RingNet: A Memory-Oriented Network-on-Chip Designed for FPGA

Jakub Siast, Adam Łuczak, Marek Domański, *Member, IEEE*

***Abstract*—In this paper, we identify the general requirements for NoCs and the general characteristics of FPGAs from leading producers. Based on the analysis provided, an FPGA-oriented NoC called RingNet is proposed. As a distinctive feature, RingNet uses communication through a centrally placed memory that aims at preventing network congestions and limiting the network buffer requirements. Optimal utilization of FPGA resources is one of the goals of RingNet development. Especially, buffers are implemented in distributed RAM available in FPGAs, and the virtual cut-through is used as an efficient switching technique for FPGA. Simulations prove guaranteed throughput, predictable latency, and fair network access provided by RingNet. Synthesis results for sample FPGAs from Xilinx, Intel, and Lattice prove the universality of RingNet. The provided analysis of NoC implementations leads to the conclusion that RingNet needs fewer resources and supports higher clock frequencies than the widely used AXI4 architecture.**

***Index Terms*—FPGA, Network-on-chip, distributed memory (LUTRAM), virtual cut-through, fairness.**

I. Introduction

NETWORK-ON-CHIP (NoC) is a widely adopted solution for the interconnect problem in large Systems-on-Chip (SoCs). Hitherto, most of research on NoCs has been related to ASICs [1] but the rapidly growing FPGA size yields also growing interest in NoCs implemented in complex FPGAs. Using NoCs as interconnections for FPGAs is not a new idea [2], [3], [4]. Although FPGAs differ from their ASIC counterparts, most of known FPGA NoCs were adopted from ASICs without considering FPGA-specific features. Therefore, the potential of NoCs has not been fully exploited for FPGAs. In this situation, older interconnecting techniques like crossbars (e.g., AXI4 Interconnect) are still in use, regardless of their poor scalability [3]. The development of new NoC architectures, better suited to FPGA, is still a challenging problem.

Manuscript received ……..; revised ……; accepted ……. . Date of publication ……; date of current version ……. .

This work was supported by the Ministry of Science and Higher Education as DS project at Chair of Multimedia Telecom. Microelectronics, Poznań University of Technology.

Authors would like to thank Synopsys Inc. for providing Synplify synthesizer for experimentation.

This paper has supplementary downloadable material available at http://ieeexplore.ieee.org. This includes Appendices I-III to Section VII, and Appendix IV to Section VI. This material is 8 MB in size.

J. Siast and M. Domański are, and A. Łuczak was with Chair of Multimedia Telecommunications and Microelectronics, Poznań University of Technology, Poznań, Poland (email: {jsiast, aluczak}@multimedia.edu.pl, marek.domanski@put.poznan.pl)

In this article, we propose a novel NoC architecture called RingNet that is well-suited to the features of contemporary FPGAs. Among other NoC architectures proposed for FPGAs, RingNet stands out with communication through a central memory and traffic load controlled by the recipient. The paper starts with a discussion on the common features of FPGAs that are important for NoC architectures (Section II). In Section III, general requirements for NoCs are discussed. Next, in Section IV we summarize the state-of-the-art in NoC designs for FPGA and point out NoC design constraints that match the FPGA features considered in Section III. In Section V we propose the idea of a novel RingNet NoC with its protocol description provided in Section VI. Through Sections VII-IX we present experimental results, compare RingNet synthesis for different FPGAs, and compare RingNet implementation with the widely-used crossbar interconnection called AXI4 Interconnect.

II. Characteristics of FPGAs

In the paper, the basic common features of modern FPGAs are studied in the context of products offered by three of the leading FPGA vendors: Xilinx Inc., FPGA department of Intel (formerly Altera), and Lattice Semiconductor Corp. We limit our considerations to devices large enough to contain a SoC. The size of an FPGA is measured in logic cells (LCs) that are equivalent to a 4-input look up table (LUT) paired with a flip-flop (FF). A memory controller, a common module of SoC, requires several thousands of LCs (e.g., DDR3 SDRAM memory controller for Artix7 requires more than 6500 LCs). The whole SoC is expected to be substantially larger, therefore we limit the considerations to a series of FPGAs with devices of more than 50,000 LCs.

The FPGA considerations are based on data sheets from the vendors [5]–[18] and summed up in Table I. The above-mentioned three manufacturers offer products similar in various aspects, all based on LUTs and FFs, and with the capacity of up to millions of LCs. One can also see that memory controllers for high-capacity SDRAM are supported as software IPs or hardware pre-engineered blocks. Each FPGA contains memory blocks (BRAM, with capacity from 9 kb to 45 Mb) distributed across its array.

Each of the considered devices also includes distributed RAM (LUTRAM). LUTRAMs utilize LUTs with limited hardware added for supporting on-the-fly LUT reprogramming. In the considered FPGAs, from 11% to 50% of LUTs can be used as RAM. The smallest available LUTRAM configurations have the depth of 16 or 32 words, depending on the producer. Different numbers of LUTRAM bits are available per utilized LUT; on average, 32 bits per an LUT for Intel devices, 48 for

TABLE I
COMPARISON OF FPGAS

Vendor	Device name	Year of announcement	Technology [nm]	Embedded ARM processor	Number of logic cells	FPGA basic cell feature	Supported external memory controller	Size of a block memory (BRAM)	The smallest LUT-based distributed memory (LUTRAM) configurations with separate read and write port	Portion [%] of LUTs usable as RAM
Intel (Altera)	Stratix 10	'13	14	yes	378k — 5.5M	6-input LUT + 2 FF [18]	hard DDR4-2666	20kb & 45Mb	20×LUT as 20b×32-word deep or 10b×64-word deep RAM (32 bits per LUT)	25
	Arria 10	'13	20		160k — 1.15M			20kb		27 — 50
	Cyclone 10	'17	20	no	85k — 220k		hard DDR3-1866			21
	Stratix V	'10	28		236k — 952k		soft DDR3-1600			50
	Arria V	'11	28	yes	75k — 504k		hard DDR3-1066	20kb & 10kb		25 — 50
	Cyclone V	'11	28		25k — 301k		hard DDR3-800	10kb		24 — 35
Xilinx	Zynq UltraScale+	'15	16	yes	83k — 914k	6-input LUT + 2 FF	soft DDR4-2666	36kb or 2×18kb & 288kb	4×LUT as 3b×64-word deep or 6b×32-word deep RAM (48 bits per LUT)	30 — 50
	Virtex UltraScale+	'15		no	690k — 2.9M					
	Kintex UltraScale+	'15			205k — 915k					
	Virtex UltraScale	'14	20		783k — 5.5M		soft DDR4-2400	36kb or 2×18kb		14 — 22
	Kintex UltraScale	'14			318k — 1.5M					14 — 47
	Virtex7	'10	28		326k — 2M		soft DDR3-1866			28 — 43
	Kintex7	'10		yes (in Zynq)	66k — 478k					
	Artix7	'10			13k — 215k		soft DDR3-1066			
	Spartan7	'16		no	6k — 102k		soft DDR3-800			
	Spartan6	'09	45		4k — 147k		hard DDR3-800	18kb or 2×9kb		24 — 50
Lattice	ECP5	'14	40	no	12k — 84k	4 input LUT + 1/0 FF	soft DDR3-800	18kb	6×LUT as 4b×16-word deep RAM (10.7 bits per LUT)	50
	LatticeECP3	'09	65		17k — 149k					11 — 15
	LatticeECP2	'06	90		6k — 95k		soft DDR2-533			

Xilinx and 10.7 for devices from Lattice. As pointed out by a source related to Xilinx [3], RAM-capable LUTs are well spread over an FPGA and their potential should be considered in NoC and SoC designs for Xilinx FPGAs. All the considered FPGAs include LUTRAM, so this conclusion should be extended to all of them.

