VOLUMETRIC DEGENERATIVE ROUTING FOR 3D NETWORK-ON-CHIP

A Thesis
Submitted to the Graduate Faculty
of the
North Dakota State University
of Agriculture and Applied Science

By

Druhin Bala

In Partial Fulfillment of the Requirements
for the Degree of
MASTER OF SCIENCE

Major Department:
Electrical and Computer Engineering

November 2014

Fargo, North Dakota

North Dakota State University

Graduate School

Title

VOLUMETRIC DEGENERATIVE ROUTING FOR 3D NETWORK-ON-CHIP

By

Druhin Bala

The Supervisory Committee certifies that this ***disquisition*** complies with North Dakota State University's regulations and meets the accepted standards for the degree of

MASTER OF SCIENCE

SUPERVISORY COMMITTEE:

Dr. Chao You

Chair

Dr. Jacob Glower

Dr. Kendall Nygard

Approved:

12/1/2015

Date

Dr. Scott C. Smith

Department Chair

ABSTRACT

As we reach the limits of scaling down of circuits, Three Dimensional Integrated Circuits (3D ICs) offer a very promising opportunity to keep on increasing the processing capacities and speed. In a Multi-Processor System-on-Chip (MPSoC) based embedded system with Network-on-chip (NOC) as the communication architecture, routing of the traffic among the Processing Elements (PEs) contributes significantly to the global latency, throughput and energy consumption. Almost all prior studies have focused on 2D NOC designs. The field of 3D integration is relatively new and has emerged to provide an alternate solution for high performance computation.

This paper introduces a new routing algorithm which aims to improve performance characteristics of conventional existing algorithms. Volumetric Degenerative Routing, as proposed in this paper, reduces maximum delay by as much as 40%.

ACKNOWLEDGEMENTS

This thesis represents my endeavor during the past few years in my graduate study life. Though the final result is far away from perfection, it deserves my devotion and hardwork.

I would first express my appreciation to Dr. Chao You, who has been acting as my advisor with my research, helping me from the very beginning all the way till this very day, entertaining my questions on research, entrepreneurship and life in general. I am also very thankful for Dr Kendall Nygard and Dr. Jacob Glower for being my committee members and supervising my final examination. I would also like to thank Dr. Barabanov. All of you have made a deep impact in my studies through your classes and guidance and I will be forever grateful.

Last but not least, my parents, my brother, my close friends Tanvi, Maximilian, Ryan, Stephan and lab mates thank you for your support.

DEDICATION

To all of my teachers and professors I have had the pleasure of meeting in my life.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS

PE……………………………………..Processor Element

2D……………………………………..2 Dimensional

3D……………………………………..3 Dimensional

NOC…………………………………..Network-on-Chip

TSV……………………………………..Through Silicon Via

VDR…………………………………..Volumetric Degenerative Routing

CH……………………………………..Cluster Head

ZXY…………………………………..Z direction, X direction, Y direction

XYZ…………………………………..X direction, Y direction, Z direction

1. INTRODUCTION

1.1. Overview

The speed of our processors has always been dependent on our ability and ingenuity to shrink the size of the transistors. The scaling has enabled us to accommodate the growing needs of component count and complexity of calculations. However, we are reaching the limits of how small we can build our transistors. Moving towards sub-20nm technology has a significant challenge to the design and manufacturing techniques. One of the greatest challenges of the present day is obtaining sub-20 nm CMOS technology and a higher computing power of our chips.

Secondly, we have to consider the number of components we can fit in one chip without having the performance suffer from other factors like power generation and heat consumption. Now, as the number of components keep on increasing, the architecture of the interconnect network comes into play and impacts the performance and heat generation of the system as a whole. Bus-based systems are no longer dependable architectures for System-on-chips because they are not massively scalable and do not provide efficient parallel integration, low global latency and low throughput. Network-on-chip (NOC) is a suitable successor to bus-based systems.

This has necessitated us to think of novel ways of thinking about the architecture itself. 3D integration Network-on-Chip has been identified as a suitable successor to meet the demand for higher performance chips.

1.2. Network-on-Chip

Network-on-chip is a communications paradigm where different components of an integrated circuit like processors and memory are connected using a public network with switching packets on a hop-by-hop basis. The NOC public transportation network and each of the multiple point-to-point data links are interconnected by routers. Traditionally, integrated circuits have dedicated point-to-point connections with one wire reserved for one signal. With the public network NOC interconnects provide for high-bandwidth, scalability, better performance, simpler design, lower power, noise and predictable speed.

