

IEEE PROCEEDINGS

COMPUTERS AND DIGITAL TECHNIQUES

Volume 142, Number 6, November 1995

UNIVERSITY OF MICHIGAN

JAN 17 1996

ENGINEERING LIBRARY



**HONORARY EDITORS**Prof. E.L. Dagless (*University of Bristol*)Prof. C.R. Jesshope (*University of Surrey*)**PAPERS**

Papers should be addressed to the Managing Editor, *IEE Proceedings-Computers and Digital Techniques*, Publishing Department, Institution of Electrical Engineers, Michael Faraday House, Six Hills Way, Stevenage, Herts. SG1 2AY, United Kingdom; telephone 01438 313311; telex 825578 IEESTV G; facsimile 01438 742849. Notes for the guidance of authors may be found on the inside back cover.

Papers of outstanding merit are eligible for IEE premiums.

SUBSCRIPTIONS

There are eleven *IEE Proceedings* Titles. Six issues of each Title are published each year.

Annual subscription prices for 1995 (paper or microfiche, including delivery):

1 Title	£340.00
2 Titles	£520.00
3 Titles	£630.00
4 Titles	£756.00
5 Titles	£835.00
6 Titles	£900.00
7 Titles	£938.00
8 Titles	£968.00
9 Titles	£990.00
10 Titles	£1010.00
11 Titles	£1034.00

Each additional Title subscription in excess of eleven Titles: £94.00.

Single copy: £58.00. Airmail supplement: £20.00 per Title.

Orders should be addressed to IEE Publications Sales Department, Michael Faraday House, PO Box 96, Stevenage, Herts. SG1 2SD, United Kingdom.

Telephone: 01438 313311; telex 825578 IEESTVG; facsimile: 01438 742792.

ADVERTISING

Inquiries should be sent to the IEE Advertising Department at the Stevenage address.

COPYRIGHT AND COPYING

This publication is copyright under the Berne Convention and the Universal Copyright Convention. All rights reserved. Apart from any copying under the UK Copyright, Designs and Patents Act 1988, Part 1, Section 38, whereby a single copy of this article may be supplied, under certain conditions, for the purposes of research or private study, by a library of a class prescribed by The Copyright (Librarians and Archivists) (Copying of Copyright Material) Regulations 1989: SI 1989/1212, no part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means without the prior permission of the copyright owners. Permission is, however, not required to copy abstracts of individual contributions on condition that a full reference to the source is shown.

Single copies may be made for the purpose of research or private study.

Authorisation to photocopy items for internal or personal use, or for the internal or personal use of specific clients, is granted by the Institution of Electrical Engineers for libraries and other users registered with the Copyright Clearance Center Transactional Reporting Service, provided that the base fee of \$10.00 per copy is paid directly to the Copyright Clearance Centre Inc., 21 Congress Street, Salem, MA 01970, USA (1350-2387/95 \$10.00). This consent does not extend to other kinds of copying, such as copying for general distribution, for advertising or promotional purposes, for creating new collective works, or for resale.

Multiple copying of the contents of this publication without permission is always illegal.

The IEE is not as a body responsible for the opinions expressed by individual authors in *IEE Proceedings-Computers and Digital Techniques*.

The IEE is a member of the Association of Learned & Professional Society Publishers.

© 1995: The Institution of Electrical Engineers.

Secretary of the IEE: J.C. Williams, PhD, FEng, FIEE

Managing Editor: Gill Wheeler; Executive Editor: Julia Hubble

Typeset by: Pindar plc, York, United Kingdom

Printed by: Black Bear Press Limited, Cambridge, United Kingdom



IEE PROCEEDINGS

COMPUTERS AND DIGITAL TECHNIQUES

November 1995

Volume 142

Number 6

UK ISSN 1350-2387

ICDTEA 142 (6) 377-440

Contents

	page
Anatomy of a simulation backplane. M. Zwolinski, C. Garagate, Z. Mrcarica, T.J. Kazmierski and A.D. Brown	377
Adaptive wormhole routing in tori with faults. S. Chalasani and R.V. Boppana	386
Deadlock-free wormhole routing algorithms for star graph topology. C.P. Ravikumar and A.M. Goel	395
Technique to minimise area overhead for delay-driven clustering. C. Yeh and Y.-Y. Gu	401
Effects of technology mapping on fault-detection coverage in reprogrammable FPGAs. K. Kwiat, W. Debany and S. Hariri	407
Design and hardware architectures for dynamic Huffman coding. L.-Y. Liu, J.-F. Wang, R.-J. Wang and J.-Y. Lee	411
Novel design of arithmetic coding for data compression. J. Jiang	419
Concurrent error detection in array multipliers by BIDO. T.-H. Chen, Y.-P. Lee and L.-G. Chen	425
Spatial bounding of complex CSG objects. A. Sanna and P. Montusch	431