FPGAs, unlike ASICs, have a predefined network of programmable signal pathways connecting LUTs. Each LUT in FPGA implements a logic function with a limited number of logic inputs. For Intel and Xilinx products, functions of up to 6 inputs can be implemented in a single LUT. For Lattice products, the number of inputs is limited to 4. As the number of inputs for required logic exceeds the number of single LUT inputs, it needs to be realized as multiple layers of LUTs connected with pathways. Additional layers of LUTs reduce the maximum clock frequency. In order to obtain the required design frequency, special care needs to be taken not to exceed the critical number of layers. It is not the case for ASIC designs, where the maximum clock frequency can be balanced with flexible lengths of pathways.

It can be concluded that the considered FPGAs differ from their ASICs counterparts. The key advantages of the considered FPGAs are highly available distributed RAM and support for high-capacity SDRAM, whereas the discussed frequency limitation is the main FPGA constraint. The identified advantages and constraints should be taken into account in the development of NoCs for FPGA.

III. REQUIREMENTS FOR NOC AND RELATED PREVIOUS WORKS

An NoC consists of switches connected with links. Processing elements (PEs) are connected to NoC using network interfaces. Switches can realize two switching techniques: packet-switching and circuit-switching. The paper considers only packet-switched networks that are characterized by high link utilization and are widely used for FPGA [3], [19]. In packet-switched NoC, a single message is divided into packets, and sent across the network. The packets are further divided into flow control units called flits. A flit is a portion of data usually transferred at one clock cycle between connected switches [20].

A number of requirements for NoCs have already been identified [4], [21], [22]. The obvious requirements are:

a) Utilization of resources (power, silicon area, LUTs and FFs) should be minimized.

b) High data throughput should be offered.

c) Latency should be limited.

There are some other important requirements that are addressed in few references only:

d) Fairness of network access should be guaranteed, i.e., all network interfaces should experience throughput proportional to their relative request rates and the same latency [22].

e) Network should be reliable, i.e., it should be deadlock-free, whereas the requirements of the minimum throughput and the maximum latency should be met [21]. Another kind of reliability is fault tolerance required by some applications [23], but fault tolerance is not considered in this article.

Network reliability can be affected by congestions [20]. Congestions lead to throughput and latency fluctuations, thus

the average throughput is below the theoretical maximum value, and the average latency is increased, especially under high network load conditions.

According to [1], [4], [24], [25], and [49], the main aspects influencing throughput, latency and probability of congestions are the following: switching technique, topology together with the routing algorithm, and the size of buffers. For each of those aspects, a number of techniques have been developed to meet the requirements for NoCs:

a) The switching technique. Two frequently used techniques are the wormhole and the virtual cut-through [20]. The wormhole is congestion-sensitive and may result in low network utilization, but can provide low latency, and requires smaller buffers than the virtual cut-through. The latter requires larger buffers but provides low latency without limiting network utilization due to congestions.

b) The topology and the routing algorithm define paths in the network and influence the loads of individual links and switches. Uneven loads result in bottlenecks in the network that may cause congestions and unfair network access, leading to throughput and latency fluctuations [19].

Topologies and routing algorithms that prevent congestions have already been investigated in the literature. In [19], the congestions are limited by spreading traffic across NoC evenly by using adaptive routing. Another approach is to use multiple physical link/network (MP) topologies [21], [27]. Redundant physical links increase throughput, reduce bottlenecks and prevent congestions.

c) The size of buffers has been shown to have a major impact on throughput, latency and occurrence of congestions [24].

Determining the buffer size is not trivial in the case of *a priori* unknown traffic load generated by processing elements (PEs) and unknown ability of PEs to accept packets from the network buffers. In the references, several mechanisms were proposed to determine the buffer size. The simplest method is to set the size of buffers to hold as many packets as can be generated. In the case of *a priori* unknown traffic load, this worst-case approach results in unnecessarily large buffers [28].

More advanced methods of determining the buffer size exploit the statistics of the traffic load. Those statistics need to be explicitly provided [48], or a dedicated traffic load monitoring technique needs to be used [24], [25], [26], [29].

In [24], the buffer size is adjusted iteratively in consecutive SoC implementations. The monitoring module collects traffic statistics that are used to adjust the buffer size accordingly, and the estimated size is used in the next implementation. Multiple SoC implementations are time-consuming, therefore, in [25] simplified PEs are emulated and traffic statistics are collected faster. In [26] an NoC simulator is proposed to estimate traffic statistics prior to NoC implementation. The above-mentioned mechanisms [24]–[26] need training data and are sensitive to any change in the traffic pattern.

In [29], the total size of memory in a switch is constant, but based on the measured traffic load at each switch output, the memory is assigned between output buffers adaptively, during runtime. Nevertheless, even in [29], it is pointed out that this mechanism is not dedicated for FPGA due to its complexity.

Commonly, networks distinguish the types of transmitted data and assign separate buffers to these different types. This technique is called virtual channels (VCs) [27], [30], [31], [49]. In NoCs using VCs, it is common to guarantee throughput and latency just for critical types of data, like control messages. Therefore, the aforementioned buffer size determining techniques are only applied to buffers used by the critical data types. Buffers used by non-critical data can be optimized to reduce the memory cost.

In the NoC proposed in this paper, the buffer size does not depend on the traffic load. It is the opposite, and the traffic load is controlled to utilize the fixed-size network buffers without causing congestions. Details are provided in Section V.

IV. State-of-the-art in Interconnections for FPGA

A number of NoCs, crossbars and bus architectures are proposed for ASICs in the references, and some of them are adopted for FPGAs. For the sake of brevity, we focus only on FPGA-oriented crossbars and NoCs.

The AMBA AXI4 [32] is a version of the AMBA crossbar recommended by ARM Ltd for FPGAs. Originally, AXI4 was used to communicate with ARM cores embedded in many devices (see Table I). AXI4 modules, especially a switch called AXI4 Interconnect, are available as Intellectual Property (IP) cores in the Xilinx Vivado Design Suite and Intel Quartus. Many other IP cores with AXI4 interface are available. Therefore, an FPGA-based SoC using AXI4 Interconnect and AXI4-compatible Processing Elements can be developed rapidly. For this reason, AXI4 is widely used in FPGAs [3], [4]. Despite its popularity, AXI4 Interconnect has drawbacks, e.g., in [3] long connections were pointed out as the main drawback, and the application of NoC instead of AXI4 Interconnect was suggested for Xilinx FPGAs.