But as systems grow larger and larger to hundreds of cores the performance of an NOC starts to decline. As the number of cores increase, the number of hops needed to reach the destination node increases. The length of the minimal path increases. As a consequence, a secondary problem of high latency is incurred. Furthermore, the performance when it comes to heat dissipation falls with the center of the NOC often creating Hotspots. The proposed solution to this has been 3D Integrated Network on Chips.

1.3. 3-D Network-on-Chip

The previous section talked about Network-on-chip in 2-dimensions (2D). 2D NOC architectures have been well studied and researched over the last few years. However, a 3-dimensional Network-on-Chip is a very new topic of research with immense possibilities. They are an attractive option to existing 2D NOCs because they offer -

1. Enhanced functionality
2. The ability to encapsulate different technologies.

A 3D Network-on-Chip is made by stacking layers of integrated chips and connecting the layers with vertical Through-Silicon-Via (TSV) interconnects. These interconnects pass completely through a silicon wafer or die. Most studies on 3D Network-on-chips have been done through simulations because the manufacturing techniques for such are still evolving to meet the precision standards. Currently the different approaches to creating the vertical interconnects are:

Wire bonded - Done at the die level with a vertical pitch of 35 to 100mm. This has been the most prevalent approach so far. Individual dies are connected with wires in a stack. One of the major shortcomings of this approach is these wire bonds can only be done at the chips outer edges, and as a consequence it limits the density of the chips that we can pack in. The manufacturing of wire-bonded 3D NOCs are stressful on the chips because of the heat and pressure involved. Metallic pads are often used to minimize the stresses during manufacture and to keep the integrity of the chips.

Microbump (Face-to-face) - Done at the die level with a vertical pitch of 10-100mm. This technique uses solder or gold bumps to make the connections. These microbumps are made on the surface of the chip. This has a few major advantages over the wire-bonded technique -

a. It offers a higher density of vertical interconnects.
b. The physical stresses on the die is far less.

By creating face-to-face microbump bonds, the distance between two dies are greatly minimized at the same time.

Contactless (Inductive or Capacitive) - Done at the die level with a vertical pitch of 50 to 200mm. This technique connects two different chips with either capacitive or inductive coupling. The manufacturing process is simpler and more inexpensive than the previous two. The biggest drawback in this technique is that it requires the two dies to be only face-to-face when bonded and hence is limited to only two dies. Also, the distance between the two dies must be small enough that the coupling has a strong enough effect for a signal to be transmitted.

Through Silicon Via : Done at the wafer level with a vertical pitch of 50mm.

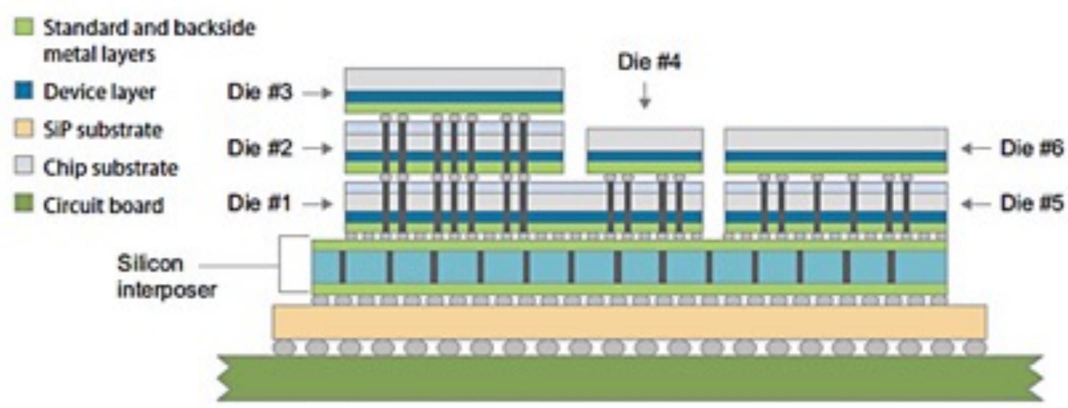

Figure 1. 3D TSV interconnect bonding [1]

Through Silicon Vias is the most promising of the various approaches, however the cost is also the most. The first pair of wafers are stacked face-to-face. The next wafers are then placed as back-to-face or again face-to-face according to the number of wafers being stacked and the orientation of the system. The advantages provided by 3D Network-on-Chips are manifold -

1. Smaller form factor
2. Reduced wire length
3. Improved bandwidth and throughput

1.4. Motivation

Now, as the number of components start to increase, the routing algorithm used to transfer flits becomes important and plays a major part in reducing global delay, power consumption, heat dissipation and throughput. Before we begin, let us also take a look at the two basic types of routing algorithm -

1. Static
2. Dynamic

Static routing is lightweight and fast. Each of the routers employed in a static routing algorithm has a fixed table that it looks up when propagating flits forward. The advantages are that it is easy and fast to implement with little or no overhead. The disadvantage is that it does not take into account if links are broken or network congestion.

