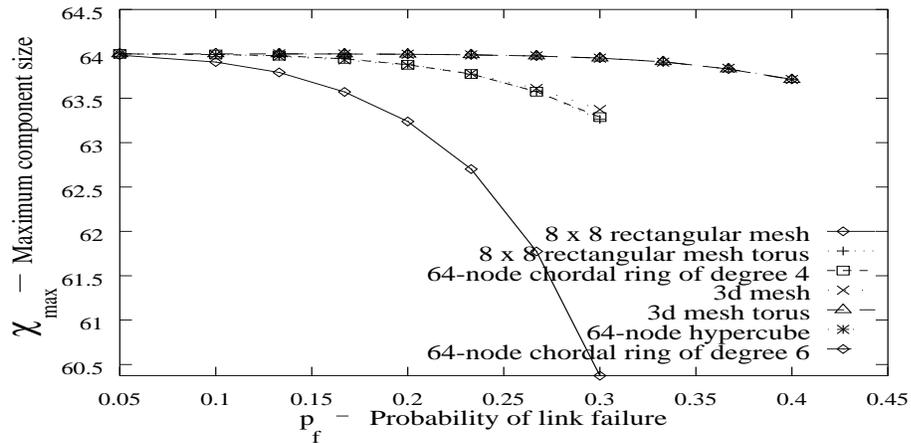


Fig. 2. Maximum component size versus p_f

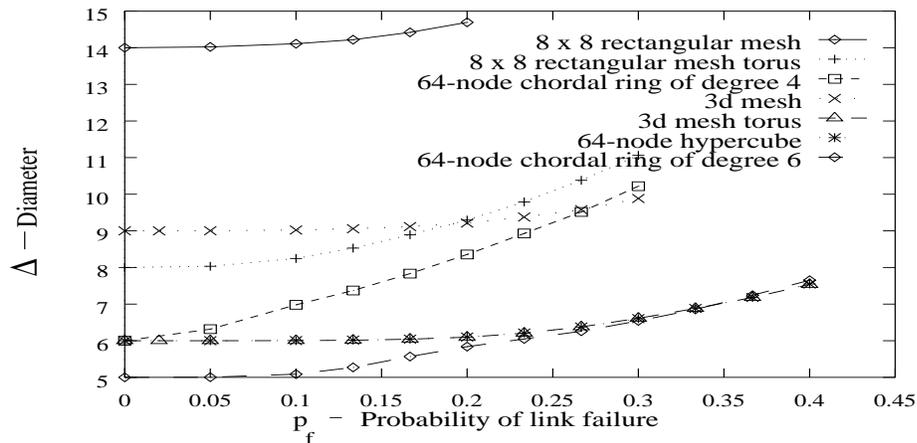


Fig. 3. Diameter versus p_f

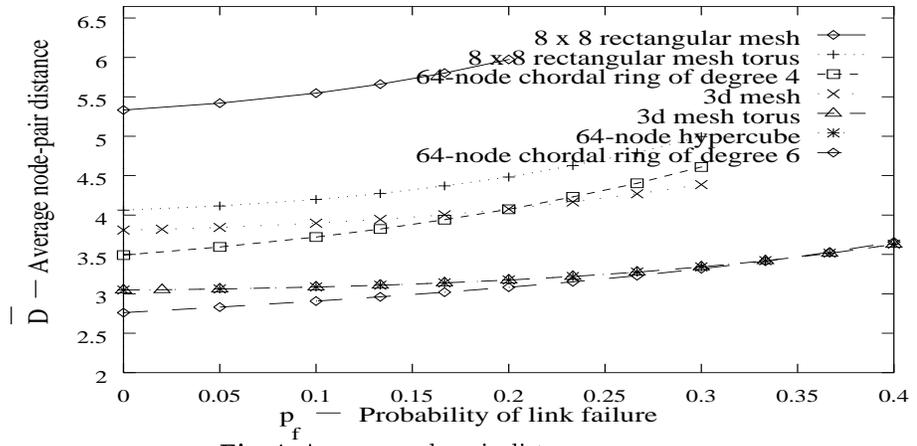


Fig. 4. Average node-pair distance versus p_f

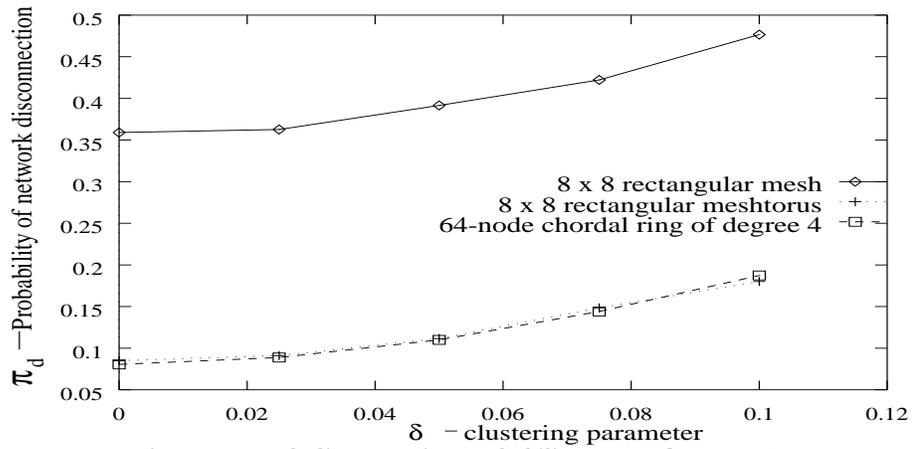


Fig. 5. Network disconnection probability versus δ ; $p_f = 0.1$

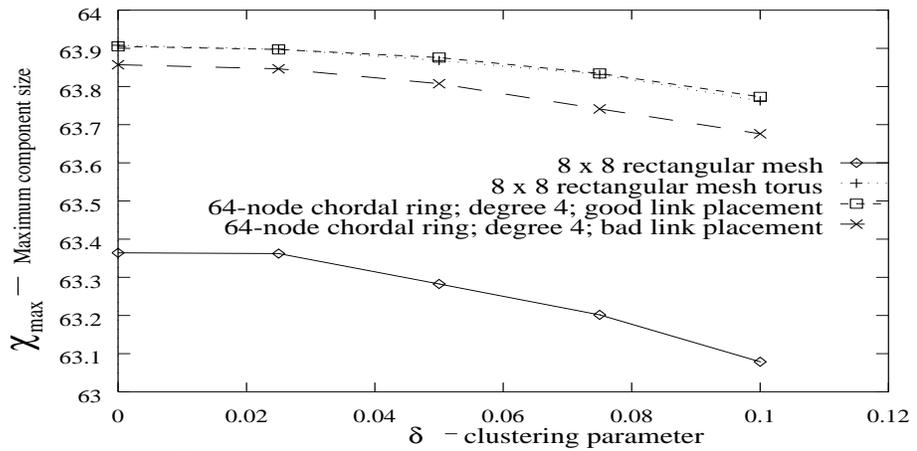


Fig. 6. Maximum component size versus δ ; $p_f = 0.1$