Two types of NoCs for FPGA have already been proposed in the literature: soft, with infrastructure implemented using general FPGA resources, and hard, using hardware switches embedded into the FPGA fabric as an additional resource [33]–[35]. Hard NoCs may limit the flexibility of FPGA and have not been implemented in chips that are available, and therefore only soft NoCs will be considered in this paper.

A few soft NoC architectures were proposed recently for FPGAs: Hoplite [36]–[41], RAR-NoC [19], CONNECT [4], [31], [42], [3] and LinkBlaze [43].

The above-mentioned works use different approaches to the requirements for NoC (cf. Section III). All the works consider the need to limit the usage of FPGA resources. Most of them dismiss adaptive routing techniques, and exploit static routing, willing for a smaller control logic [42], [4], [36], [3]. Only in RAR-NoC [19] is the usage of adaptive routing considered instead of static, and the increased throughput is reported at the cost of the switch size increase by 11%.

Buffers can consume a significant part of the network resources, therefore the authors of Hoplite [36] propose an extreme approach and implement a toroid NoC without buffers. Recently, Hoplite-based NoCs were proposed with the aim of lowering Hoplite average latency by improved routing [37], [38], or changing the topology from toroid to butterfly fat tree

[41]. The small size of Hoplite NoC is paid for by no guaranties for packet latency and lack of reliability. Using no buffer may not be justified for FPGAs that provide easily available distributed RAM (LUTRAM).

In [19] and [42], it is proposed that the buffer size be limited by implementing wormhole switching that requires smaller buffers than virtual cut-through. Xilinx-related authors pointed out [3] that LUTRAMs in Xilinx FPGAs have some generally defined minimum depths, so most of the LUTRAM capacity may be wasted by wormhole switching with shallow buffers. Virtual cut-through is proved [3] to be an appropriate switching technique for a system with moderate size packets, i.e., shorter than the depth of LUTRAM-based buffers.

In the paper describing the CONNECT switch [31], the LUTRAM potential is underlined, and the buffers are implemented strictly using LUTRAM. The configurability of the depth and width of the CONNECT switch is an advantage of the implementation, but it is also pointed out that even fixed but appropriate buffer depth can improve the predictability of resource utilization.

Based on [3] and [31], one can conclude that using buffers with depth equal to the depth of LUTRAM, using packets of length shorter than the buffer depth and employing virtual cut-through switching can result in resource-efficient NoC. Therefore, we implement those ideas in RingNet NoC.

The maximum operating frequency and the average throughput and latency are reported for all the above-mentioned networks in their source papers. The respective reports are provided for individual FPGA device types [36], [19], [3], [42], or for entire FPGA device lines [4]. The lack of reports for FPGAs from different vendors may be an obstacle in determining the usability of NoCs.

Another NoC designed especially for FPGA is LinkBlaze [43] that is dedicated for UltraScale+ devices from Xilinx. In particular, the switching logic was optimized for 6-input LUTs, and the switches were placed manually in the array with the aim of connecting them with global pathways. This device-aware network design results in a high-frequency NoC with its throughput higher than obtained for Hoplite and CONNECT when implemented on UltraScale+ FPGA. LinkBlaze is an example of NoC optimized for one line of FPGAs, whereas in this article we look for more universal NoC.

A few of the presented NoC proposals focus on the requirement of NoC reliability. RAR-NoC [19] uses a traffic monitor to control the routing algorithm implemented with the aim of reducing congestions. In [4] and [3], switches that support VCs are implemented. Still, even though fairness is an essential reliability parameter, it is out of the scope of most propositions. Only in [4] did the authors point out the importance of fairness, but gave no numerical results for their NoC. In this paper, we will provide a fairness analysis for RingNet, together with resource utilization and maximum operating frequency for various FPGAs to prove its usability.

Out of the above-mentioned NoCs, only the authors of LinkBlaze [43] consider FPGAs' support for SDRAM. SDRAM is a crucial component of many SoC projects [33] and utilizes a substantial part of NoC throughput. In [20] an example of an ASIC SoC for an AVC decoder is presented. In the system, the total communication traffic to/from SDRAM is much higher than the one required for communication between other processing elements (PEs). This type of memory-oriented SoCs are the target application for RingNet described in the next section.

V. RingNet Architecture Description

Most NoC proposals focus on switch architecture only. In this paper, we propose a complete NoC architecture and a protocol for a memory-oriented SoC which we call RingNet. In the development of RingNet, the conclusions from Sections II-IV are taken into account.

A. Primary Ideas of the Proposal

1) As described in Section III, determining the buffer size is often complex due to the *a priori* unknown traffic load. In RingNet, the traffic load is controlled by a destination processing element (PE), which guarantees that packets injected into the network can be accepted by the destination PE without congestions. This way, fixed-size network buffers can be utilized.

With the aim of providing the traffic control mechanism, we disallow direct communication between PEs. In RingNet, all traffic goes through one of the system buffers: System Memory (e.g., external SDRAM with a memory controller implemented in FPGA) or the Reflector. The System Memory is used as a data buffer, whereas the Reflector is a dedicated network buffer for control messages. The Reflector is introduced because it serves functions that are not supported by an ordinary memory controller. Among its functions, the Reflector informs a PE about data waiting to be read from a buffer. For two PEs to communicate, the first PE writes to a system buffer and the second PE reads from it to complete the communication. What is important, in RingNet, a PE requests data from system buffers when it is ready to accept it. Still, sending data to the system buffers is not a matter of any restriction and a sending PE can work at its own pace. The Reflector and the memory controller are the only devices in RingNet that need to be prepared for traffic with *a priori* unknown pattern.

RingNet has two distinctive properties among other NoCs proposed for FPGA. These are traffic load controlled by a destination Processing Element, and communication through system buffers.

2) Considering the conclusions from Section IV, we propose virtual cut-through switching for efficient utilization of LUTRAM. In Section IV, it was also concluded that packets should be small enough to fit into the available LUTRAM. For the considered FPGAs, it means packets should be shorter than 32 flits (see Table I).

3) The usage of logic functions with a low number of inputs can give a high clock frequency implementation in FPGA (cf. Section II). In [4], it is shown that a low number of ports results in a high frequency switch. In fact, a 3-port switch is the smallest switch usable in a network (see also [43], and [33]). In order to maximize the frequency of the network, we use 3-port switches, called the Leaf Interface (LI) and Root Interface (RI).

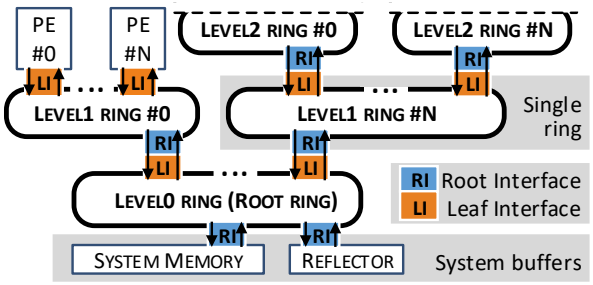

Fig. 1. RingNet topology.