In dynamic routing each of the routers calculates the next node in runtime. Various strategies could be used about this calculation -

1. It could be done at every single turn that flits need to be forwarded
2. It could be done according to a specific time cycle
3. The controller that is keeping an eye out for broken links pushes a broadcast update to all the routers as and when it gets to know.

The advantage to dynamic routing is that it is fault-tolerant and can keep track of network congestion. The disadvantages are that the overhead is greatly increased and it is susceptible to deadlocks or livelocks which could take the system into a never-ending cycle. In this thesis, we introduce a new static routing algorithm called Degenerative Routing Algorithm which aims to improve the performance of the 3D NOC.

2. RELATED WORK

There has been very little work in this field [5] - [8] where the latter two are adaptive algorithms and VDR is a static routing algorithm. Parischa et al. [7] focuses on reliability of the flit communication in an NOC. Ville Rantala et al.[8] tries to predict congestion spots and divert traffic elsewhere across the network. Importance in these adaptive routing has been placed on arrival rate with given link faults and there is no stress on the delay that can be produced in the network. Previous work on routing algorithms have mainly focused on a 2D NOC architecture [9]. There is the most standard case of XY algorithm in 2D NOC. XY algorithms route flits along the X-axis, until they reach the destination PE x coordinate. Next, the flits are routed in the Y-axis until they reach the destination PE. In a 3D NOC architecture, the base case is the ZXY or XYZ algorithm [10]. In ZXY, the flits are routed first along the Z-axis upto the layer of the corresponding PE and then XY routing is performed in that layer. Other conventional approaches include performing routing along the Z-axis to the required layer and then performing West-First, North-Last, NegativeFirst, Odd-Even algorithm. Viswanathan et al. put forward a new architecture for 3D NOC and a hierarchical routing scheme to transfer flits [11]. The architecture is that each node in a layer is a Cluster Head (CH) and is connected to four PEs. For a flit to reach any local PE, it must pass through a CH. The routing algorithm in this paper proposes a hierarchical scheme where a flit is transferred to the desired layer, then the desired cluster head and finally the intended PE. However, such approaches do not efficiently use network links and lead to the generation of unwanted hotspots and forwarding data through the same routers in the network. Some routers end up being always busy and some remain idle. VDR is a new approach to routing and it offers significantly better results.

3. ALGORITHM DESIGN FOR VOLUMETRIC DEGENERATIVE ROUTING

A 3D NOC may be comprised of N arbitrary number of components or processing cores. Let this number N be obtained as :

$$N = X \times Y \times Z$$

where X,Y and Z represent the number of rows, columns and layers respectively in the 3D NOC. The smallest and basic building block of any such 3D NOC is a $1 \times 1 \times 1$ cubic lattice, called as Base Cubes (BC). Volumetric Degenerative Routing aims to create 3D diagonal routes by propagating through these cubic lattices, reducing the search space of the 3D NOC after every iteration. It is important to try to follow diagonal routes in the 3D space as these cover nodes with greatest diversity and the intersecting or overlapping nodes are also minimised. Diversity as in [7] is defined as the number of paths available from any given node to the destination. It is to be noted that the in a 3D NOC, the nodes with greatest diversity are always towards the center of the structure. The algorithm strives to make diagonal paths by making the following moves -

1. Traverse the Z axis by one node and then the X axis by one node.
2. Traverse the Z axis by one node and then the X axis by one node.

These two turn procedures are called alternatively to propagate through the cubic lattices and hence the 3D NOC in a diagonal fashion. However, if the flits reach the target integrated chip layer in the vertical direction, it makes a conventional XY routing to reach the target node. Using VDR, the number of common routers used, while generating routes between source-destination PE pairs, are greatly reduced. This contributes to reduced global delay, reduced maximum delay and better utilization of network bandwidth.

3.1. Pseudocode

The Pseudocode for the algorithm is described in this section.

Step 1: Get current PE ID and Destination PE ID. FlagZ=0, FlagXY=0.

Step 2: If Current PE is in the same layer as Destination PE, perform XY routing.

