

THESIS FOR THE DEGREE OF LICENTIATE OF ENGINEERING

DDRNoC: Dual Data-Rate Network-on-Chip

AHSEN EJAZ



Division of Computer Engineering
Department of Computer Science & Engineering
Chalmers University of Technology
Gothenburg, Sweden, 2018

DDRNoC: Dual Data-Rate Network-on-Chip

AHSEN EJAZ

Ioannis Sourdis	Advisor	Prof. at Chalmers University of Technology
Vassilios Papaefstathiou	Co-Advisor	Post Doctoral researcher at FORTH-ICS
Per Stenström	Examiner	Prof. at Chalmers University of Technology
Giorgos Dimitrakopoulos	Discussion Leader	Asst. Prof. at Democritus University of Thrace

Copyright ©2018 Ahsen Ejaz
except where otherwise stated.
All rights reserved.

Technical Report No 180L
ISSN 1652-876X
Department of Computer Science & Engineering
Chalmers University of Technology

Contact Information:

Division of Computer Engineering
Department of Computer Science & Engineering
Chalmers University of Technology
Se-412 96, Gothenburg, Sweden
Telephone: +46 (0)31-772 10 00
<http://www.chalmers.se/cse/>

Author's e-mail: ahsen@chalmers.se

Printed by Chalmers Reproservice,
Gothenburg, Sweden 2018.

DDRNoC: Dual Data-Rate Network-on-Chip

Ahsen Ejaz

*Department of Computer Science & Engineering
Chalmers University of Technology, Sweden*

Abstract

Networks-on-Chip (NoCs) are becoming increasingly important for the performance of modern multi-core system-on-chip. For various on-chip networks with virtual channel (VC) flow control, the slow control logic (VC and switch allocation logic) of the NoC routers limits the NoC clock period while their datapath (switch and link) possesses significant slack. This slack results in wasted performance potential of the datapath, limits the saturation throughput of the network and reduces its energy efficiency. The aim of this thesis is to improve NoC performance by eliminating this slack and removing control logic from the router critical path. To this end, this thesis presents the Dual Data-Rate (DDR) network architecture called the DDRNoC. It utilizes the NoC datapath twice within a clock cycle to forward flits at DDR. This not only exploits the slack present in the datapath but also requires a clock with period twice the datapath delay, thus removing the shorter control logic from the critical path. This enables the DDRNoC to achieve throughput higher than single data-rate networks. Moreover, the DDRNoC also employs lookahead signalling to reduce end-to-end packet latency. FreewayNoC, an extension to the DDRNoC supplements the DDRNoC with simplified pipeline stage bypassing to reduce the zero-load latency of packets in the network.

Implementation of the DDRNoC and FreewayNoC architectures require redesign of the switch allocation (SA) mechanism to resolve contention among competing flits by granting up to two flits access to each switch input and output port per clock cycle. It further requires separate paths for the propagation of lookahead control signals. FreewayNoC also requires implementation of multiple checks to guarantee conflict-free bypassing of the SA stage.

Physical implementation results using 28nm process technology show that DDRNoC and FreewayNoC have 5% and 15% area overhead, respectively, compared to a simple 3-stage network with VCs. Performance evaluation shows that for a 16×16 mesh network, FreewayNoC supports 25% higher throughput compared to current state-of-the-art NoC, ShortPath. Moreover, FreewayNoC achieves a zero-load latency which scales better than ShortPath and equally well with an ideal network that has no control overheads. For application driven traffic, FreewayNoC reduces average packet latency by 18% compared to ShortPath. Alternatively, low voltage implementation of the DDRNoC and FreewayNoC can be used to conserve power and improve energy efficiency at the cost of higher packet latency.

Keywords

Network-on-Chip, On-Chip Interconnect, System-on-Chip, Dual Data-Rate, Multiprocessor System-on-Chip, Chip Multiprocessors

List of Publications

Appended publications

This thesis is based on the following publications:

- [A] Ahsen Ejaz, Vassilios Papaefstathiou and Ioannis Sourdis, “DDRNoC: Dual Data-Rate Network-on-Chip”, *ACM Transactions on Architecture and Code Optimization (TACO)*, vol. 15, no. 2, June 2018.
- [B] Ahsen Ejaz, Vassilios Papaefstathiou and Ioannis Sourdis, “FreewayNoC: a DDR NoC with Pipeline Bypassing”, To appear *In Proceedings of the 12th IEEE/ACM International Symposium on Networks-on-Chip (NoCS)*, Torino, Italy, 2018.

Acknowledgment

I would like to express my gratitude to my PhD supervisor Ioannis Sourdis for giving me the opportunity to pursue a PhD degree under his supervision. Ioannis has been very patient and supportive during my research and has kept me motivated since I started on this path to a PhD degree.

I am grateful to Vassilis Papaefstathiou, my co-supervisor for his valuable comments, based on his deep insight of the field, during our meetings.

I would also like to thank my PhD research examiner, Per Stenström, and PhD follow-up committee members Jan Jonsson, Agneta Nilsson and Wolfgang Ahrendt for their constructive feedback and support over the years.

I am also very grateful to Lars Svensson for his useful critiques of this research work and for his support with setting up the physical design tools needed for this research.

Place-and-Route of the DDRNoC had been a daunting and a time consuming task, and it would have taken much longer had it not been for the useful impromptu discussions with Kevin Cushon, Christoffer Fougstedt, Erik Ryman and Victor Åberg. Thanks guys.

I would also like to thank the administrative staff of the department, Eva Axelsson, Monica Månhammar, Marianne Pleen-Schreiber, Nadja Johansson, Lars Norén, Michael Morin and Rune Ljungbjörn for their help and support.