B. Secondary Ideas of the Proposal

4) A tree-of-rings topology is used (see Fig. 1) with the System Memory and the Reflector connected to the ring at the root of the network tree, whereas PEs are connected to the rings at higher levels of the tree. The ring topology is recommended for FPGAs by several authors [33], [31], and it is one of just few topologies that can be constructed using 3-port switches. Moreover, the ring topology allows the network to spread over the whole FPGA area.

On the other hand, in NoCs with ring topology, latency increases proportionally to the number of PEs, and for a high number of PEs, this latency may be unacceptably high. Such latency can be reduced with the use of a mixed topology of smaller rings connected to a tree.

Both the tree and ring topologies are easy to scale in FPGAs without reducing the maximum clock frequency [31]. For the tree-of-rings topology there is one path between the root and a leaf device, so static routing may be efficiently used, which simplifies the RingNet logic.

5) The choice of topology affects the parameters of RingNet, especially it limits network throughput. The throughput of a ring is a function of the flit width and the clock frequency. The width of a flit in RingNet is constant, so the maximum ring throughput is determined by the maximum clock frequency for a given FPGA. To overcome this limitation, the multiple physical network (MP) technique is used. A single ring at any level of the tree can be replaced with multiple rings connected in parallel. The traffic is spread evenly between the parallel rings, and the throughput is multiplied. The traffic in RingNet aggregates in a root ring and this level can also be multiplied to meet the throughput of the attached SDRAM memory.

The usage of the MP technique makes the throughput of RingNet controllable. In this situation, the System Memory throughput becomes an obvious limitation of our approach. Nevertheless, the System Memory load can be reduced by connecting additional memory buffers at any ring. PEs attached to a common ring can exchange data using the additional buffers. Each such buffer is connected to a ring through its own RI and can utilize the block RAMs available in all considered FPGAs.

6) Flits in RingNet have 8 data bytes with additional 8 bits of byte-enable (BEN). Like in AXI4, BEN bits indicate which data bytes from a flit are valid. One of the goals of RingNet development is to provide throughput equal to or higher than the throughput of SDRAMs supported by FPGAs. While using flits with 8 bytes of data, the throughput of RingNet is compared with the throughput of the supported SDRAMs:

- The slowest SDRAM from Table I is DDR2-533, which provides throughput of 4.3 Gbps for an 8-bit interface, whereas a single RingNet ring running at a moderate frequency of 75MHz already exhibits the throughput of 7 Gbps.
- The most demanding DDR4-2666, supported in Xilinx UltraScale+ series and Intel 10 series, offers 192 Gbps for a 72-bit interface. For comparison, five parallel RingNet rings can transfer 233 Gbps when running at the clock frequency of 500 MHz that is easily achievable for FPGAs.

7) Considering the requirements for fair network access (Section III), we propose a flow control mechanism for RingNet with the access controlled locally, at the level of each ring (cf. Section VI).

VI. Elements and Protocol of the RingNet Network

This section provides an overview of the implementation of example RingNet components and protocol.

In this NoC, there are two physical channels. The first one transports packets form a processing element (PE) to a system buffer (System Memory or Reflector). The second one transports packets from a system buffer to a PE. PEs are connected at the leaves of the network tree, whereas the system buffers are connected at the root of the tree, therefore, the first channel is called Leaf-to-Root (L2R) and the second one is called Root-to-Leaf (R2L).

The main elements of RingNet are two types of network interfaces: Leaf Interface (LI) and Root Interface (RI). The LI is used to connect a PE to a network. The RI is used to connect a system buffer. As depicted in Fig. 1, a combination of RI and LI is used to connect rings at different levels of a network tree. Both RI and LI are 3-port switches and they insert packets to a ring using a 2×1 multiplexer and accept packets from a ring using 1×2 demultiplexers. Detailed descriptions of the LI and RI are provided in supplementary materials as Appendix IV.

Each ring in RingNet has one L2R channel and one R2L channel (see Fig. 2). A flow control mechanism for the L2R

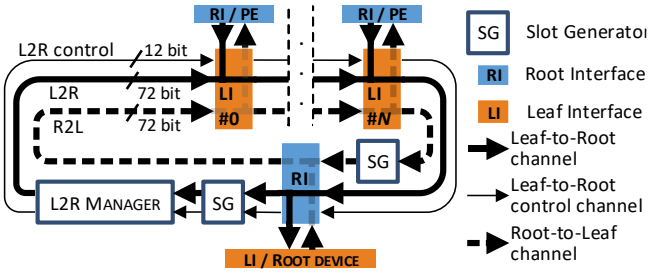

Fig. 2. A single RingNet ring.

channel uses an additional L2R control channel.

As already mentioned, the rings can be connected in parallel in order to increase the overall network throughput, e.g., at certain critical levels of the tree. The multiplied rings can be connected to other levels of the tree using a dedicated adapter. Fig. 3 depicts an adapter for the L2R channel. The adapter is placed between LIs and RIs of the rings when at least one ring is multiplied. An individual adapter is used for L2R and R2L

channels. It is built of 2×1 multiplexers and 1×2 demultiplexers and LUTRAM-based 32-flit deep buffers.

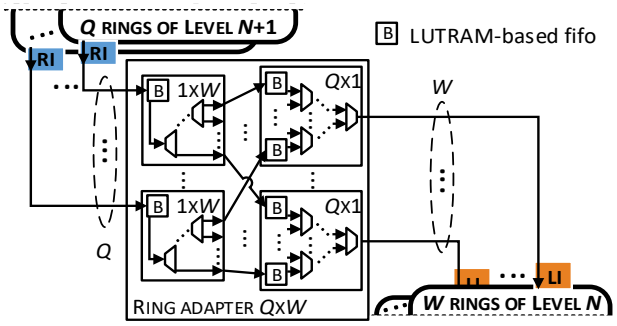

Fig. 3. Ring adapter for Leaf-to-Root (L2R) channel.

For RingNet, a protocol is developed that provides communication between PEs, limits congestions, provides fair network access and supports packet priorities. The protocol covers 4 layers of the OSI model, from the data link layer to the session layer. Higher OSI layers are out of scope of the RingNet protocol.

A. Data Link Layer

The data link layer specifies the structure of flits and packets, and defines the ring access protocol.

As recommended in Section V, packets should be shorter than the depth of the buffers. For the FPGAs from Table I, the most shallow LUTRAM is 32-word deep. Therefore, 32 and 64-flit deep buffers are used in RingNet. We choose packets with two possible lengths of 2 and 9 flits. The purpose for the two lengths is described in Subsection C. The first flit of a packet is the header that encapsulates control information for each protocol layer (especially the routing information). The following 1 or 8 flits transport data. In RingNet flits, 64 bits are transmitted with additional 8 bits of byte enable (BEN).

