Information Sciences 180 (2010) 664–678

Contents lists available at ScienceDirect

Information Sciences

journal homepage: www.elsevier.com/locate/ins

A new performance measure for characterizing fault rings in interconnection networks

F. Safaei [a,*], A. Khonsari [b,c], M.M. Gilak [a]

[a] Department of Electrical and Computer Engineering, Shahid Beheshti University, Tehran, Iran
[b] Department of Electrical and Computer Engineering, University of Tehran, Tehran, Iran
[c] School of Computer Science, Institute for Research in Fundamental Sciences (IPM), Tehran, Iran

ARTICLE INFO

Article history:
Received 9 July 2007
Received in revised form 24 October 2009
Accepted 26 October 2009

Keywords:
Fault-tolerance
Fault rings
Torus
Interconnection networks
Adaptive routing
Performance analysis

ABSTRACT

One of the fundamental issues in parallel computers is how to efficiently perform routing in a faulty network where each component fails with some probability. Adaptive fault-tolerant routing algorithms in such systems have been frequently suggested as a means of providing continuous operations in the presence of one or more failures by allowing the graceful system degradation. Many algorithms involve adding buffer space and complex control logic to the routing nodes. However, the addition of extra logic circuits and buffer space makes nodes more liable to failure and less reliable. Further, if the shape of fault pattern is confined, then many non-faulty nodes will be sacrificed and hence their resources are wasted. This is clearly an undesirable solution and motivates solutions that provoke efficient use of non-faulty nodes. One such approach to reducing the number of functional nodes that must be marked as faulty is based on the concept of fault rings to support more flexible routing around rectangular fault regions. Before such schemes can be successfully incorporated in networks, it is necessary to have a clear understanding of the factors that affect their performance potential. In this paper, we propose the first general solution for computing the probability of message facing the fault rings with and without overlapping in the well-known torus networks. We also conduct extensive simulation experiments using various fault patterns, the results of which are used to confirm the good accuracy of the proposed analytical models.

1. Introduction

Fault-tolerant algorithms have recently attained renewed interest due to the widespread deployment of parallel computers [6], cluster-based systems [2], mobile systems [1], sensor networks [17], and networks on chip (NoC) [19]. An extensive amount of research has been carried out on fault-tolerant routing in interconnection networks [3–5,7–11,16,18,20,23–25,27]. Typically, additional routing restrictions and/or network resources are required to ensure deadlock freedom of such algorithms in the vicinity of faults [4,9,7,11,16,17,25]. Additionally, faults may be grouped into some *fault sets* that occupy rectangular (block) regions such that the boundary of the rectangle has only fault-free nodes and links and the interior of the rectangle contains all the faulty links and nodes which correspond to that particular fault set. This fault model is superior to convex fault regions proposed by Chien and Kim [9] because it deals more effectively with faults along the network boundary. For each fault region in a network with faults, it is feasible to connect the fault-free components around the region to form a ring known as the *fault ring* (f-ring for short) for that fault.

* Corresponding author. Tel.: +98 21 22287013; fax: +98 21 22828687.
E-mail addresses: f_safaei@sbu.ac.ir (F. Safaei), ak@ipm.ir (A. Khonsari), mo.gilak@mail.sbu.ac.ir (M.M. Gilak).

0020-0255/$ - see front matter © 2009 Elsevier Inc. All rights reserved.
doi:10.1016/j.ins.2009.10.017

Many researchers [4,5,7,8,16,20,23,27] have utilized the concept of *f-rings* in order to make the existing routing algorithms tolerate multiple fault regions without disabling any healthy nodes. Messages are routed under a regular adaptive routing until they face an f-ring. Then, depending on the relative position of the destination, messages are routed clockwise or counterclockwise around the f-ring. The f-rings provide alternate paths to messages blocked by faults which are used to route messages around rectangular fault regions. An f-ring can be constructed for each rectangular fault region by using the methods as described in [4,5,7,8,16,20,23,27]. Thus, in order to assess the trade-offs and benefits of such routing schemes accurately, we must be able to investigate the characteristics of the fault rings.