I am also thankful to my colleagues in the Department of Computer Science and Engineering, Muhammad Waqar Azhar, Evangelos Vasilakis, Prajith Ramakrishnan Geethakumari, Albin Eldstal-Damlin, Stefano Ribes, Mehrzad Nejat, Petros Voudouris, Madhavan Manivannan and Pedro Petersen Moura Trancoso for our various discussions and for creating a positive work environment.

I would also like to thank my two office mates, who are also my seniors, Alirad Malek and Stavros Tzilis, for keeping me motivated and entertained.

This work has been partly funded by EUROSERVER project (grant agreement 610456), EuroLab-4-HPC project (grant agreement 371610) and Horizon 2020 Programme ECOSCALE project (grant agreement 671632).

Finally, none of this would have been possible without the unfaltering support, encouragement and prayers of my parents, my siblings and my wife. I am really grateful to have them by my side.

Ahsen Ejaz
Göteborg, September 2018

Contents

Abstract	iii
List of Publications	v
Acknowledgement	vii
1 Introduction	1
1.1 Problem Statement: Control Logic Limits NoC Performance . .	2
1.2 Thesis Objectives: NoC Performance Defined only by its Datapath	2
1.2.1 Throughput	3
1.2.1.1 Related Work	4
1.2.1.2 The Dual Data-Rate Network-on-Chip	4
1.2.2 Latency	6
1.2.2.1 Related Work	6
1.2.2.2 Control Forwarding	6
1.2.2.3 Pipeline Stage Bypassing	6
1.2.3 Energy Efficiency	7
1.2.3.1 Related Work	7
1.2.3.2 Clock Tree Power Reduction	8
1.2.3.3 Low Voltage Implementation	8
1.3 Contributions	9
1.4 Thesis Outline	10
2 DDRNoC: A Dual Data-Rate Network-on-Chip	17
2.1 Introduction	18
2.2 Related Work	19
2.3 The DDRNoC Architecture	21
2.3.1 Router Datapath	22
2.3.2 Timing	23
2.3.3 Zero Load Latency Analysis	26
2.3.4 DDRNoC Control	26
2.3.4.1 Virtual Channel Allocation	26
2.3.4.2 Switch Allocation	27
2.3.4.3 Flow Control and Minimum Buffer Size	27
2.3.5 Discussion	28
2.4 Implementation and Experimental Setup	29
2.5 Evaluation Results	31

2.5.1	Implementation Results	31
2.5.2	Evaluation using Synthetic Traffic	32
2.5.3	Evaluation using Application-Driven Traffic	34
2.5.4	System-Level Evaluation	34
2.6	Comparison	37
2.7	Conclusion	42
3	FreewayNoC: A DDR NoC with Pipeline Bypassing	49
3.1	Introduction	50
3.2	Related Work	51
3.3	The FreewayNoC Router Architecture	52
3.3.1	Router Datapath	53
3.3.2	Timing	55
3.3.3	Zero-Load Latency Analysis	56
3.3.4	Router Control	57
3.3.4.1	Combined VC and Switch Allocation	57
3.3.4.2	Pipeline Bypassing	58
3.3.4.3	Next Route Computation	60
3.3.4.4	Flow Control and Minimum Buffer Size	60
3.3.4.5	Control Path analysis	60
3.4	Evaluation	61
3.4.1	Experimental Setup	61
3.4.2	Implementation Results	62
3.4.3	Performance Evaluation	63
3.4.4	Energy Efficiency	63
3.5	Conclusions	66

Chapter 1

Introduction

Chip Multiprocessors (CMP) are one of the most promising solutions for supporting the continuous need for single chip performance improvement. Shrinking transistor geometries still allow more cores to be integrated on a die. These increasing number of cores with increasingly complex workloads place stringent requirements over the underlying communication fabric. On one hand, applications that exhibit small transfers at low loads are often more sensitive to network latency [1]. On the other hand, systems that run concurrent scale-out applications are more demanding, push the network close to its saturation point, and are more sensitive to network throughput [1, 2]. Moreover, power constraints prevent chips from fully utilizing all these cores at their maximum performance potential [3]. The on-chip interconnection network is a critical component for power efficiency as well [4] since roughly a sixth of the available chip power budget, if not more, goes to its interconnects [5–11]. As a consequence, the design of high-performance and low-power networks-on-chip (NoCs) is essential for many-core scaling.

Since the inception of NoCs in the early 2000s, the research community has developed various techniques to improve NoC performance while trying to reduce its power and silicon footprint. Many among them improve network throughput employing new allocation techniques [12], wider datapaths, richer network topologies [13], or multiple subnetworks [14–17]. Others attempt to lower packet latency by modifying the network topology [18, 19], the router architecture [10, 20, 21] or the routing algorithm [22]. Whereas some of these techniques also result in improved energy efficiency of the interconnect, many other techniques just focus on lowering the power consumption of the NoC by power-gating idle routers [23, 24], using low swing signalling [25] or implementing dynamic voltage and frequency scaling (DVFS) capabilities in NoC routers [8]. These networks require complex control logic to implement various allocation schemes which allocate available network resources to packets and flits propagating through the network. For conventional networks, the path delay of this control logic is longer than the datapath delay which leads to the NoC clock period being defined by its control logic. The datapath in such a case possesses some slack. Compared to an ideal network which operates at a clock rate defined only by the datapath delays, current 2D-mesh NoC architectures with virtual channel (VC) flow control still offer lower performance. This is

because an ideal network delivers more packets per node as a consequence of improved datapath utilization achieved by removing the slack present in the datapath. Higher network throughput obtained from increased datapath utilization also improves the performance of the MPSoC when compared to a conventional network.