Step 3: If FlagZ=0 and current PE is not in same layer as destination PE, then Step 4; else Step 6.

Step 4: Make FlagZ=1. If Current PE is above destination PE then forward flit to the immediate PE in bottom layer. Go to Step 6.

Step 5: Transmit flits to the immediate PE in the above layer.

Step 6: If FlagXY=0, FlagZ=1 and Current X-coord (co-ordinate) is not equal to Destination X-coord then Step 7 else Step 9.

Step 7: FlagXY=1, FlagZ=0.

Step 8: If current X-coord¡destination X-coord then forward flits to West, else forward to East. Go to step 12

Step 9: If FlagXY=1, FlagZ=1 and Current Y-coord is not equal to Destination Y-coord then Step 10, else Step 12.

Step 10: FlagXY=0, FlagZ=0.

Step 11: If current Y-coord¡destination Y-coord then forward flits to North, else forward to South.

Step 12: Go to Step 2 and repeat until current PE is equal to destination PE

3.2. Comparison between XYZ and VDR

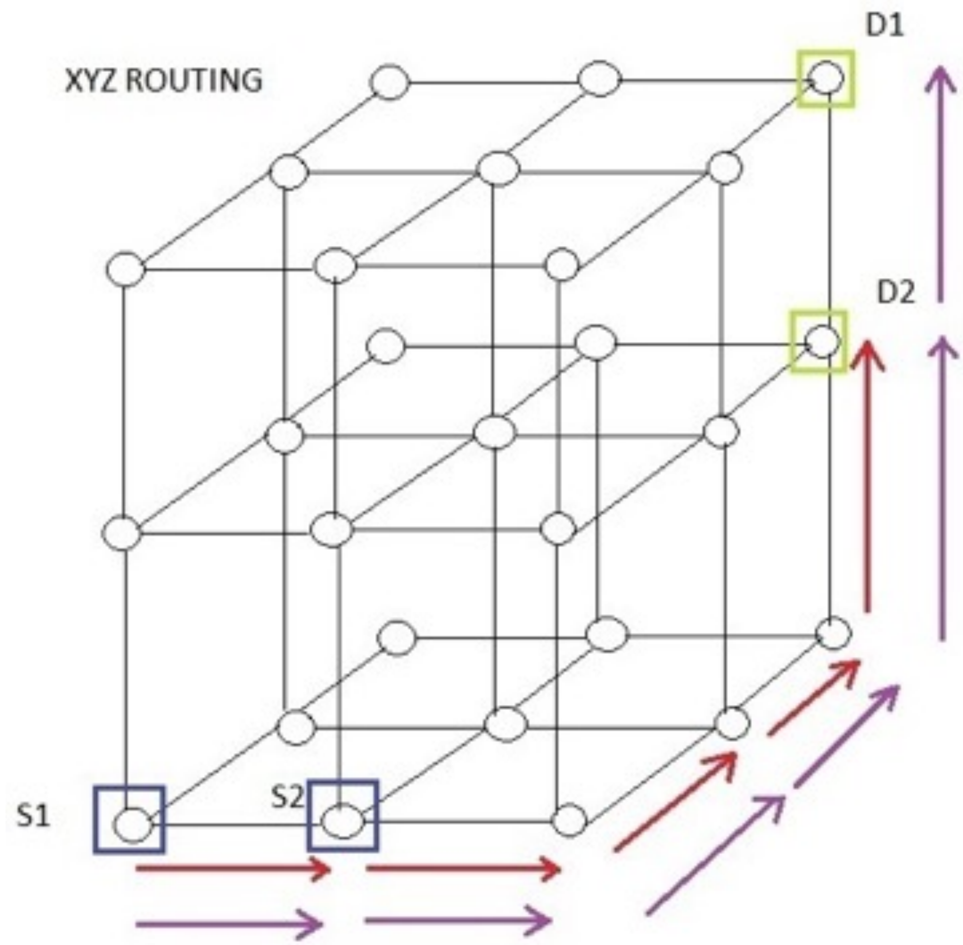

Figure 2. Example of routing with XYZ algorithm

Figure 2 depicts the case of a 3 × 3 × 3 3D NOC and the case of XYZ routing. S1 and S2 are sources and D1 and D2 are their respective destination PEs. Next, the following case is examined : S2 starts to transmit flits before S1 and continues to do so even when S1 wants to transmit. This would mean that flits from S1 would have to wait, until S2 is done transmitting its own packets. Moreover, the paths to be taken by the flits for S1-D1 and S2-D2, are equivalent except for one additional node PE/router for S1-D1. This means that the flits from S1 would have to wait at every node until the flits generated from S2 have been successfully forwarded from any particular node on that route. Evidently, this would lead to an inherent lag in the whole architecture.