Several recent studies [1–3,6,10,16–19,24] have addressed fault-tolerance in a diverse range of systems and applications. Most of the performance evaluation studies on the functionality of these systems, however, have been solely based on simulations. The limitations of simulation-based studies are that they are highly time-consuming and expensive [22]. An alternative form of simulation approach is the analytical modeling. The significant advantage of analytical models over simulation is that they can be used to obtain performance evaluation for large systems which may not be feasible to study using simulation due to the excessive computational demands. The study presented in this paper is a performance evaluation, both analytic and experimental, of an adaptive routing system in torus networks. The goal is to calculate the probability of message facing the fault rings when an adaptive routing scheme is used. To the best of our knowledge, this research is the first attempt to determine the probability of message facing the fault rings as a means of examining the relative performance merits of adaptive fault-tolerant routing algorithms.

In this paper, we investigate the rectangular fault patterns which are suitable for modeling faults at the board level in networks, particularly in the torus topology. Our approach employs the theoretical results of matrices, algebra, and combinatorics to calculate the probability of occurrences of facing the f-rings. Analytical expressions are firstly derived for 2-D torus with non-overlapping f-rings and then the model is extended for overlapping f-rings. Our research effort is a step towards providing a versatile analytical tool for characterizing fault rings. Such tools could greatly help in the study of network performance in the presence of large number of faults.

The remainder of the paper is structured as follows. Section 2 provides relevant preliminary information and insight into fault-tolerance at the node as well as network level. In Section 3, we derive, in detail the probability of message facing the fault rings in the torus network. This section also describes an extension of the proposed analytical modeling approach in order to deal with the case of overlapping fault rings. Section 4 validates the analytical equations using extensive simulations. Finally, Section 5 summarizes the work reported in this paper and presents possible directions for future work.

2. Background

In this section, we define the terminology associated with the interconnection networks. For the sake of completeness, some of the definitions given below are reiterated from previous works [3,4,7,10–12,16,18,22,24–26].

2.1. The torus network

A topology can be represented as a graph $\mathbb{T} = (N, E)$ where N is the set of vertices and $E \subseteq N \times N$ is the set of edges of $\mathbb{T}$. We suppose that N is the set of routing nodes and E is the set of channels.

Definition 1. A link (or channel) is a point to point communication medium, connecting two nodes, with buffers at each end. A link is understood to operate in a full duplex mode.

To elaborate, we use the term link for a full duplex link, offering each communication device to both, send and receive messages; in this case the link has a transmitter and a receiver buffer at both ends. A channel e from node x to node y is represented as $\langle x, y\rangle$. If all links in the topology are bidirectional, then any $\langle x, y\rangle \in E$ implies that there is exactly as many as $\langle y, x\rangle \in E$.

Definition 2 *[4,12,24,26].* A topology $\mathbb{T}$ is disconnected by faulty nodes if there are non-faulty nodes u and v in $\mathbb{T}$ such that no fault-free path from u to v can be found in $\mathbb{T}$. Otherwise, it is connected.

Topologies can be regular or irregular [12,25]. Regular topologies are characterized by a predictable structure, and common examples are the (k, n)-*torus* (also known as *k-ary n-cubes*), meshes, and multistage networks such as Clos networks [12]. As an example of a predictable structure, (k, n)-*torus* have k_i nodes along each dimension i, where $k_i = k$ for all i. A node $\mathscr{A}$ is represented as a coordinate with n dimensions, $\left(\sigma^{\mathscr{A}}_{n-1}, \sigma^{\mathscr{A}}_{n-2}, \ldots, \sigma^{\mathscr{A}}_{2}, \sigma^{\mathscr{A}}_{1}\right)$ where $0 \leqslant \sigma^{A}_{i} \leqslant k_i - 1$ for $0 \leqslant i \leqslant n-1$. All nodes have the same number of neighbors, and two nodes $\mathscr{A}$ and $\mathscr{B}$ are neighbors if and only if $\sigma^{\mathscr{A}}_{i} = \sigma^{\mathscr{B}}_{i}$ for all $i,\ 0 \leqslant i \leqslant n-1$ except one, j, where $\sigma^{\mathscr{A}}_{i} = \left(\sigma^{\mathscr{B}}_{i} \pm 1\right) \bmod k$. The use of the modulo-operator represents wrap-around links [25]. A (k, n)-*torus* has k^n nodes. If $k = 2$, each node has n neighbors, one in each dimension. If $k > 2$, each node has $2n$ neighbors two in each dimension.

2.2. Routing algorithms

The topology describes the nodes that are connected to each other while the routing algorithm describes which nodes and links are visited by packets on their journey from source to destination. As before, definitions are used for emphasizing on new or unfamiliar terms. This section presents the terminology of routing algorithms.