This introductory chapter provides a brief overview of the work presented in this thesis. The remainder of this chapter is organized as follows: Section 1.1 presents the main problem statement of this research. Section 1.2 discusses the objective of this thesis. Section 1.3 summarizes the contributions of this thesis.

1.1 Problem Statement: Control Logic Limits NoC Performance

A NoC router has two sets of components: router datapath and control logic. Router datapath is composed of input/output link wires, VC buffers, crossbar and output registers. It is used to either store flits in NoC routers due to contention or unavailability of required resources or to move them forward through the network towards their destination. The control logic, on the other hand, is responsible for managing the flow of flits propagating through the network by allocating router datapath resources to them. It also resolves contention among flits competing to use shared router resources at the same time. It comprises next route computation logic (NRC), VC allocators (VA) and switch allocators (SA). NoC designs typically use pipelined architectures which divide the datapath and control logic into two or more stages to improve NoC clock frequency and its throughput. The performance of current NoCs with VC-flow control is limited by their control logic and is far from ideal performance that a NoC could achieve if its datapath were the only factor.

As shown in literature and confirmed by our experiments, the critical path of existing 2D-mesh routers with multiple VCs, is in its control logic [12, 26–28]. More specifically, the clock period of a typical 3-stage (NRC/VA/SA, switch traversal (ST), link traversal (LT)) router [26] is dictated by the VA stage of the router pipeline. Even for the current state-of-the-art NoC router with VC-flow control, ShortPath [21], a significant fraction of the clock period is based on the router control logic. Consequently, the datapath of these routers possess considerable slack. It is used only for a fraction of the clock cycle to forward the flit and remains idle otherwise. This reduces network datapath utilization, wastes link bandwidth and hinders the NoC from achieving its maximum performance potential by limiting its saturation throughput and increasing packet latency.

1.2 Thesis Objectives: NoC Performance Defined only by its Datapath

In this thesis, we target the slack in the datapath of VC based 2D-mesh NoC architectures to increase NoC datapath utilization. We aim at designing a NoC able to achieve performance (both throughput and latency) defined solely by the network datapath delays. The implications of such a network on throughput,

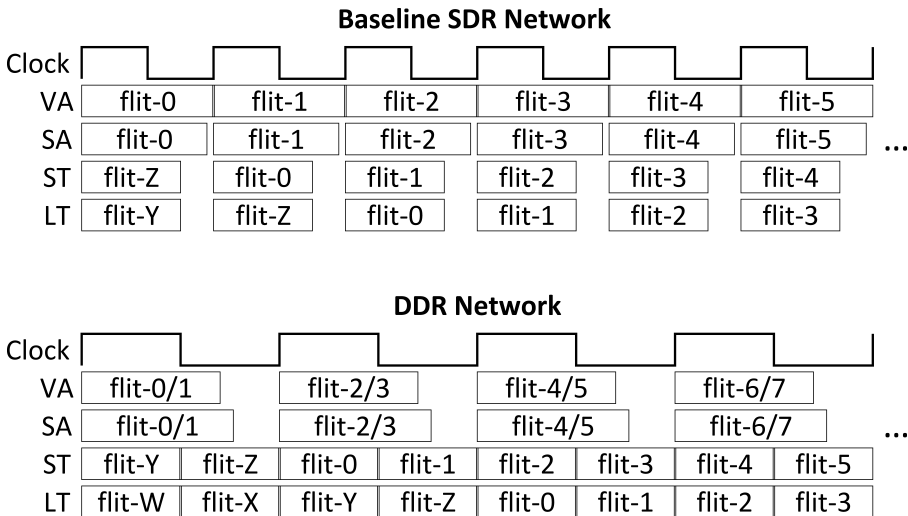


Figure 1.1: For a conventional SDR network, the clock period is completely (or partly) determined by the control logic. The DDR network extends the clock period to twice the datapath delay, implements control logic capable of allocating datapath resources to at most two flits per cycle and enables the two datapath stages (ST and LT) to propagate two flits per cycle. The clock period now is defined only by the datapath delays. This improves datapath utilization and enhances network throughput.

latency and energy efficiency are presented below along with a brief description of related works.

1.2.1 Throughput

One of the factors limiting the throughput of single data-rate (SDR) networks is the under-utilization of its datapath during a clock cycle. It is caused by an imbalance in the datapath and the control logic path delays of conventional SDR NoC routers as shown in figure 1.1. This is because the complex allocation schemes employed within the control logic of these routers either become the critical path or a part of the critical path sequentially shared with the datapath. In both these cases the clock period of the resulting NoC router is greater than the datapath delay and leads to datapath under-utilization during the clock cycle.

This thesis presents the DDRNoC network architecture which, contrary to a conventional network, has a clock period equal to twice the delay of the slowest datapath stage (ST or LT) and allows the router datapath to propagate up to two flits in a single cycle as shown in figure 1.1. This dual data-rate (DDR) transmission of flits in the network ensures that the datapath can potentially be utilized for the complete duration of a cycle, serving a flit during both high and low clock phases. For a balanced router datapath, where the ST and LT delays are almost the same, this removes the slack from the router datapath and shifts it to the control logic as shown in figure 1.1. We use this slack and modify the SA of the router to allocate input and output ports of the switch

to up to two flits. This significantly improves the NoC saturation throughput as flits are now routed at a rate defined by the datapath while the flow control logic still has some available slack.

1.2.1.1 Related Work

Various existing NoC architectures attempt to either reduce or exploit the imbalance between the router datapath and control aiming to improve performance. Many use architectural techniques such as combined allocation or pipeline VA and SA to reduce allocation complexity and the critical path of the flow control logic [10, 21]. Others employ low level implementation techniques like time stealing (or retiming), currently supported by CAD tools, to balance the pipeline of a router and have a clock period equal to the average delay of all (3) router stages [29]. Nevertheless, for all these techniques, the control logic either ends up with a delay larger than the datapath delay or it adds to the datapath delay causing a single clock cycle to be sequentially shared by a part of both the control logic and the datapath. Both these cases result in a router critical path which includes the control logic, requires a clock with a period larger than the datapath delay and leads to NoC datapath under-utilization. The router architecture we present in this thesis allows the control logic neither to define nor share the clock period with the datapath.