Adaptive wormhole routing in tori with faults

S.Chalasan
R.V.Boppana

Indexing terms: Adaptive routing, Deadlocks, Fault-tolerant routing, Multicomputer networks, Message routing, Performance evaluation, Torus networks, Wormhole routing

Abstract: The authors present a method to enhance wormhole routing algorithms for deadlock-free fault-tolerant routing in tori. They consider arbitrarily-located faulty blocks and assume only local knowledge of faults. Messages are routed via shortest paths when there are no faults, and this constraint is only slightly relaxed to facilitate routing in the presence of faults. The key concept used is that, for each fault region, a fault ring consisting of fault free nodes and physical channels can be formed around it. These fault rings can be used to route messages around fault regions. We prove that, at most, four additional virtual channels are sufficient to make any fully-adaptive algorithm tolerant to multiple faulty blocks in torus networks. Simulation results are presented for a fully-adaptive algorithm showing that good performance can be obtained with as many as 10% links faulty.

1 Introduction

Point-to-point torus and related networks are being used in many recent experimental and commercial multicomputers and multiprocessors [1–3]. A (k,n) -torus network has an n -dimensional grid structure with k nodes (processors) in each dimension such that every node is connected to two other nodes in each dimension by direct communication links. The *wormhole* (WH) switching technique by Dally and Seitz [8] has been used in many recent multicomputers [1–4]. In the WH technique, a packet is divided into a sequence of fixed-size units of data, called *flits*. If a communication channel transmits the first flit of a message, it must transmit all the remaining flits of the same message before transmitting flits of another message. To avoid deadlocks among messages, multiple virtual channels are simulated on each physical channel and a predefined order is enforced on the allocation of virtual channels to messages.

For fault-free networks, important issues in the design of a routing algorithm are high throughput, low-latency message delivery, avoidance of deadlocks,

livelocks and starvation and ability to work well under various traffic patterns. Given a network with faults, our approach is to use the existing network rather than recreate the original network using spare nodes and links. Therefore, for networks with faults, a routing algorithm should exhibit the following additional features: graceful degradation of performance and ability to handle faults with only a small increase in routing complexity and local knowledge of faults.

The well-known e -cube or dimension-order routing algorithm is an example of nonadaptive routing algorithms, since always a particular path is used in routing messages between a pair of nodes even when multiple shortest paths are available. With the e -cube, even a single fault disrupts communication between multiple pairs of nodes. With increase in adaptivity, a message is more likely to find a less congested path or fault-free path. Therefore, the issue of adaptivity, the extent of choice in selecting a path between a pair of nodes in routing a message, plays an important role in designing fault-tolerant routing algorithms.

1.1 Description of the problem and results

We present a technique to enhance minimal, fully-adaptive routing algorithms for fault-tolerant routing in tori. A minimal fully-adaptive algorithm routes messages along any of the shortest paths available. We consider routing methods that use only local knowledge of faults. We assume that faulty processors are confined to one or more rectangular blocks.

For each fault region, there exist one or more paths that pass through fault-free nodes and links and encircle the fault. For a fault in a 2D torus, there is an undirected ring of fault-free nodes and links; we refer to this ring as *fault-ring*. In this paper, we show that fault rings can be used to route messages around the fault regions using only local knowledge of faults and without introducing deadlocks and livelocks.

1.2 Related results

Adaptive, fault-tolerant routing algorithms for WH and virtual cut-through switching techniques has been the subject of extensive research in recent years [5–10]. Reddy and Freitas [11] use global knowledge of faults, spare nodes, and routing tables to investigate the performance limitations caused by faults. Gaughan and Yalamanchili [12] use a pipelined circuit-switching mechanism with backtracking for fault-tolerant routing. These two results are applicable to networks with arbitrarily-shaped faults. Our interest in this paper, is to design fault-tolerant wormhole routing algorithms that can be applied with local knowledge of faults. One important criterion is that the fault-free performance should not be sacrificed for fault-tolerant routing.

© IEE, 1995

IEE Proceedings online no. 19952079

Paper first received 5th January 1995 and in revised form 10th May 1995

S. Chalasan is with the Electrical & Computer Engineering Department, University of Wisconsin-Madison, Madison, WI 53706-1691, USA

R.V. Boppana is with the Computer Science Division, The University of Texas at San Antonio, San Antonio, TX 78249-0664, USA

There are no previous results specifically on fault-tolerant wormhole routing in tori. Often, the results developed for meshes [5,8,13] can be extended to tori with suitable modifications, since meshes and tori are closely related. The wraparound links in tori lead to extra deadlock possibilities, however. Therefore, if the results developed for meshes are applied with few changes, the number of virtual channels required to avoid deadlocks may be doubled [5]. Furthermore, meshes have edges, for example, the top row in a 2D mesh, and faults on edges are complicated to handle [13]. But this case never arises in tori, since they are node symmetric. Hence, extending efficient mesh routing techniques to tori in a straight forward manner may not necessarily yield efficient routing algorithms for the latter networks.

