Resilient routing implementation in 2D mesh NoC

Rimpy Bishnoi [a,*], Vijay Laxmi [a], Manoj Singh Gaur [a], Mark Zwolinski [b]

[a] Department of Computer Science and Engineering, MNIT Jaipur, India
[b] Mark Zwolinski University of Southampton, Southampton, United Kingdom

A R T I C L E   I N F O

A B S T R A C T

With the rapid shrinking of technology and growing integration capacity, the probability of failures in Networks-on-Chip (NoCs) increases and thus, fault tolerance is essential. Moreover, the unpredictable locations of these failures may influence the regularity of the underlying topology, and a regular 2D mesh is likely to become irregular. Thus, for these failure-prone networks, a viable routing framework should comprise a topology-agnostic routing algorithm along with a cost-effective, scalable routing mechanism able to handle failures, irrespective of any particular failure patterns. Existing routing techniques designed to route irregular topologies efficiently lack flexibility (logic-based), scalability (table-based) or relaxed switch design (uLBDR-based). Designing an efficient routing implementation technique to address irregular topologies remains a pressing research problem. To address this, we present a fault resilient routing mechanism for irregular 2D meshes resulting from failures. To handle irregularities, it avoids using routing tables and employs a few fixed configuration bits per switch resulting in a scalable approach. Experiments demonstrate that the proposed approach is guaranteed to tolerate all locations of single and double-link failures and most multiple failures. Also, unlike uLBDR it is not restricted to any particular switching technique and does not replicate any extra messages. Along with fault tolerance, the proposed mechanism can achieve better network performance in fault-free cases. The proposed technique achieves graceful performance degradation during failure. Compared to uLBDR, our method has 14% less area requirements and 16% less overall power consumption.

1. Introduction

Network-on-Chip (NoC) is a scalable and modular communication paradigm proposed to overcome the limitations of traditional bus-based interconnects [1]. Regular network structures like 2D meshes are usually preferred by NoC designers owing to their simplicity and perfect physical layout on a 2D surface. As scaling of CMOS technology approaches nanometre scales, the reliability of on-chip interconnects becomes a design concern, as any failure may cause the entire system to fail in non-fault tolerant designs [2,3].

In traditional network literature, the most popular fault classes are transient faults and permanent faults [3]. Transient faults include temporary failures (bits-errors in a physical channel) that often occur at random and unpredictable times. Faults occurring due to neutrons and alpha particles are classed as transient faults. Other faults are permanent, including failures that are not temporary and that result in irreversible physical changes. These are caused by poor design, including incorrect specifications, manufacturing defects, component wear-outs, random device effects, broken wires, time-dependent dielectric breakdowns, electromagnetic interference (EMI), etc. [4].

Although, permanent faults are not as common and frequent as transient faults, our focus in this paper is restricted only to permanent failures. The reason is that transient faults can be detected and corrected locally by methods such as cyclic redundancy checking or forward error correction [4]. Moreover, if it occurs frequently, a transient fault can be modelled as a time-limited permanent fault and preventive measures taken for permanent faults are also applicable to this fault class. On the other hand, the effects of permanent faults are irreversible, and it is not always possible to repair or replace the components on a chip [3].

Furthermore, due to decrease of inter-wire spacing in modern on-chip interconnects, short switch-to-switch links in regular structures are more prone to noise sources such as crosstalk, EMI, radiation, process variations and so on. This fact makes links more susceptible to failures [4]. More specifically, our interest is restricted only to the permanent failures of underlying physical links and switches of a 2D mesh NoC. Failures may also be present in other components of a system, however, such as in computational components (core), storage components (cache slice), power supplies, and clocks. The failure of computational/storage components may result only in degraded system performance. A failure in any of the NoC components (link or switch) will be more harmful as the NoC provides the communication backbone to connect multiple components on the chip. Indeed, such a failure may become a single point of failure, potentially causing the entire system to fail.

* Corresponding author.
 E-mail addresses: rimpybishnoi@gmail.com (R. Bishnoi), vlgaur@gmail.com,
gaurms@gmail.com (M.S. Gaur), mz@ecs.soton.ac.uk (M. Zwolinski).

The unpredictable locations of these failures may also harm the regularity of the underlying topology. More specifically, a regular 2D mesh is more likely to become irregular due to these failures. Thus, a NoC must continue its operation even if the network becomes irregular.

In the event of failure, it is possible that the current routing function is unable to offer a path between each pair of nodes. Hence, it needs to be replaced with one offering full connectivity. To ensure this requirement, it is necessary that the underlying routing implementation framework should provide the flexibility to reconfigure the old routing function with a new routing function. This expectation demands a flexible routing framework for NoCs, which can efficiently support any irregularity generated in the initial regular structures.

A finite-state machine (FSM) based implementation [5], is very efficient in terms of both area and latency but is topology- and routing-dependent. Such methods may not work in the event of any failure-induced irregularity. Implementations based on forwarding tables (source, distributed) [6,7,8] can be used to support any irregular topology. Though tables provide the flexibility to work with any irregular topology, they do not scale well in terms of area, power, and performance.