Moreover, NoC research community has also proposed router architectures which employ DDR flit traversal in different parts of the NoC. Such NoC routers either use double-pumped crossbars [6] or DDR links [16, 17, 30]. These designs employ DDR over a part of the router's datapath to improve either area, power or performance but they do not achieve to eliminate the slack in the router datapath, resulting in its under-utilization. Contrary to these works, our approach increases network throughput by enabling DDR flit propagation over all parts of the router datapath to route flits at a rate determined entirely by the datapath delays (the switch and the link traversals) rather than by the control logic delays.

1.2.1.2 The Dual Data-Rate Network-on-Chip

In this thesis we describe the design of the DDRNoC which improves datapath utilization by receiving and sending flits at DDR. At the input port, DDRNoC routers require dual-edge triggered (DET) VC registers which can store incoming flits at both the falling and the rising clock edges. However, the available 28nm technology standard cell libraries, and for that matter most of the commercially available technology libraries, do not provide standard cells for DET flip-flops. This limitation is overcome by designing a specialized triggering mechanism for the VC registers which uses forwarded control signals to selectively trigger the destination VC register of an incoming flit at either positive or negative clock edge. This triggering mechanism exploits the fact that even if an input port receives two flits (belonging to same or different VC) in a cycle at DDR, they will always be stored in different VC buffer registers. Thus, if a VC buffer register has to register an incoming flit undergoing LT during the positive or negative clock phase it only needs to be triggered either at the falling or the rising clock edge, respectively, but never on both clock edges per cycle. Moreover, the same problem of registering flits on falling and rising clock edges

also exists at the output port, where up to two flits might need to be registered per cycle after the ST stage and before the LT stage. Here we use two registers, one negative and one positive edge triggered, to store flits which undergo ST during positive and negative clock phases, respectively.

Once the flits are registered in the VC buffers of the input port, the control logic of such a router should be capable of allocating each datapath resource to at most two flits per cycle. These datapath resources include downstream VCs for new packets and input/output ports of the local router switch. It is important to note that whereas the datapath of the DDRNoC (crossbar and link) can be used twice in a single cycle, the control logic (VA and SA) uses a complete cycle to generate the allocation decisions. We analyze the performance impact of various VC allocation schemes including a maximum matching VC allocator (based on augmenting path algorithm [26]) and observe minor performance improvements when compared to the increase in complexity of the respective allocators. So, the DDRNoC routers use a simple VA scheme based on separable allocators (similar to our baseline NoC) opting for energy efficiency and area without sacrificing much performance.

Once a downstream VC is allocated to a packet, its flits enter the ST stage and propagate from the input port to the output port of a router where they are registered. A switch capable of propagating two flits per input/output port requires an allocator which firstly resolves contention among competing switch allocation requests by distributing them over the two clock phases. If contending requests still exist then up to two requests per switch port should be granted based on some local priority while the remaining should be rejected. On the other hand, if there is only one packet requesting both input and output ports of the switch (i.e. no contention exists), then the SA should allow two flits of that requesting packet to undergo ST at DDR while ensuring at least two flits and two downstream VC credits will be available in the router at the time of ST (in next cycle). More importantly, the critical path of such an allocator should be less than twice the delay of the longest datapath stage (ST or LT). To this end we propose a novel speculative dual-grant output-first separable SA with round-robin input and output priorities per port which not only facilitates a contention-free DDR utilization of the switch, but also has a critical path that is less than twice the longest datapath stage (ST or LT). Moreover, the VA and SA logic operates on same allocation requests in parallel. Hence, the DDR SA speculatively grants allocation requests, which are also requesting a downstream VC in the same cycle, while giving higher priority to those requests which already possess a downstream VC from an earlier grant.

The next stages after SA is ST and LT. In order to implement the two SA decision corresponding to two clock phases, 2:1 multiplexers are added to the control of the router datapath (which includes VC buffers, input/output multiplexers of the crossbar and inter-router links) to reconfigure it during each clock phase. This allows the router switch and link to be sequentially shared by two flits in a single cycle.

Moreover, extra logic and link wires are also needed to decrement VC credit counters by two and convey two freed credits to the upstream router, respectively, in the case of DDR flit transmission. Further details of the DDRNoC are discussed in Chapter 2 (paper I) of this thesis.

1.2.2 Latency

Whereas the DDR traversal of flits over the NoC datapath significantly improves the throughput of the network, the latency of packets in the network also doubles. Although research shows that some applications gain more in performance from increase in the NoC throughput [31] while others are more sensitive to latency. In any case, this thesis supplements the DDR flit traversal in a router with low cost techniques which reduce zero-load latency of packets.

1.2.2.1 Related Work

On the packet latency front, some existing works use lookahead signaling to inform the downstream router of incoming flits one clock cycle earlier. This enables the downstream router to complete VA and SA for incoming flits before their arrival [10, 32, 33], effectively hiding the cycle required for these allocation stages per hop. Many other works also allow bypassing of allocation stages in the absence of contention usually under low network load [21, 33, 34]. Whereas these works only consider NoC architectures which route flits at SDR, we also analyze the application of these techniques in routers where both the control and the datapath allow flits to be routed at DDR. Moreover, we implement these techniques while ensuring that the critical path remains within the datapath of the router.