In terms of adaptivity and performance comparisons, the results by Dally and Aoki [8] are the most relevant to ours. With the dimension-reversal schemes of Dally and Aoki, a message may lose its adaptivity, if its number of dimension reversals equals the number of virtual channel classes. A message that has lost adaptivity is routed by the ϵ -cube algorithm and is not guaranteed to be delivered to its destination if there are faults in the network. Thus the number of virtual channels needed and the number of faults tolerated is highly dependent on the number of virtual channels and the location of faults. In contrast, our algorithms can tolerate any number and combination of rectangular faulty blocks with simple logic, and require only four virtual channels more than that required for the original adaptive algorithm. (Throughout this paper we indicate the number of virtual channels on per physical channel basis). This result compares well with our earlier result that four extra virtual channels are sufficient for routing in meshes with faults [13].

2 Preliminaries

A (k,n) -torus (also called k -ary n -cube) has n dimensions, numbered from 0 to $(n-1)$, and $N = k^n$ nodes. Each node is uniquely indexed by an n -tuple in radix k . Each node is connected via communication links to two other nodes in each dimension. The neighbours of the node $x = (x_{n-1}, \dots, x_0)$ in dimension i are $(x_{n-1}, \dots, x_{i+1}, x_i \pm 1, x_{i-1}, \dots, x_0)$, where addition and subtraction are modulo k . A link is said to be a wraparound link if it connects two neighbours whose addresses differ by $k-1$ in dimension i , $0 \leq i < n$. A (k, n) -mesh is a (k, n) -torus with the wraparound connections missing. We assume that each communication link represents two unidirectional physical communication channels. The link between nodes x and y is denoted by $\langle x, y \rangle$. To simplify presentation, we discuss the concept of fault-rings for two dimensional (2D) tori. We label the sides of a 2D torus as North, South, East and West.

2.1 Fault model

We consider both node and link faults. All the links incident on a faulty node are considered faulty. We assume that faults are permanent and nonmalicious faults. Therefore, messages are generated by and for nonfaulty processors only. We develop fault-tolerant algorithms, for which it is sufficient if each nonfaulty processor knows the status of the links incident on it.

A *fault set* is defined as the set F of faulty nodes and

faults and two link faults in the two-dimensional network shown in Fig. 1. We assume that faults in a 2D torus have *rectangular* shapes. A set F of faulty nodes and links in a 2D torus is said to have a rectangular shape if there is a rectangle in the torus such that: (a) there are no faulty components on the boundary of the rectangle, (b) the interior of the rectangle[†] includes all faulty components in F and (c) the interior of the rectangle contains no component that is not present in F .




Fig. 1 Three fault regions and their associated fault rings in a 6×6 torus

For example, the set $\{(3,3), (3,4), (4,3), (4,4)\}$ of faulty nodes shown in Fig. 1 is rectangular, since the interior of the rectangle – with corners $(2,2), (2,5), (5,2)$, and $(5,5)$ – includes all faulty components in F and no non-faulty component (recall that a processor fault implies that all links incident on it are faulty). However, the set of faulty links $\{\langle (1,1), (1,2) \rangle, \langle (1,2), (2,2) \rangle, \langle (2,2), (2,1) \rangle, \langle (2,1), (1,1) \rangle\}$ in a 6×6 torus is not rectangular, since any rectangle with nonfaulty elements on its boundary contains at least one element not in F . The faulty link $\langle (1,2), (2,2) \rangle$ is an example of a rectangular fault region, since the interior of the rectangle with corners $(1,1), (1,3), (2,1)$, and $(2,3)$ contains only the faulty link. The faulty link $\langle (0,0), (0,1) \rangle$ in Fig. 1 is considered rectangular; the rectangle that covers the faulty link has processors $(1,0), (1,1), (5,1)$ and $(5,0)$ as its corners.

An f -region is the fault region of the torus given by a block-fault. Under the *block-fault* model, the fault-set in a 2D torus can be written as a union of disjoint smaller fault sets, each of which denotes an f -region. For example, the fault set F in Fig. 1 is in fact the union of three disjoint f -regions $\{(3,3), (3,4), (4,3), (4,4)\}$, $\{\langle (0,0), (0,1) \rangle\}$ and $\{\langle (1,2), (2,2) \rangle\}$. We assume that faults do not disconnect the network.

There are many reasons to consider block faults. First, they model several common fault scenarios such as faults of isolated nodes and links and consecutive nodes in a row or column. Second, an arbitrarily-shaped fault can be modelled as a block-fault, albeit by labelling some nonfaulty processors and/or links as

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.