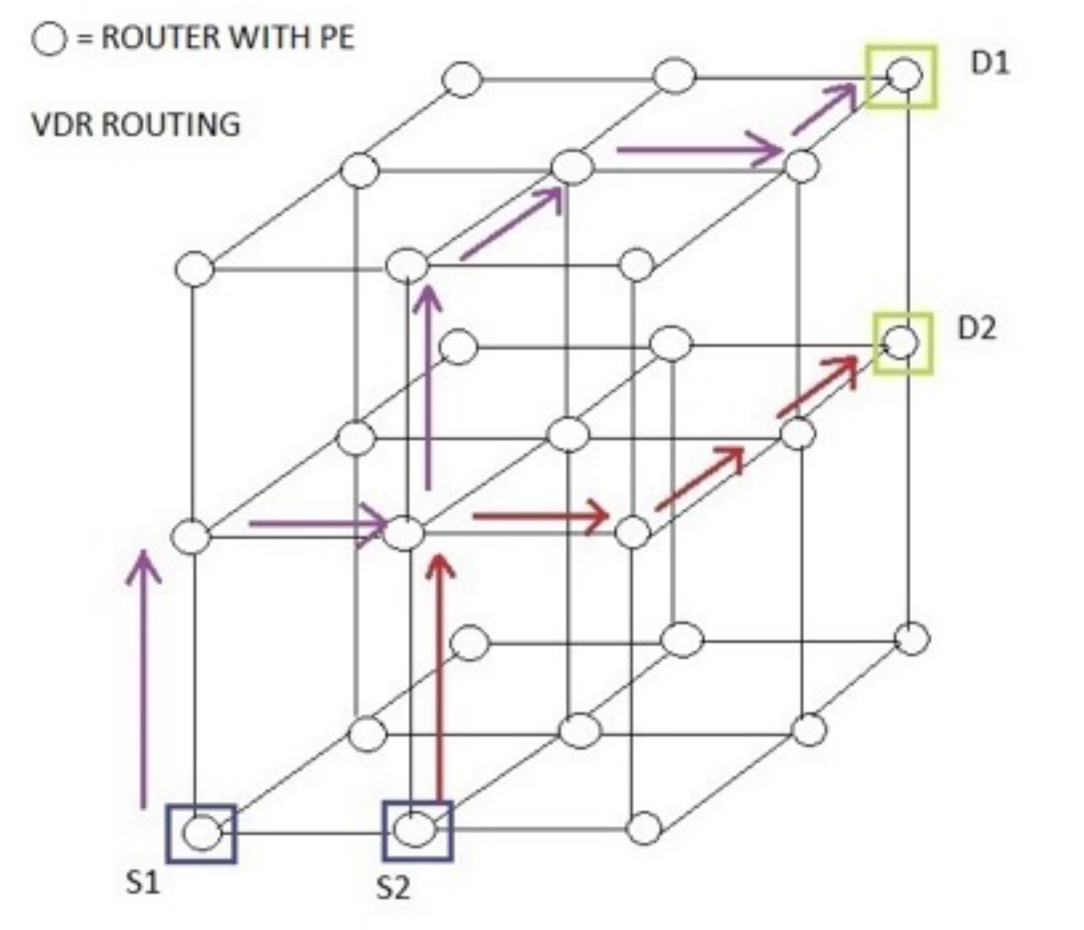

Figure 3. Example of routing with VDR

Figure 3. depicts the same 3 × 3 × 3 3D NOC but with VDR routing scheme. The diagram illustrates the path that the flits take according to VDR. As shown in the figure, the routers employed by the VDR routing algorithm to forward the flits from S1-D1 are different from the S2-D2 path, except for one common node. The max delay and global average delay is reduced because there is just one common router/PE that coming into play unlike in XYZ routing. Otherwise, the PEs utilized by the routing algorithm to forward the flits from S1-D1 are

totally different from that of S2-D2. As such, the flits from S1 do not need to wait to get transmitted. This eventually lead to reduced global delay and better utilization of network bandwidth. VDR provides increasingly better results with an increase in 3D NOC size as shown in the results section. This discussion can be easily extended to a comparison with ZXY, ZYX or YXZ routing algorithm, in the same manner.