1.2.2.2 Control Forwarding

Similar to a simple SDR network [26], the flits in the DDRNoC also require three cycles per hop (VA/SA, ST, LT). However, longer clock period of the DDRNoC router compared to an SDR router results in a higher zero-load latency for the flits in DDR network. Consequently, under low network traffic, packets encounter high latency in the DDRNoC when compared to any SDR network. In order to overcome this limitation, lookahead signaling (called control forwarding) is implemented to transmit and register control signals at the downstream router one cycle before their corresponding flits and initiate VA/SA in the downstream router in advance. By the time the (up to) two incoming flits (per input port) are registered in the input VC buffers of the downstream router, their allocation decisions are already made which, if successful, allow the flits to initiate ST in the next cycle. This overlaps LT stage of the current router with VA/SA stage of the downstream router and hides the cycle a flit needs at each hop for allocation. For DDRNoC, control forwarding reduces packet zero-load latency by one cycle per hop. Control forwarding is described in more detail in Chapter 2 (paper I) of this thesis.

1.2.2.3 Pipeline Stage Bypassing

This thesis also presents the FreewayNoC network architecture (an improvement over the DDRNoC architecture) which uses simplified pipeline stage bypassing to further reduce zero-load packet latency. This technique targets end-to-end packet latency at low traffic loads when packets encounter reduced contention on their way to the destination. By allowing the incoming flits at a router input port to bypass the SA stage in the absence of contention and immediately

undergo ST and thereafter LT if downstream buffer resources are also available, the amount of time a flit spends in the network is further reduced.

Various checks are required before a flit is allowed to bypass SA stage to ensure conflict-free ST and LT. We reduce the complexity of the bypassing check logic and remove the need for arbitration among bypassing requests by only allowing flits to bypass SA if they are propagating straight through a router (i.e. N to S, E to W and vice versa). Although this restricts the benefits of pipeline stage bypassing, for most of the routing algorithms the number of turns a flit takes, under low network loads, during its journey are only two or three (to/from local port and an XY turn). In any case, pipeline bypassing requires bypassing checks before the control forwarding logic which include control switch traversal (CST) and control link traversal (CLT). These bypassing checks should happen in sequence and in the same cycle as CST and CLT. This control path can potentially become the critical path and extend the DDRNoC clock period, which would limit performance and goes against the key objectives of this thesis. So, three techniques are used to ensure that this path does not affect the clock cycle time of the NoC router, i) low latency bypass paths are provided for switch traversal of the bypassing control signals ii) most of the checks have very low complexity and delay and are parallelized and iii) for bypassing checks with longer delays, speculative CST of the bypassing control signals is implemented to provide more time for the check logic to complete while the bypassing signal is cancelled at the output port in case of mis-speculation.

Furthermore, to make pipeline bypassing of SA stage viable, free downstream VCs have to be allocated to bypassing head flits of new packets. FreewayNoC accomplishes this by assigning free downstream VCs with available credits to the head flits of new packets when they are registered at the output port of a router after ST stage and before the LT stage. SA only grants requests by head flits if there are available VCs with credits in their required output direction. Further details of the FreewayNoC architecture are presented in Chapter 3 (paper II) of this thesis.

1.2.3 Energy Efficiency

The breakdown of Dennard scaling in mid 2000's has made power and energy efficiency of on-chip components a vital design parameter. NoCs are one of the major sources of power consumption, consuming in many systems 10% to 25% of the total power budget [5–11]. In the absence of low power and energy efficient on-chip interconnects, on-chip cores may have to be either dimmed down or turned off to ensure fault free on-chip communication. Hence, the performance of many-core on-chip systems depend not only on the performance of NoCs but also on their energy efficiency and power consumption.

The DDR network architecture trades power for throughput, without improving packet latency, or trades latency for power savings without compromising throughput.

1.2.3.1 Related Work

Many existing techniques attempt to exploit network traffic fluctuations to conserve power or improve performance. For instance, power can be saved by

turning off network parts that are idle [35–37]. It is common to have unevenly loaded network parts as applications tend to generate traffic that fluctuates both in space and time [38–41]. To exemplify, channel utilization in PARSEC benchmarks varies from zero to 43% [38, 40].

Power gating has been applied in various granularities to save static power on idle resources. In the past, it has been suggested to power off virtual-channel buffers [35, 36], links [37] or even complete routers when they are not serving any traffic. Power gating usually negatively affects the latency and throughput of a network because powered off modules require multiple cycles (usually 6 to 12 clock cycles [42–45]) to turn on again and start serving traffic. Clock gating is another option for saving power on idle buffers, when they are not power gated. This prevents clock signal propagation to inactive registers and reduces clock power.

One of the most common approaches for adjusting network efficiency to its particular traffic is the use of Dynamic Voltage and Frequency Scaling (DVFS) [29]. Research shows that the clock distribution in the NoC consumes 20% of the total NoC power [46, 47]. On one hand, Voltage and Frequency can be reduced to save power at the cost of higher latency when the throughput requirements (and network injection rate) are low. On the other hand, in a power constraint design Voltage and Frequency could be increased to boost performance for a short period of time at the cost of higher power consumption. DVFS can be applied to an entire network leading to a simple design, which misses however some power saving opportunities as it cannot exploit spatial traffic fluctuations [48]. Whereas when applied separately to each router or groups of routers, it reduces power more effectively but costs additional latency and power for crossing different Voltage Frequency Islands (VFIs) [49].

