J. Parallel Distrib. Comput. 101 (2017) 92–102

Contents lists available at ScienceDirect

J. Parallel Distrib. Comput.

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

A new power efficient high performance interconnection network for many-core processors

Faiz Al Faisal [a,*], M.M. Hafizur Rahman [b], Yasushi Inoguchi [c]

[a] School of Information Science, JAIST, Nomi-Shi, Ishikawa, Japan
[b] Department of Computer Science, KICT, IIUM, Malaysia
[c] Research Center for Advanced Computing Infrastructure, JAIST, Nomi-Shi, Ishikawa, Japan

HIGHLIGHTS

- A new power efficient interconnection network has been introduced as '3D-TESH'.
- 3D-TESH shows about 45∼52% better static performance than 3D-Torus.
- Achieves similar dynamic performance in comparing with 3D-Torus even at on-chip level.
- 3D-TESH requires about 14.81% less power usage than 3D-Torus (16 384 nodes).
- Guarantees about 41% less router power usage than 5D-Torus (on-chip level).

ARTICLE INFO

Article history:
Received 22 July 2015
Received in revised form
8 September 2016
Accepted 14 November 2016
Available online 27 November 2016

Keywords:
Interconnection network
3D-TESH network
Routing algorithm
Static network performance
Power estimation

ABSTRACT

Next generation high performance computing will most likely depend on the massively parallel computers. The overall performance of a massively parallel computer system is heavily affected by the interconnection network and its processing nodes. Continuing advances in VLSI technologies promise to deliver more power to individual nodes. However, the on-chip interconnection networks consume up to 50% of the total chip power and off-chip bandwidth is limited to the maximum number of possible out going physical links. In addition, the long wiring and low performance of communication network overwhelm the benefit of parallel computer system whereas it increases total cost. In this paper, we propose a new interconnection network that reduces the problems of high power consumption, long wiring length and low bandwidth issues. We have measured the static network performance and required power consumption of our proposed '3D-TESH' interconnection network and compared the performance with other networks at different levels of hierarchy such as inter-chips, inter-nodes and inter-cabinets. 3D-TESH network has achieved about 52.08% better diameter and about 45.71% better average distance than the 3D-Torus network with 12.61% less router power usage at on-chip level. Furthermore, 3D-TESH requires about 41% less router power usage than 5D-Torus at the on-chip level.

1. Introduction

Sequential computers are no longer sufficient to meet the increased computational demand and already reached into the saturation point due to the scaling difficulties of uniprocessor architectures. In addition, high performance computer (HPC) has become the main focus for next generation of computers. Even today, the molecular research in health, nuclear energy, organic simulation and weather forecasting highly depends on parallel computers. Hence, the need for massively parallel computers (MPC) is increasing progressively. MPC with thousand of nodes is commercially available with tera-flops performance. However, efforts have been made to build MPC systems with millions of nodes for peta-flops or for exa-flops performance. To facilitate millions of nodes, the interconnection networks are one of the key elements [8]. Interconnection network acts as a communication path between one node to another. However, to leap into the next level, there are few key constraints such as low network performance, high cost-performance ratio, high power consumption, low throughput and latency [26].

* Correspondence to: School of Information Science, JAIST, 923-1211 Nomi-shi, Ishikawa, Japan.
E-mail addresses: ffaisal@jaist.ac.jp (F.A. Faisal), hafizur@iium.edu.my (M.M.H. Rahman), inoguchi@jaist.ac.jp (Y. Inoguchi).

http://dx.doi.org/10.1016/j.jpdc.2016.11.007
0743-7315/

In MPC systems, the number of physical links is also a major concern [24]. Hence, the design of interconnection network has become a vital issue. Early research for interconnection networks was focused on the conventional networks. However, conventional network like the torus network [22] shows better performance than the mesh network [25] due to the tori connections, but consumes more electrical power than the mesh network due to the increased wraparound channels. One of the recently adapted network is the TOFU interconnection network [3], which requires 10 outgoing links for each node; causing a higher cost than the other networks and the scaling at higher level increases the power usage. Hence, it is important to find a new interconnection network by considering the key constraints as well as reducing the overall power consumption than the other existing networks.

Hierarchical Interconnection networks (HIN) [10] are a cost-effective way to interconnect a large number of processing nodes. Though many hypercube based HINs have been available, many HIN does not have the higher scalability. In some earlier research, HIN network like TESH [16,13] was introduced and showed better results than the conventional networks due to its high communications performance. However, TESH had failed to meet the high performance demand for higher scalability. In this paper, we introduce a new interconnection network named as '3D-TESH', and we are able to show that 3D-TESH has achieved better performance over millions of nodes and better than that of TESH due to its multi-dimensional structure at the basic module.

The remaining of the paper describes about the architecture of 3D-TESH network, reviews the routing algorithm, shows the static and dynamic performance analysis and then finally evaluates the power estimation for 3D-TESH.