3.3. Avoiding Deadlocks and Livelocks

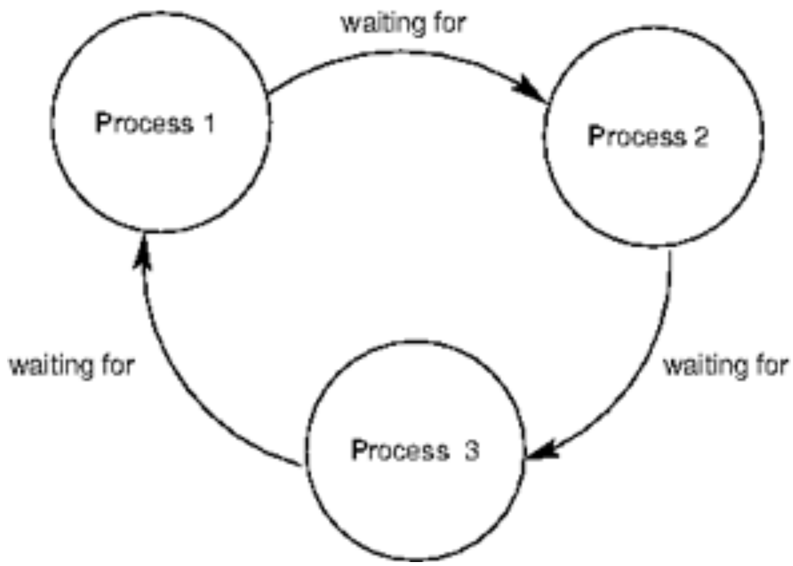

Figure 4. A deadlock situation with two or more competing actions waiting to finish

The deadlock problem was first cited in W.J. Dally's work in [12]. The deadlock problem in wormhole networks has been exhaustively worked upon [14] and [15]. VDR uses a Global Routing Table which is stored in each router of every PE. The table also has a list of invalidated entries for non-existent channels. Such nonexistent channels will exist in border and corner and side tiles of a 3D NOC. VDR employs virtual channels and buffers in routers in its implementation. Deadlock in VDR is further avoided by assigning each channel of any PE a unique number and allocating channels to packets in order. Furthermore, VDR is a dimension ordered routing scheme, where each flit of a packet is routed in one dimension at a time. The flits reach the proper coordinate in the designated dimension before proceeding to the next. A combination of the above factors and the enforcing of a strict order on the dimensions traversed, deadlock-free quality is guaranteed. VDR is deadlock free as is XYZ or ZXY routing.

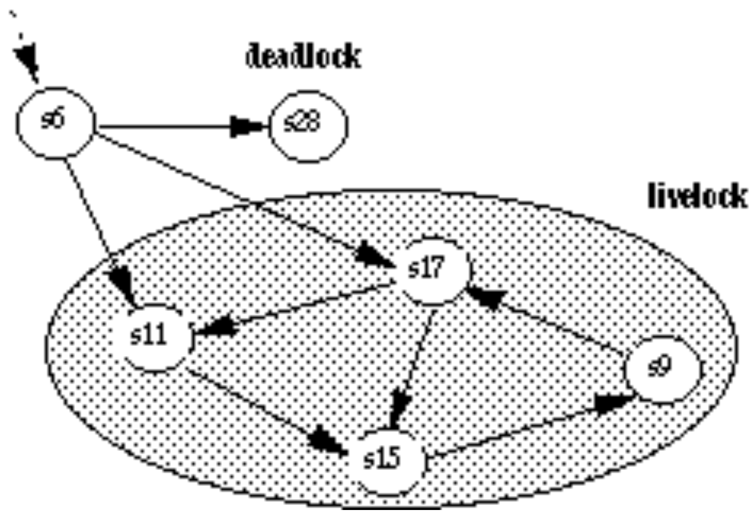

Figure 5. Situation of Livelock. The nodes s11, s17, s15 and s9 depict the case of livelock

A livelock is similar to a deadlock, except that the states of the processes involved in the livelock constantly change with regard to one another, none progressing. Lovelock is a special case of resource starvation; the general definition only states that a specific process is not progressing.

Livelock-free quality is guaranteed because this is a deterministic routing algorithm. Each router has its own global routing table. Whenever a flit arrives, the router looks up which output channel to use to forward the data by using the precomputed Global Routing Table. As such, flits are always reach their destination PE and avoid livelocks.

4. SIMULATION AND PERFORMANCE

The simulations were done on an Intel Core 2 Duo processor running at 2.35 GHz running Xubuntu. VDR and other algorithms were tested on a SystemC [16] based cycle accurate 3D Mesh simulator that was made by modifying NOXIM simulator [17].