1.2.3.2 Clock Tree Power Reduction

Since the DDRNoC operates at half the clock frequency required by the slowest datapath stage, the power consumption of the clock tree is also reduced. It further reduces clock distribution power by clock gating the datapath registers. These registers also include the input VC buffer registers which can be triggered to store incoming flits at either negative or positive clock edges. Automatic clock gating of these VC buffer registers is not supported by the available synthesis tools. In order to implement fine grained clock gating for each of the input VC registers, a low cost triggering mechanism is designed and implemented per VC register which only triggers the destination VC register of the incoming flit at the required clock edge while the remaining VC registers remain clock gated (discussed in more detail in section 1.2.1.2). Other registers in the datapath or the control logic are either positive or negative edge triggered and can be automatically clock gated by the synthesis tools. These optimizations serve to reduce dynamic power consumption of the clock tree and improve energy efficiency of the interconnect.

1.2.3.3 Low Voltage Implementation

We also perform low voltage physical implementation of the DDRNoC to analyze its performance and energy characteristics when powered by a 0.95V supply instead of 1.1V. We observe that throughput similar to the baseline

SDR NoC [26] operating at 1.1V can be obtained from a low voltage DDRNoC with significant power savings and improved energy efficiency. This provides an interesting design point for energy efficient DDR networks. This is discussed in more detail in Chapter 2 (paper I) of this thesis.

1.3 Contributions

In summary, the contributions of this thesis are the following

(i) design a NoC router for which the throughput is based on the routing rate defined by the longest datapath stage (ST or LT)

(ii) for all, but the turning hops, it has a latency per hop at best equal to the datapath delay ($2 \times \max(ST, LT) \approx ST + LT$)

(iii) improve the energy efficiency of the NoC by extracting higher performance while reducing the operating clock frequency

More precisely, we design SystemC simulators as well as RTL implementation of NoC routers which support DDR flit traversal, lookahead signaling and pipeline stage bypassing while targeting previously mentioned objectives. We evaluate NoC performance using the SystemC simulators while the RTL implementation is synthesized and placed-and-routed (P&R) using commercial 28nm process technology to obtain NoC area and operating frequency. Furthermore, power analysis is performed simulating post-P&R netlists with back-annotated delays to analyze the energy characteristics of the DDRNoC and the FreewayNoC.

We evaluate these network designs in the two papers briefly discussed below.

Paper I presents the first DDRNoC network architecture which uses a double-pumped datapath. It is based on the observation that conventional SDR 2D mesh interconnects with VC flow control have significant slack on their datapath stages while the control logic defines the critical path. DDRNoC uses this slack offering two flits to share the same datapath within a cycle at DDR. Both the allocation logic and the datapath are modified to facilitate the conflict free DDR transmission of flits. Furthermore, lookahead signaling is used to hide the cycle required for switch and VC allocation. Paper I also presents comparison of performance and energy characteristics of the DDRNoC with various SDR networks.

- While operating at the same voltage as the baseline, an 8×8 2D-mesh DDRNoC supports up to 27-45% higher throughput compared to SDR NoCs and is up to 27% better than a network with only DDR links.
- At low loads the average latency of packets in the DDRNoC is up to 50% higher than the SDR NoCs for synthetic traffic. This is mainly because of longer clock period of the DDRNoC and pipeline stage bypassing capabilities of some SDR networks.
- The DDRNoC reduces energy per transferred bit, energy delay product (EDP) and energy throughput ratio (ETR) due to a slower clock and higher network throughput.

- Furthermore, a low-voltage DDRNoC implementation reduces power consumption by up to 40%, at the cost of 26-40% higher latency, still however offering similar or better throughput when compared to a high-voltage implementation of a conventional SDR NoC. This low voltage implementation further improves energy efficiency, offering a substantially better energy-performance trade-off.

Paper II presents FreewayNoC, an enhancement of the DDRNoC, which not only routes packets at DDR but also allows simplified pipeline stage bypassing to reduce zero-load latency packets encounter in the DDRNoC.

- Paper II shows that at high traffic loads, FreewayNoC matches or improves the throughput of the DDRNoC. Hence, the throughput of a 16×16 2D-mesh FreewayNoC is about 25% better than the VC based SDR networks.
- FreewayNoC achieves a zero-load latency that scales to the number of hops better than current state-of-the-art NoCs and equally well with an ideal network that has no control overheads. This is because, for a balanced router datapath, per hop packet latency with FreewayNoC is at best equal to the sum of the ST and LT delays. This enables FreewayNoC to achieve lower packet latency than existing networks at higher number of hops per packet and larger network sizes.

1.4 Thesis Outline

The remainder of this thesis is organized as follows:

The Chapter 2 of this thesis presents Paper I. It provides the design and evaluation of the DDRNoC network architecture. The DDRNoC utilizes DDR flit traversal to improve network throughput and lookahead signalling to reduce packet latency.

In Chapter 3, of this thesis presents Paper II. It provides the design and evaluation of the FreewayNoC network architecture. FreewayNoC extends the DDRNoC architecture with simplified pipeline stage bypassing to reduce packet zero-load latency.