2. Related works

Research for traditional interconnection networks have already been closed. Traditional networks are designed with flat structures, have the problem of poor network performance and require high electrical power for large scale systems. In contrast, modern supercomputers solely depend on hierarchical network connections for off-chip levels between chip to chip, node to node and so on. However, an 80-tile teraFLOPS processor arranged as 8×10 2-D array of floating-point cores having interconnected through the 2D-mesh network and operated at 4.27 GHz with 65-nm CMOS process requires 97 W of electrical power [27]. In spite of on-chip connections, off-chip connections also have a major impact on system power usage. In traditional networks, such off-chip connections are not being optimized. On the other hand, HINs could be a possible solution for these concerns due to their hierarchical structure pattern at different levels of hierarchy.

HINs are useful and cost-effective over the conventional networks. However, many hierarchical interconnection networks have already been introduced earlier, such as TESH [16,13], HTN [11], TTN [14], and STTN [12,4]. As we already discussed about TESH network, the TESH network has not been a better choice than the others. However, TTN and STTN networks, based on 2D-Torus network at the on-chip level, also have the problem of poor network performance at high-scalability over one million of nodes. As the processing power of single core is limited, the scalability for networks is also an important issue for next generation exa-scale systems. Furthermore, HTN considers the 3D-Torus network at the on-chip level, which also shows much poor network performance in comparison with 5D-Torus network. Nevertheless, Friedman et al. [24] confirm that 3D NoC is the most desirable architecture for next generation supercomputers due to its 40% and 36% performance improvement and a decrease of 62% and 58% in power dissipation over traditional 2D NoC with network size of $N = 128$ and $N = 256$ (through the reduced vertical node to node connection). Under those circumstances, the key motivation for our research is to consider a 3D NoC based HIN. Hence, in this paper, we introduce a new hierarchical interconnection network 3D-TESH.

3. Architecture of 3D-TESH network

Hierarchical interconnection networks are desirable over the traditional networks. 3D-TESH is a HIN which consists of multiple basic modules (BMs) that are hierarchically interconnected for higher levels.

Definition. A BM of 3D-TESH(m, L, q) network, by definition is built using 2^m number of 2D-TESH$(2^m \times 2^m)$ basic modules, where m is a positive integer, has L levels of hierarchy and q is used for the inter-level connectivity.

In 3D-TESH, a BM, which is also similar to 3D-Mesh network, consists of $(2^m \times 2^m \times 2^m)$ connected processing elements (PEs) having (2^m) rows and $(2^m \times 2^m)$ columns. Fig. 1 shows the BM for 3D-TESH with $m = 2$ which defines the size of the BM as $(4 \times 4 \times 4)$. Similarly, if $m = 3$, then the size of the BM becomes $(8 \times 8 \times 8)$ network with 512 processing nodes.

3.1. Basic module

Each $(2^m \times 2^m \times 2^m)$ BM of a 3D-TESH network has 2^{2m+2} free ports for its higher level of interconnections. In 3D-TESH network for each higher level of interconnection, a BM uses $2^m \times 4 \times (2^q) = 2^{m+q+2}$ of its free links, where $2(2^{m+q})$ free links for vertical interconnections and $2(2^{m+q})$ free links for horizontal interconnections. Here, $q \in 0, 1, \ldots, m$ is defined as the inter-level connectivity; $q = 0$ leads to the minimum inter-level connectivity, whereas $q = m$ leads to the maximum inter-level connectivity. As from Fig. 1a $(4 \times 4 \times 4)$ BM has $2^{2\times2+2} = 64$ free ports. If we choose $m = 2$ and $q = 0$, then $(2^{2+0+2}) = 16$ of the free ports and their associated links are used for each higher level interconnection, 8 for horizontal and 8 for vertical connection. Fig. 1 also shows 8 horizontal connections and 8 vertical connections for each higher level network (Level-2 – Level-5). However, if the maximum value of L is 2 or 3, the number of vertical_in, vertical_out, horizontal_in, horizontal_out connections are more than four; this means that the inter-level connectivity (q) is increasing. Similarly, if $m = 2$ and $L = 3$, then 8 links can be used for each vertical_in, vertical_out, horizontal_in, horizontal_out connections.

3.2. Higher level network

Higher level of 3D-TESH network is built by the recursive interconnection of the immediate lower level of sub-networks. Fig. 2 illustrates the higher-level interconnection of 3D-TESH. A Level-2 network can be formed by $(2^{2\times2})$ 16 BMs (16 Level-1 3D-TESH network). Similarly, a Level-3 3D-TESH network can be formed by interconnecting 16 Level-2 sub-networks. As we have considered m is 2, then the number nodes at Level-2 network can be defined by $N = (2^{2mL} \times 2^m)$. Hence, the number of nodes at Level-2 network is $N = (2^8 \times 2^2) = 1024$.

The maximum level for 3D-TESH can be built by a $(2^m \times 2^m \times 2^m)$ BM is $L_{\max} = 2^{m-q} + 1$. If the inter-level connectivity, $q = 0$ and $m = 2$, $L_{\max} = 5$; Level-5 is the highest possible level. This paper also considers $m = 2$ and $q = 0$. Hence, we showed the static network performance up to Level-5. The maximum number of nodes in each level of network can be defined as $N = (2^{2mL} \times 2^m)$. If $m = 2$ and $L = 2$, then $N = 1024$. Similarly, a Level-3 3D-TESH network consists of $N = 2^{2\times2\times3} \times 4 = 16\,384$ nodes, which is equal to 16 Level-2 3D-TESH networks. Table 1 generalizes the various aspects of 3D-TESH.