4.1. Comparison with ZXY, West first, North last, Negative first and Odd-even

The first set of results compares VDR to ZXY which is the base case. It also provides a comparison with West first, North last, Negative first and Odd-even. For the last five cases, the flits are propagated along the Z axis first and then the algorithms are executed.The tests were conducted keeping the following constant:

1. A 3 x 3 x 3 NOC architecture.
2. Simulation done on 100000 clock cycles for each algorithm.
3. Random traffic distribution.
4. Probability of re-transmission of flits 0.01.
5. Poisson packet injection rate 0.01.

Global average delay and max delay are measured in clock cycles.

Table 1. Comparison of Algorithm type, global average delay, max delay and energy

Algorithm	Packets	Flits	Global Avg. Delay	Max Delay	Energy(mJ)
VDR	26786	160353	9.92892	66	256.519
ZXY	26640	159659	10.0542	83	257.064
West first	26734	160147	10.0243	83	257.762
North last	26644	159722	10.1152	77	257.649
Negative first	26653	159701	10.1348	80	257.804
Odd even	26672	159809	10.0767	93	260.755

4.2. Comparison with varying Packet Injection Rate and Traffic Distribution

This next section compares the performance of VDR with respect to ZXY routing when the Packet Injection Rate (PIR) and Traffic distribution is changed. For these tests, Poisson and Bursty type PIR were used. The traffic distributions schemes used were Random and Transpose. The following were kept constant during the tests:

1. A 3 x 3 x 3 NOC architecture.
2. Simulation done on 100000 clock cycles for each algorithm.
3. Probability of re-transmission of flits 0.01.

Global average delay and max delay are measured in clock cycles.

Table 2. Bursty packet injection rate = 0.4 (Max=1) and random traffic distribution

Algorithm	Flits	Global Avg. Delay	Max Delay	Energy (mJ)
VDR	768108	42016.4	85803	1547.78
ZXY	767545	42014.3	86316	1551.81

Table 3. Poisson packet injection rate = 0.4 (Max=1) and transpose traffic distribution

Algorithm	Flits	Global Avg. Delay	Max Delay	Energy (mJ)
VDR	1039502	41076.9	96870	1782.24
ZXY	1039498	41118.0	97125	1782.80

4.3. Comparison with ZXY with varying network architecture

The following set of results compare VDR to ZXY with varying network architecture. The network architecture is varied as - a. 5 x 5 x 5 3D NOC b. 10 x 10 x 10 3D NOC c. 15 x 15 x 15 3D NOC. The tests were conducted keeping the following constant:

1. Simulation done on 100000 clock cycles for each algorithm.
2. Random traffic distribution.
3. Probability of re-transmission of flits 0.01.
4. Poisson Packet injection rate 0.01.

Global Avg. delay and max delay are measured in clock cycles.

Table 4. 5 x 5 x 5 3D NOC architecture

5 x 5 x 5	Flits	Global Avg. Delay	Max Delay	Throughput (flits/cycle/ip)	Energy (mJ)
VDR	743287	15.9953	100	0.0600636	1849.54
ZXY	743688	16.0603	168	0.060096	1860.88

Table 5. 10 x 10 x 10 3D NOC architecture

10 x 10 x 10	Flits	Global Avg. Delay	Max Delay	Throughput (flits/cycle/ ip)	Energy (mJ)
VDR	5937569	34.6511	296	0.0599754	28.208
ZXY	5935550	34.6875	400	0.0599551	28.210

Table 6. 15 x 15 x 15 3D NOC architecture

15 x 15 x 15	Flits	Global Avg. Delay	Max Delay	Throughput (flits/cycle/ ip)	Energy (mJ)
VDR	20035940	69.1123	1727	0.0599654	142.886
ZXY	20042788	69.4266	1949	0.05999859	142.958

5. CONCLUSION

Routing in a many core or many component 3D NOC becomes essential in achieving high performance. The field of 3D Network-on-chips is vast. I have attempted to shed some light on one of its core parts - routing. The results conclusively demonstrate the promise of VDR in providing a reduced global average delay and reduced maximum delay in comparison to other traditional algorithms. Future research can revolve around streamlining and optimizing VDR for better results, a fault-tolerant VDR and a hierarchical VDR scheme to reduce the area footprint required. I envision that newer concepts will include hybrid interconnects like combining Photonic routing with TSV interconnects and using wireless connectivity also to connect the various components.