Bibliography

- [1] A. Psathakis, V. Papaefstathiou, N. Chrysos, F. Chaix, E. Vasilakis, D. Pnevmatikatos, and M. Katevenis, “A systematic evaluation of emerging mesh-like cmp nocs,” in *ANCS*, 2015, pp. 159–170.
- [2] P. Lotfi-Kamran, B. Grot, and B. Falsafi, “Noc-out: Microarchitecting a scale-out processor,” in *2012 45th Annual IEEE/ACM Int. Symp. on Microarchitecture*, Dec 2012, pp. 177–187.
- [3] H. Esmailzadeh, E. Blem, R. St. Amant, K. Sankaralingam, and D. Burger, “Dark silicon and the end of multicore scaling,” in *ISCA*, 2011, pp. 365–376.
- [4] S. Borkar, “How to stop interconnects from hindering the future of computing!” in *2013 Optical Interconnects Conf.*, May 2013, pp. 96–97.
- [5] M. B. Taylor, J. Kim, J. Miller, D. Wentzlaff, F. Ghodrati, B. Greenwald, H. Hoffman, P. Johnson, J.-W. Lee, W. Lee *et al.*, “The raw microprocessor: A computational fabric for software circuits and general-purpose programs,” *IEEE micro*, vol. 22, no. 2, pp. 25–35, 2002.
- [6] Y. Hoskote, S. Vangal, A. Singh, N. Borkar, and S. Borkar, “A 5-ghz mesh interconnect for a teraflops processor,” *IEEE Micro*, vol. 27, no. 5, pp. 51–61, 2007.
- [7] S. Bell, B. Edwards, J. Amann, R. Conlin, K. Joyce, V. Leung, J. MacKay, M. Reif, L. Bao, J. Brown *et al.*, “Tile64-processor: A 64-core soc with mesh interconnect,” in *IEEE Int. Solid-State Circuits Conf.*, 2008, pp. 88–598.
- [8] P. Salihundam, S. Jain, T. Jacob, S. Kumar, V. Erraguntla, Y. Hoskote, S. Vangal, G. Ruhl, and N. Borkar, “A 2 tb/s 6 x 4 mesh network for a single-chip cloud computer with dvfs in 45 nm cmos,” *IEEE J. of Solid-State Circuits*, vol. 46, no. 4, pp. 757–766, April 2011.
- [9] J. Howard, S. Dighe, S. R. Vangal, G. Ruhl, N. Borkar, S. Jain, V. Erraguntla, M. Konow, M. Riepen, M. Gries, G. Droege, T. Lund-Larsen, S. Steibl, S. Borkar, V. K. De, and R. V. D. Wijngaart, “A 48-core ia-32 processor in 45 nm cmos using on-die message-passing and dvfs for performance and power scaling,” *IEEE Journal of Solid-State Circuits*, vol. 46, no. 1, pp. 173–183, Jan 2011.
- [10] B. K. Daya, C.-H. O. Chen, S. Subramanian, W.-C. Kwon, S. Park, T. Krishna, J. Holt, A. P. Chandrakasan, and L.-S. Peh, “Scorpio: A 36-core research chip demonstrating snoopy coherence on a scalable mesh noc with in-network ordering,” in *Int. Symp. on Computer Architecture*, ser. ISCA '14, 2014, pp. 25–36.
- [11] B. Bohnenstiehl, A. Stillmaker, J. J. Pimentel, T. Andreas, B. Liu, A. T. Tran, E. Adeagbo, and B. M. Baas, “Kilocore: A 32-nm 1000-processor computational array,” *IEEE Journal of Solid-State Circuits*, vol. 52, no. 4, pp. 891–902, 2017.

- [12] S. Rao, S. Jeloka, R. Das, D. Blaauw, R. Dreslinski, and T. Mudge, "Vix: Virtual input crossbar for efficient switch allocation," in *Design Automation Conf. (DAC)*, 2014, pp. 103:1–103:6.
- [13] J. Kim, J. Balfour, and W. Dally, "Flattened butterfly topology for on-chip networks," in *MICRO-40*, 2007, pp. 172–182.
- [14] J. Balfour and W. J. Dally, "Design tradeoffs for tiled cmp on-chip networks," in *Int. Conf. on Supercomputing*, 2006, pp. 187–198.
- [15] F. Gilabert, M. Gómez, S. Medardoni, and D. Bertozzi, "Improved utilization of noc channel bandwidth by switch replication for cost-effective multi-processor systems-on-chip," in *NOCS*, 2010.
- [16] A. Psarras, S. Moisisdis, C. Nicopoulos, and G. Dimitrakopoulos, "Rapidlink: A network-on-chip architecture with double-data-rate links," in *2016 IEEE International Conference on Electronics, Circuits and Systems, ICECS*, 2016, pp. 93–96.
- [17] —, "Networks-on-chip with double-data-rate links," *IEEE Trans. on Circuits and Systems*, vol. 64, no. 12, pp. 3103–3114, 2017.
- [18] B. Grot, J. Hestness, S. W. Keckler, and O. Mutlu, "Express cube topologies for on-chip interconnects," in *HPCA*, Feb 2009, pp. 163–174.
- [19] C. H. O. Chen, N. Agarwal, T. Krishna, K. H. Koo, L. S. Peh, and K. C. Saraswat, "Physical vs. virtual express topologies with low-swing links for future many-core nocs," in *NOCS*, 2010, pp. 173–180.
- [20] M. Hayenga and M. Lipasti, "The NoX router," in *MICRO-44*, 2011.
- [21] C. N. Anastasios Psarras, Ioannis Seitaniadis and G. Dimitrakopoulos., "ShortPath: A Network-on-Chip Router with Fine-Grained Pipeline Bypassing," *IEEE Trans. on Computers*, vol. 65, no. 10, pp. 3136–3147, 2016.
- [22] J. Kim, D. Park, T. Theocharides, N. Vijaykrishnan, and C. R. Das, "A low latency router supporting adaptivity for on-chip interconnects," in *Design Automation Conf. (DAC)*, 2005, pp. 559–564.
- [23] H. Farrokhbakht, M. Taram, B. Khaleghi, and S. Hessabi, "Toot: an efficient and scalable power-gating method for noc routers," in *2016 Tenth IEEE/ACM International Symposium on Networks-on-Chip (NOCS)*, 2016, pp. 1–8.
- [24] L. Chen, D. Zhu, M. Pedram, and T. M. Pinkston, "Power punch: Towards non-blocking power-gating of noc routers," in *2015 IEEE 21st International Symposium on High Performance Computer Architecture (HPCA)*, 2015, pp. 378–389.
- [25] S. Park, M. Qazi, L. S. Peh, and A. P. Chandrakasan, "40.4fj/bit/mm low-swing on-chip signaling with self-resetting logic repeaters embedded within a mesh noc in 45nm soi cmos," in *2013 Design, Automation Test in Europe Conference Exhibition (DATE)*, 2013, pp. 1637–1642.