Fig. 4 depicts a pattern of time slots for header flits and data flits circulating through interfaces around a ring. Those flit slots are organized into long and short slots for long and short packets, respectively. Slots are produced by the Slot Generator (SG) depicted in Fig. 2. The Leaf Interface (LI) and the Root Interface (RI) can populate those slots with packets.

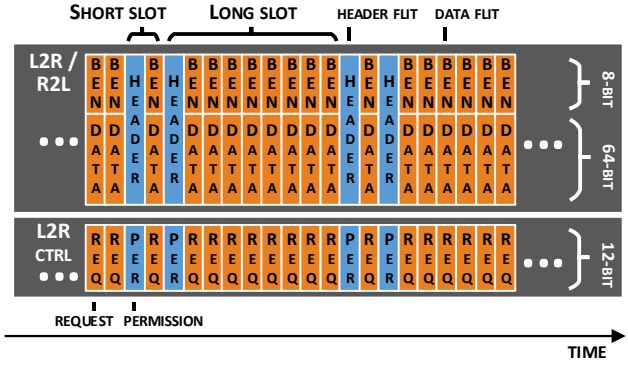

Fig. 4. Time slots at R2L, L2R and L2R control channels, passing through an interface.

The access to the Leaf-to-Root (L2R) channel slots is controlled by the L2R Manager. In order to access this channel, LI sends a request to the L2R Manager and obtains permission. The requests and permissions are sent using the L2R control channel. The L2R Manager keeps the requests in LUTRAM-based buffers, and grants the permissions based on their order of arrival. This strategy is sufficient to guarantee fair access to the L2R channel for all PEs in terms of the average latency and granted throughput (see simulation results in Section VII).

The flow control mechanism used for the L2R channel also guarantees fair access to the Root-to-Leaf (R2L) channel. Therefore, no flow control for the R2L channel is needed, and no additional resources are used. Details will be given in Subsection C.

RingNet supports 4 packet priorities, ordered from the highest priority 3 to the lowest priority 0. Each packet and the associated request has an assigned priority. The L2R Manager has individual buffers for requests of different priorities.

The LIs buffer packets, send requests, and insert buffered packets onto the L2R ring after acquiring permission. Sending a request and obtaining permission requires some additional time, so a constant flow of packets from one LI may be impossible. To overcome this problem, multiple packets can be buffered in LI and multiple requests can be sent without receiving permission. This overlap lets a single LI exploit the full ring throughput. The buffer used in LIs utilizes LUTRAM with depth of 64 words. It provides enough capacity for 6 short packets and 5 long packets.

According to the virtual cut-through switching, the packets from the L2R channel leave a ring through RI only if the RI has buffer space available for a whole packet. Otherwise, the packet is rejected by the RI and starts to cycle around the ring until enough space is available. A circulating packet is recognized by the L2R Manager, and no new permission is granted until all packets that circulate on L2R leave the ring. This way, all packets should finally leave the ring with limited differences in latency. Detailed description of the L2R Manager is provided in supplementary materials as Appendix IV.

B. Network Link Layer

The network layer protocol defines RingNet addressing. Different addressing is used for the L2R and the R2L channels. For the L2R channel, 37-bit memory addressing is used, therefore up to 128 GB can be addressed in the network; both the Reflector and System Memory have assigned memory address spaces recognized by Root Interfaces (RIs). For the R2L channel, the packet address contains five (each 4-bit wide) numbers. Each number identifies the Leaf Interface (LI) that should accept the packet at a certain RingNet tree level. The addressing used in the R2L channel limits the number of network tree levels to 5 and the number of network interfaces connected to a single ring to 15. The applied addressing scheme limits the number of PEs that can be connected to the RingNet network to 759,375.

C. Transport Layer

The transport layer defines logical channels and describes the communication between PE and system buffers.

In RingNet, packets are transported to and from memory-mapped system buffers, therefore the RingNet

protocol supports basic memory operations. The memory read or write operation is encoded in a packet header, thus creating a read or write packet. If a packet is sent from PE via the L2R channel to a system buffer, a response packet has to be sent back to the PE by the system buffer through the R2L channel. For the read operation, the response packet contains the data read from the memory, whereas for the write operation the response packet is a write operation acknowledgement.

Logical channels are defined for RingNet. In each logical channel, a certain memory operation and packets of a certain length are used. Most transfers are realized by two logical channels:

a) Logical write channel. PE sends a long write packet through the L2R physical channel, with data to be stored to the memory. A short response packet is sent back by a system buffer through R2L as a write acknowledgement.

b) Read logical channel. PE sends a short read packet through the L2R physical channel, and a long packet with data read from the memory is sent back from a system buffer through the R2L physical channel.

Other types of logical channels are used by the session layer. What is very important is that for read and write logical channels a single response packet is sent via R2L in response to the L2R packet. Therefore, the same number of packets are transferred through L2R and R2L channels. This way, fair access to the L2R channel guarantees fair access to the R2L channel as well, despite the fact that for R2L no requests-permissions mechanism is used explicitly. The proposed flow control prevents congestions in both physical channels, still utilizing resources in the L2R channel only.

Theoretical throughput T_{RW_MAX} of the logical read and write channels, expressed in bits per clock tick (bpt) can be calculated according to the given formula:

$$T_{RW\_MAX} = R \cdot 64 \cdot 8 / 11 \text{ [bpt]}, \quad (1)$$

where 64 is the number of data bits in a flit, the 8/11 factor is the share of flits carrying write data at the L2R physical channel, and the share of flits carrying read data at the R2L physical channel, and R is the number of parallel rings used at the root ring.

D. Session Layer

The transport layer describes the communication between PEs and system buffers, but not between individual PEs. The communication between PEs is finally possible at the session layer. An example of communication between two PEs using the System Memory and the Reflector is depicted in Fig. 5. The example illustrates two transactions of sending data from PE1 to PE2.

PE1 starts the transaction by sending data to the System Memory. Then, PE1 informs PE2 that data in the System Memory is ready to be processed by sending an *event* message through the Reflector. An event is a short packet with an address in a range reserved for the Reflector. The events are buffered in the Reflector and sent to the destination PE one at a time. After

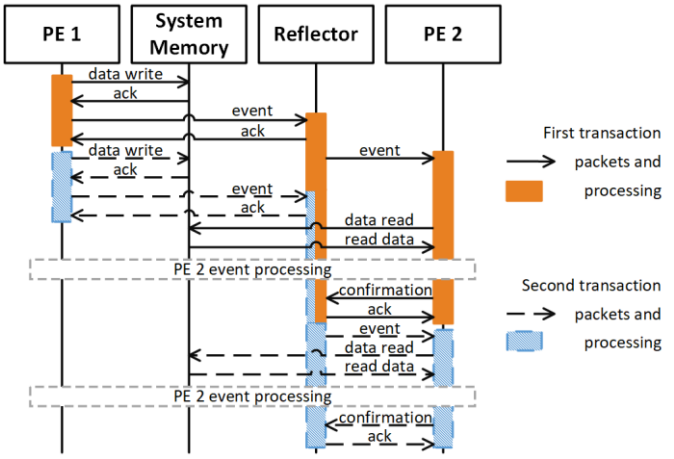

Fig. 5. Sequence diagram for two transactions between PEs.

processing data, PE2 sends the event confirmation to the Reflector which ends the transaction.