Recently, uLBDR (Universal Logic Based Distributed Routing) [9] has been proposed as a scalable and efficient routing implementation method to support irregular 2D mesh topologies. It combines the scalability of traditional FSM based implementations and the flexibility of table-based implementations. Despite its scalability and flexibility, uLBDR adds some limitations to the switch design that restrict its applicability to VCT (Virtual Cut-Through) switching only and it requires a particular arbiter. In addition to this, uLBDR also increases the network traffic by replicating messages [9].

1.1. Motivation and objectives

Most routing implementation techniques that have been proposed to efficiently capture the routing function for irregular topologies are either table based [6,7,8] or require a complex constrained switch design [9]. This fact motivates us to design an efficient routing implementation technique to address irregular topologies.

The main objective of this paper is to propose a fault-tolerant routing implementation that can handle the irregularities resulting from failures in regular 2D meshes. The aim is to provide the routing framework with the flexibility to reconfigure the routing function when the topology becomes irregular. In addition to this, the proposed implementation should achieve high performance (low latency, high throughput) while keeping the cost and complexity low.

1.2. Contributions

To achieve the aforementioned objective, we propose a routing implementation technique. Our proposal is capable of handling any irregularity induced in a 2D mesh because of any number of 1-link and 2-link and most multiple link failures. The novelty of our proposal lies in the fact that it neither utilizes any routing tables nor imposes any restrictions on switch design. Additionally, it is deadlock and livelock free. Having these properties, the proposed method does not require any additional hardware (virtual channels). While developing an efficient routing implementation, the proposed mechanism is designed to achieve the following:

1. Improved fault coverage: This is a measure of the reliability of the proposed approach. It is defined as the percentage of irregular topologies (generated due to link failures) supported by a particular routing implementation. The proposed implementation guarantees full coverage in the case of any number of 1-link and 2-link failures provided there exists a deadlock-free path from source to destination. We have also extended the scope of the proposed mechanism to handle multiple failures.

2. Performance: In the case of a regular 2D mesh structure (without any failures), this approach maintains low latency and achieves high throughput similar to the baseline technique. In the cases of irregular structures being generated due to failures in an initial 2D mesh, the proposed approach gracefully degrades the network's performance.

3. Minimized area and power overhead: For applicability in the NoC domain, a low value for both area and power is necessary while keeping the design scalable. The proposed method requires respectively 14% and 16% less area and power than other state-of-the-art logic-based solutions proposed for irregular topologies and it keeps the design scalable.

The rest of the paper is organized as follows: In Section 2 we discuss the related work. In Section 3, we describe our proposed mechanism. We present experimental results analysis in Section 4. Finally, concluding remarks are presented in Section 5.

2. Related work

Considerable research has been carried out on resilient NoCs. In this section, we discuss previous work targeting permanent failures (either links or switches). We classify solutions based on their routing implementation mechanism, indicating whether they support fault tolerance or not.

1. FSM based implementations: An FSM based implementation [5] of a routing algorithm is very efficient in terms of both performance and area but is topology sensitive. Any failure in the network might convert a regular topology into an irregular one. There is a large body of work on fault tolerance based on FSM implementations providing only partial support for irregular topologies. They are not able to handle all possible sets of irregular topologies derived from a specific set of failures. Some of them utilize virtual channels to implement fault-tolerant routing algorithms.
The adaptive routing algorithm proposed by Linder and Harden [10] supports single node failure by doubling the number of virtual channels along one dimension, resulting in underutilized resources. The reliable router [11], proposed by Dally et al. for 2D meshes, can support irregular topologies derived from a single link or node failure. To achieve this, it also utilizes a large number (five for each physical link) of virtual channels.
Similarly, several other prior works [12,13] utilize additional virtual channels to provide fault tolerance and result in improved coverage support. In [14], Glass and Ni showed how to modify turn model-based routing algorithms to provide $(n$-$1)$ fault tolerance for n-dimensional meshes without using any virtual channels.

2. Table based implementations: Implementations based on forwarding tables (either at source or distributed) are not sensitive to topology change and offer the flexibility to support any irregular topology derived from any set of link failures. Schonwald et al. [15] proposed a table-based Force-Directed Wormhole Routing (FDWR) which is based on hop distance to the destination from the current switch. Though it supports all irregular topologies as tables are deployed, it results in a large packet processing time at routers. Large table size is also an issue with FDWR
Feng et al. [8] proposed a Fault Tolerant Deflection Router (FTDR) able to handle permanent and transient faults. They also proposed an improvement over the basic FTDR algorithm, named FTDR-H, which tries to reduce the table size by dividing the network into regions. However, the problem of increasing the number of tables with network size remains the same. A few other works on fault tolerance [6,7] also employ tables either at the source or a router. A number of resilient routing algorithms based on network reconfiguration have also been proposed such as Immunet [16], Vicis [17] and ARIADNE [18].
Immunet [16] routes packets towards their destination using a fully adaptive routing algorithm. A new ring is used for network reconfiguration and deadlock freedom that deterministically routes