- [26] W. Dally and B. Towles, *Principles and Practices of Interconnection Networks*. Morgan Kaufmann Publishers, 2004.
- [27] C. Nicopoulos, V. Narayanan, and C. R. Das, *Network-on-Chip Architectures - A Holistic Design Exploration*, ser. Lecture Notes in Elect. Eng. Springer, 2010, vol. 45.
- [28] H. Kim, A. Vitkovskiy, P. V. Gratz, and V. Soteriou, "Use it or lose it: Wear-out and lifetime in future chip multiprocessors," in *MICRO-46*, 2013, pp. 136–147.
- [29] A. K. Mishra, R. Das, S. Eachempati, R. Iyer, N. Vijaykrishnan, and C. R. Das, "A case for dynamic frequency tuning in on-chip networks," in *MICRO-42*, 2009, pp. 292–303.
- [30] J. Xu, W. Wolf, and W. Zhang, "Double-data-rate, wave-pipelined interconnect for asynchronous nocs," *IEEE Micro*, vol. 29, no. 3, pp. 20–30, May 2009.
- [31] A. Bakhoda, J. Kim, and T. M. Aamodt, "Throughput-effective on-chip networks for manycore accelerators," in *MICRO-43*, 2010, pp. 421–432.
- [32] A. Kumary *et al.*, "A 4.6tbits/s 3.6ghz single-cycle noc router with a novel switch allocator in 65nm cmos," in *2007 25th International Conference on Computer Design*, 2007, pp. 63–70.
- [33] S. Park, T. Krishna, C. H. Chen, B. Daya, A. Chandrakasan, and L. S. Peh, "Approaching the theoretical limits of a mesh noc with a 16-node chip prototype in 45nm soi," in *DAC Design Automation Conference 2012*, June 2012, pp. 398–405.
- [34] T. Krishna *et al.*, "Swift: A swing-reduced interconnect for a token-based network-on-chip in 90nm cmos," in *IEEE ICCD*, 2010, pp. 439–446.
- [35] H. Matsutani, M. Koibuchi, D. Ikebuchi, K. Usami, H. Nakamura, and H. Amano, "Ultra fine-grained run-time power gating of on-chip routers for cmps," in *NOCS*, 2010, pp. 61–68.
- [36] G. Kim, J. Kim, and S. Yoo, "Flexibuffer: Reducing leakage power in on-chip network routers," in *Design Automation Conf. (DAC)*, 2011, pp. 936–941.
- [37] G. Michelogiannakis and J. Shalf, "Variable-width datapath for on-chip network static power reduction," in *NOCS*, 2014, pp. 96–103.
- [38] N. Barrow-Williams, C. Fensch, and S. Moore, "A communication characterisation of splash-2 and parsec," in *IEEE Int. Symp. on Workload Characterization (IISWC)*, 2009, pp. 86–97.
- [39] L. Chen and T. M. Pinkston, "Nord: Node-router decoupling for effective power-gating of on-chip routers," in *MICRO-45*, 2012, pp. 270–281.
- [40] R. Hesse, J. Nicholls, and N. E. Jerger, "Fine-grained bandwidth adaptivity in networks-on-chip using bidirectional channels," in *IEEE/ACM NOCS*, May 2012, pp. 132–141.

- [41] V. Soteriou and L. s. Peh, "Exploring the design space of self-regulating power-aware on/off interconnection networks," *IEEE Trans. on Par. and Distr. Syst.*, vol. 18, no. 3, pp. 393–408, 2007.
- [42] L. Chen and T. M. Pinkston, "Nord: Node-router decoupling for effective power-gating of on-chip routers," in *MICRO-45*, 2012, pp. 270–281.
- [43] R. Das, S. Narayanasamy, S. K. Satpathy, and R. G. Dreslinski, "Catnap: energy proportional multiple network-on-chip," in *ACM SIGARCH Comp. Arch. News*, vol. 41, no. 3, 2013, pp. 320–331.
- [44] H. Matsutani, M. Koibuchi, D. Wang, and H. Amano, "Run-time power gating of on-chip routers using look-ahead routing," in *Asia & South Pacific Design Aut. Conf.*, 2008, pp. 55–60.
- [45] A. Samih, R. Wang, A. Krishna, C. Maciocco, C. Tai, and Y. Solihin, "Energy-efficient interconnect via router parking," in *HPCA*, 2013, pp. 508–519.
- [46] M. A. Anders, "High-performance energy-efficient noc fabrics: Evolution and future challenges," in *IEEE/ACM NOCS*, Sept 2014.
- [47] S. Vangal, A. Singh, J. Howard, S. Dighe, N. Borkar, and A. Alvandpour, "A 5.1ghz 0.34mm² router for network-on-chip applications," in *IEEE Symp. on VLSI Circuits*, June 2007, pp. 42–43.
- [48] P. Salihundam, S. Jain, T. Jacob, S. Kumar, V. Erraguntla, Y. Hoskote, S. Vangal, G. Ruhl, and N. Borkar, "A 2 tb/s 6 4 mesh network for a single-chip cloud computer with dvfs in 45 nm cmos," *IEEE J. of Solid-State Circuits*, vol. 46, no. 4, pp. 757–766, 2011.
- [49] M. K. Yadav, M. R. Casu, and M. Zamboni, "Laura-noc: Local automatic rate adjustment in network-on-chips with a simple dvfs," *IEEE Trans. on Circuits and Syst. II: Express Briefs*, vol. 60, no. 10, pp. 647–651, 2013.