The example demonstrates two advantages of the proposed indirect PE communication. First, congestions arising due to PE overload are prevented, because the destination PE (PE2 in the example) controls its load. Second, the source PE (PE1 in the example) does not need to monitor the state of the destination PE before starting a new transaction.

The reflector has short dedicated first-in-first-out buffers (fifos) for each PE, and a general buffer for events that do not fit into the dedicated fifos. When space in the dedicated fifo is released, an event from the general buffer is transferred to the dedicated fifo.

The events transfer control information and use packets with the highest priority (3). Moreover, a dedicated virtual channel (VC) is created by reserving part of the Leaf Interface (LI) buffer just for the event packets.

The Reflector provides additional system functions, like informing about PEs connected to NoC, registering new PEs, informing about events buffered in the Reflector, alarming about the fullness of event fifos, resetting a PE, etc.

VII. Simulation Results

The proposed RingNet is simulated with the aim of testing its throughput and latency. Moreover, the following reliability aspects are tested: inter-channel dependencies, network access fairness and the priority mechanism.

The simulations use a model of the System Memory that supports unlimited throughput and introduces negligible latency, thanks to which the System Memory does not affect network performance results. There is no need to simulate various traffic patterns with various packet destinations [4], [49] because the System Memory is the destination of all the data in RingNet.

PEs are simulated with the use of packet generators (PGs). Each PG independently generates packets for the logical read and write channels with the configurable average delay D_{aver} between two consecutive packets:

$$D_{aver} = \frac{N \cdot 64 \cdot 8}{T_{RW\_MAX} \cdot L \cdot 11}, \quad (2)$$

where T_{RW_MAX} is the theoretical throughput of a network, L is the requested aggregated load generated by all PGs expressed as a percentage of the throughput T_{RW_MAX}. Eq. (2) is explained because 64 is the number of data bits in a flit, the 8/11 factor is the share of flits carrying write data at the L2R physical channel and the share of flits carrying read data at the R2L physical channel. The actual delay D is defined as the time interval that a PG waits before it sends another packet. D is an output of a random number generator with discrete uniform distribution $\mathcal{U}\{0.8D_{aver}, 1.2D_{aver}\}$ used for making the traffic more realistic.

In the experiments, RingNet with two levels of rings is simulated as depicted in Fig. 6. The size of the network is controlled using the following parameters:

- R: Multiplication degree of the root level, i.e., the number of parallel rings used at the root level (level 0).
- F: The number of 1st level rings.
- G: The number of PGs connected to a single 1st level ring.

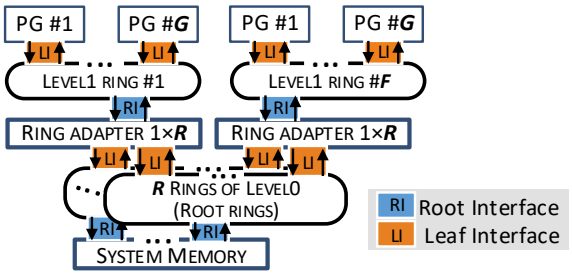

Fig. 6. Simulated RingNet topology.

Latency is measured in terms of the clock cycles that elapse between the emission of a new packet from a PG, and the acceptance of the corresponding response packet (send by the System Memory) by the PG.

The Verilator compiler [44] is used to convert Verilog code into C++ code that is further compiled by the GCC compiler and executed to perform functional simulations. The presented results summarize the simulations of over 3000 different configurations of network size and load.

A. Performance Test

In this test, the achievable throughput of RingNet is checked together with the average latency that can be expected under different conditions. The theoretical value of throughput for read and write logical channels can be calculated according to (1).

In the experiment, all packets are sent between PGs and the System Memory. Only packets with the lowest priority (priority 0) are generated. The parameters of the test are: the multiplication degree R of a root level is restricted to the interval 1 – 4, the number F of 1st level rings is from the interval 1 – 5, while the number G of PGs connected to a single 1st level ring is set in the range of 1 – 15 (up to 75 PEs are connected for F=5 and G=15). The aggregated load generated by all PGs is set in the range of 0% ÷ 100% of the theoretical throughput T_{RW_MAX} (1). The logical channel load is separately set for the read and write channels.

In Table II, the average latency for the read channel is presented for a high load (92% to 97%) and various network sizes (R=1, F=1÷5 × G=1÷15). One can see that latency increases with the increased number of connected 1st level rings and PGs. For the write channel, the measured latencies are, on average, 7 clock cycles longer than those reported in Table II. The results for the write channel and for the increased number of parallel rings used at the root level (R=2÷4) are provided in the supplementary material as Appendix I. It is tested that increasing the root ring multiplication degree by one increases the average latency reported in Table II by only 6 clock cycles.

TABLE II
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR READ CHANNEL FOR VARIOUS NETWORK CONFIGURATIONS

Number of 1st level rings (F)	Number of PGs connected to a 1st level ring (G)					
	1	2	3	4	7	15
5	141	163	169	175	201	258
4	130	151	157	163	189	245
3	129	148	155	161	185	241
2	113	134	142	145	182	225
1	95	118	120	122	147	195

In Fig. 7, the latencies from Table II are presented as a function of the number of connected PGs. The fitted latency function flattens for growing number of PG, i.e. for large networks, the latency increases slowly with the growth of the network.

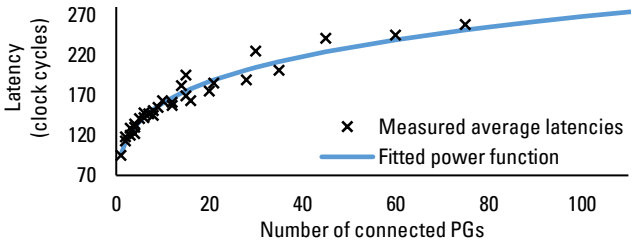

Fig. 7. Average latency (expressed in clock cycles) for read channel for various network sizes.

NoCs are often characterized by providing load-latency curves that represent packet latency as a function of network load [1], [4], [20], [22], [27], [28], [30], [36], [45]. Fig. 8 depicts load-latency curves for the logical channels of the RingNet network with five 1st level rings and 15 PGs connected at each ring (F=5 × G=15, 75 PEs connected) and 4 parallel rings used at the network root (R=4). The results for other network

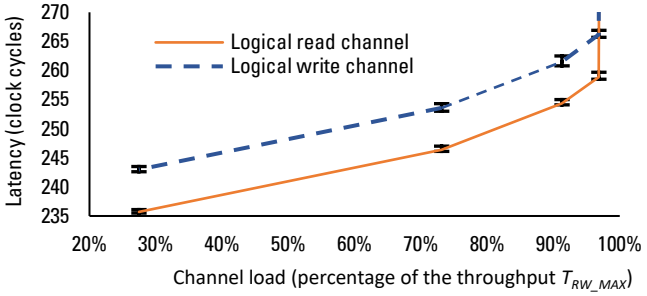

Fig. 8. Load-latency curves. These curves represent the average latency of each logical channel as a function of the channel load. Error bars represent the minimum and maximum average latencies of the channel at a given load when the second channel load changes in the range of 0% to 100%.

configurations are provided in the supplementary material as Appendix II.