REFERENCES

[1] L. P. Carloni, P. Pande, Y. Xie, Networks-on-Chip in Emerging Interconnect Paradigms: Advantages and Challenges, IEEE, Proc of 3rd International Symposium on Networks-on-Chip (NOCS), Washington DC, USA, 2009.

[2] I. Loi, S. Mitra, T. H. Lee, S. Fujita and L. Benini A Low-overhead Fault Tolerance Scheme for TSV-based 3D Network on Chip Links, IEEE, Computer-Aided Design, IEEE/ACM International Conference. 10-13, San Jose, CA Nov. 2008.

[3] R. Holsmark, S. Kumar M. Palesi and A. Mejia HiRA: A Methodology for Deadlock Free Routing in Hierarchical Networks on Chip. Networks-on-Chip, 2009. IEEE, NoCS 2009. 3rd ACM/IEEE International Symposium. San Diego USA, 10-13 May 2009.

[4] T. Dumitras and R. Marculescu, On Chip Stochastic Communication. Proc. in Design, Automation and Test in Europe - Vol. 1 Page 10790, IEEE, Washington DC, 2003.

[5] R. Nakhjavani, A. Shahabi, S. Safari and Z. Navabi University of Tehran. A novel graceful degradable routing algorithm for 3D on-chip networks, ACM, INA-OCMC 12th Proc. of the 2012 Interconnection Network Architecture: On-Chip, Multi-Chip Workshop, New York, USA 2012

[6] R.S. Ramanujam, UCSD Bill Lin, UCSD. A Novel 3D Layer-Multiplexed OnChip Network, ACM, Embedded Systems Letters, New York, USA August, 2009.

[7] S. Parischa and Y. Zou, Colorado State University, A low overhead Fault Tolerant Routing Scheme for 3D Networks on chip, Quality Electronic Design, IEEE, (ISQED), 12th International Symposium, Santa Clara, CA USA, March 2011.

[8] V. Rantala , T. Lehtonen, P. Liljeberg and J. Plosila. Hybrid NoC with Traffic Monitoring and Adaptive Routing for Future 3D Integrated Chips, TUCS, ScientificCommons repository, 2008, http://tucs.fi/publications/view/?pub_id=inpRaLeLiPl08a

[9] A. M. Shafiee, M. Montazeri, and M. Nikdast, An Innovational Intermittent Algorithm in Networks-On-Chip (NOC), World Academy of Science, Engineering and Technology 4/5/2008, http://connection.ebscohost.com/c/articles/36317040

[10] M. A. Khan and A. Q. Ansari, A Quadrant-XYZ Routing Algorithm for 3-D Asymmetric Torus Network-on-Chip, IEEE, Emerging Trends in Networks and Computer Communications (ETNCC), Udaipur India, 2011.

[11] N. Viswanathan, K. Paramasivam and K. Somasundaram. Exploring Optimal Topology and Routing Algorithm for 3D Network on Chip, American Journal of Applied Sciences, http://thescipub.com/html/10.3844/ajassp.2012.300.308, USA, 2012

[12] W. J. Dally and C. L. Seitz. Deadlock-free message routing in multiprocessor interconnection networks. IEEE Transactions on Computers. 36, (5), USA, May 1987

[13] P. Ghosh, A. Ravi, and A. Sen, An Analytical Framework with Bounded Deflection Adaptive Routing for Networks-on-Chip. IEEE, VLSI Computer Society Annual Symposium. Lixouri, Kefalonia, 2010

[14] P. Mohapatra, Wormhole Routing Techniques for directly connected Multicomputer systems, ACM, ACM Digital Library, http://dl.acm.org/citation.cfm?id=292472, New York USA, 1998.

[15] X. Lin, P. K. McKinley, A H Esfahanian, Adaptive Multicast Wormhole Routing in 2D Mesh Computers, Springer Berlin Heidelberg, Parallel Architectures and

Languages Europe, Munich, Germany 1993.

[16] SystemC, SystemC Initiative, by Accellera Systems Initiative. http://accellera.org/downloads/standards/systemc

[17] Noxim, the Network-on-Chip Simulator developed at the University of Catania (Italy). https://github.com/davidepatti/noxim

[18] D. Park, S. Eachempati, R. Das, A. K. Mishra, Y. Xie, N. Vijaykrishnan and Chita R. Das. MIRA A multi-layered on chip routing architecture, Computer Architecture, IEEE, ISCA 08. 35th International Symposium. Beijing, China, 2008

[19] W. J. Dally, B. Towles. Principles and practices of interconnection networks, Morgan Kaufmann, 2003. Book ISBN:0122007514. http://dl.acm.org/citation.cfm?id=995703