The average latency for both logical channels increases with the channel load. For a 100% channel load, all network buffers are filled, and the average latency increases drastically, which seems to be obvious. On the other hand, an increase of the channel load from 27% to 97% results in the average latency increase by only 10%. Moreover, both logical channels are independent, i.e., the load of one logical channel has a negligible impact on the average latency in the other logical channel, and does not influence its throughput.

For the presented network configuration (and for all other tested configurations), throughput T_{RW_MAX} calculated according to (1) is obtained for both the read and write channels, regardless of the multiplication degree of the root ring (R), the number of 1st level rings (F), and the number of connected PGs (G).

The performance of RingNet is compared with the performance of the state-of-the-art networks. For RingNet, an average latency of about 90 clock cycles is observed for a network with just one PE and increases to about 250 clock cycles when 75 PEs are connected. Negligible changes in latency are observed for traffic loads set in the range of 27% to 98%. For a mesh NoC with 64 PEs described in [3], an average latency of 60 clock cycles is reported. Next, a latency of about 10 clock cycles is reported for CONNECT [4] mesh with 16 PEs. The relatively high latency observed in the RingNet network is a result of two factors. First, the RingNet ring topology features relatively long routes when compared with a mesh. Next, latency is caused by the applied flow control mechanism that exchanges flow control messages before a packet can be sent. Therefore, RingNet is not recommended for designs requiring very low latency. Nevertheless, the applied flow control mechanism is demonstrated to provide a fair network access not reported for the state-of-the-art NoCs [3] and [4].

In order to compare the throughputs of RingNet and the NoCs proposed in [3] and [4], we consider a traffic scenario where all PEs send packets to just one destination, i.e., a single memory device connected to the network. In such a simplified case, the throughput of the network is at most equal to the throughput of a network interface through which the memory is connected. The throughputs of a network interface for NoCs from [3], [4], and for RingNet are proportional to the flit width and frequency of a clock. In the Section VIII, it will be demonstrated that RingNet features substantially higher maximum frequency and uses less resources for the same flit width. It means that in the traffic scenario where all PEs send packets to just one destination, RingNet provides higher throughput than the state-of-the-art NoCs from [3] and [4].

The most important conclusion drawn from the test is that the theoretical throughput T_{RW_MAX} (1) is guaranteed for RingNet, and the latency is correlated with network size, and it can be estimated during network configuration.

B. Network Access Fairness Test

In test *A*, the average parameters over all processing elements (PEs) are estimated. In real network applications, fair access to NoC for each PE can be an even more important aspect than the average parameters. In order to check the fairness, the statistics need to be gathered for each individual PE.

In the test, a RingNet network with 75 PEs is simulated (F=5 × G=15), and four parallel rings are used at the root level (R=4). PEs are simulated using Packet Generators (PGs) requesting the aggregated load of 27% to 100% of the network throughput T_{RW_MAX} (1). In this experiment, all PGs send packets to the System Memory with priority 0. For each individual PG, we collect the values of the average latency L_{PG} expressed in clock cycles, and the values of throughput T_{PG} expressed in bits per clock cycle, for both logical channels.

In Table III, the average values of L_{PG} and T_{PG}, calculated over all PGs ($\overline{L_{PG}}$ and $\overline{T_{PG}}$, respectively), are presented together with a standard deviation for those variables ($\sigma_{\overline{L_{PG}}}$ and $\sigma_{\overline{T_{PG}}}$, respectively). Low values of the standard deviations mean that all PGs are expected to exhibit the same performance expressed in terms of the average latency and throughput. These conclusions are valid for a wide range of loads. The results prove that RingNet provides fair access for each connected PE. More results that prove the fair access for other network sizes are available in the supplementary material as Appendix III.

TABLE III
TRAFFIC STATISTICS OVER 75 PGS

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	236	6	243	6	0.7	0.01	0.68	0.01
97%	259	5	267	5	2.4	0.01	2.4	0.01
100%	1157	7	1175	9	2.5	0.00	2.48	0.00

C. Latency Distribution Test

Now, the distribution of packet latency in RingNet is researched.

A network with 5 1st level rings and 15 PGs connected at each ring (F=5 × G=15) is simulated under the loads of 27% to 97%. Four parallel rings are used at the root level (R=4).

Fig. 9 depicts a histogram of packet latency for the read channel. The data are collected for a sample PE (PG connected to a 1st level ring as a 4th PE). For the write channel and for other

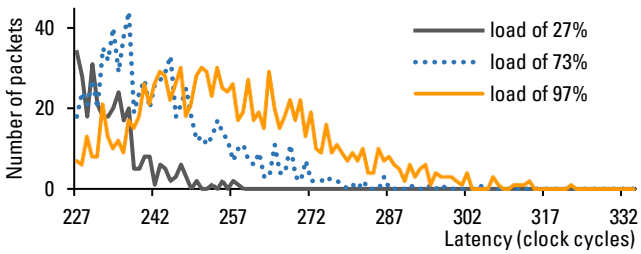

Fig. 9. Histogram of packet latency for read channel for an example of PE.

PGs, the histograms look similar, which is reasonable in view of the fact that fair network access was proved in test *B*.
The values of 227 and 322 are the minimum and maximum latency observed, respectively. For low loads, more packets experience latency close to the minimum. An increase of the load results in shifting the histogram towards higher latencies. Still, the difference between the extreme values is just 95 clock periods, which is 42% of the minimum latency.

D. Packet Prioritization Test

The proposed RingNet architecture supports packet prioritization. Higher priority packets obtain access to the Leaf-to-Root (L2R) channels first. Under high load conditions, the high priority packets should experience lower latency; also, the requested throughput should be granted starting from the highest priority requests. To check the prioritization mechanism, a test is conducted under high load conditions:

- Network with 28 PGs is tested (F=4 × G=7).
- Two parallel rings are used at the root level (R=2).
- Each PG has 3 internal sources of packets with priorities 0, 1, and 3. The source of packets with priority 0 constantly tries to send a packet, and on its own it would generate the load of 100% of the theoretical throughput T_{RW_MAX} (1). The sources of packets with priority 1 and 3 are set to request various loads.

In Fig. 10, examples of load shares are presented for sources of packets with priority 1 requesting 20% of the throughput, and sources of packets with priority 3 requesting from 0% to 100% of the throughput. For the highest priority packets, the share of throughput is always granted as requested. 20% of the throughput, requested for priority 1 packets is granted when the packets with the highest priority use less than 80% of the RingNet throughput. The example demonstrates that the lower priority packets do not influence the share of throughput granted to the higher priority packets. It needs to be emphasized that using the prioritization mechanism does not limit the aggregated throughput below the theoretical value T_{RW_MAX} (1) in any tested case.

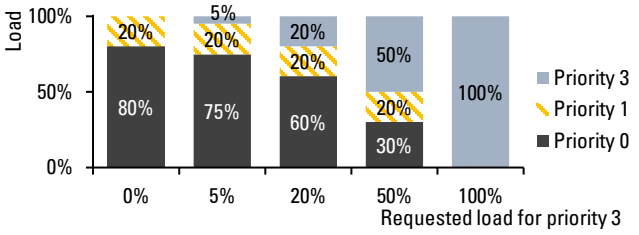

Fig. 10. Load, expressed as percentage of the theoretical throughput T_{RW_MAX} (1), granted to each priority under constant requested load for priorities 0 and 1 and increasing requested load for priority 3.

In Fig. 11, latency histograms are depicted for packets with different priorities. The presented statistics are collected for the case when packets with priority 3 and 0 utilize 5% and 95% of the network throughput, respectively. The histograms for both priorities are concentrated. The high priority packets experience moderate latencies even under the maximum network loads.
The proposed communication protocol with prioritization and flow control mechanisms was tested in a number of simulations. The results prove that the RingNet network throughput can be calculated according to (1) and the average latency can be estimated based on NoC size. Moreover, the proven fair access to RingNet network is a distinctive property among other NoCs proposed for FPGA.

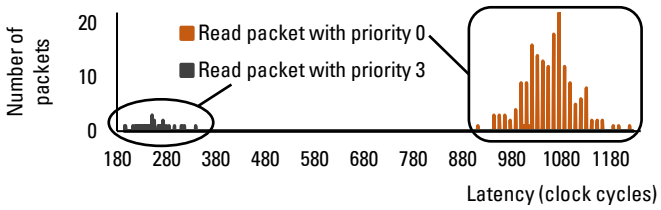

Fig. 11. Latency histogram for read channel.

In addition to the simulations, RingNet was tested in a hardware application. It was successfully used as a communication backbone for an FPGA-based depth map estimation device [46] demonstrating the above-mentioned features.

VIII. RingNet Synthesis

This section summarizes RingNet synthesis for various FPGAs and tests RingNet scalability.

RingNet is synthesized for sample FPGAs from Xilinx (Artix7, Kintex UltraScale, Virtex7), Intel (Stratix V, Arria V) and Lattice (ECP5). We synthesize individual RingNet modules and rings of different sizes that are useful for networks with a tree topology, i.e., rings with one Root Interface (RI) and from 2 to 15 Leaf Interfaces (LIs). For fair comparison, the same synthesis software – Synplify Premier 2017.03-SP1 – is used for all devices. The synthesis tools available in Synplify, like the shift register inferring and the register and logic replication [47], can balance the resource cost and clock frequency. The tools perform differently for different devices. With the aim of obtaining fair results for all devices, we disable the shift register inferring tool. Also, the register and logic replication tools are effectively disabled by setting the requested frequency to a low value of 1 MHz.

Our comparison is summarized in Table IV and reports the maximum clock frequency and the utilized resources.

The obtained frequency values are compared with the maximum frequencies specified by vendors for hardware blocks of each FPGA [5]–[15]. The frequencies for the common DSP and BRAM blocks are presented. The maximum frequencies are given for the fastest, fully pipelined configuration of each hardware block. For BRAM, two versions are given: one with simultaneous read and write to the same address handled in additional logic (RW check), and another one without this extension (no RW check). The provided frequencies are a raw estimation of what maximum frequency can be expected for a typical, high performance project implementation for each device. From Table IV, it can be seen that the maximum frequencies obtained for RingNet modules are comparable with or higher than those given for DSPs and BRAMs. Still, the results are generated for the required clock frequency set to 1MHz. If the required frequency is tuned, the Synplify synthesizer can optimize modules using techniques like register replication, and higher maximum frequencies can be obtained, gaining several dozens of MHz on average. The

TABLE IV
RESOURCES UTILIZATION AND ESTIMATED MAXIMUM FREQUENCY AFTER SYNTHESIS FOR RINGNET MODULES AND RINGS

RingNet modules, RingNet rings, and FPGA hardware blocks	FFs utilized / LUTs utilized as logic + LUTs utilized as RAM / Maximum clock frequency in MHz					
	Xilinx			Intel		Lattice
	Artix7 xc7a100tcsg324-1	Kintex UltraScale xcku060-ffva1156-3-e	Virtex7 xc7vx550tffg1158-1	Stratix5 5SGXMABK2H40C3	Arria5 5AGXBA7D6F31C6	ECP5 lfe5u_85f-8
Ring adapter 1×2	383 / 93 + 48 / 485	383 / 94 + 48 / 837	383 / 94 + 48 / 469	383 / 105 + 144 / 551	383 / 102 + 144 / 377	383 / 189 + 216 / 312
Ring adapter 2×1	462 / 184 + 96 / 459	462 / 184 + 96 / 837	462 / 184 + 96 / 486	462 / 123 + 288 / 529	462 / 127 + 288 / 345	462 / 351 + 432 / 312
L2R Manager	499 / 271 + 96 / 337	499 / 264 + 96 / 875	499 / 286 + 96 / 461	499 / 317 + 128 / 527	507 / 349 + 128 / 334	499 / 373 + 384 / 374
Slot Generator	106 / 11 + 56 / 487	106 / 11 + 56 / 837	106 / 11 + 56 / 487	106 / 14 + 168 / 983	106 / 12 + 168 / 468	106 / 16 + 126 / 372
Leaf Interface	759 / 227 + 104 / 405	759 / 218 + 104 / 807	759 / 227 + 104 / 509	759 / 145 + 166 / 524	761 / 160 + 166 / 293	761 / 391 + 450 / 235
Root Interface	648 / 125 + 48 / 469	648 / 125 + 48 / 837	648 / 125 + 48 / 483	648 / 56 + 144 / 779	648 / 57 + 144 / 381	648 / 204 + 216 / 314
Ring 2×1	2k8 / 805 + 424 / 354	2k8 / 801 + 424 / 753	2k8 / 805 + 424 / 445	2k8 / 631 + 916 / 449	2k8 / 648 + 914 / 273	2k8 / 1k3 + 1k4 / 235
Ring 4×1	4k3 / 1k2 + 632 / 363	4k3 / 1k2 + 632 / 700	4k3 / 1k2 + 632 / 460	4k3 / 934 + 1k2 / 430	4k3 / 966 + 1k2 / 250	4k3 / 2k0 + 2k3 / 235
Ring 6×1	5k8 / 1k7 + 840 / 333	5k8 / 1k7 + 840 / 763	5k8 / 1k7 + 840 / 440	5k8 / 1k3 + 1k6 / 409	5k8 / 1k3 + 1k6 / 262	5k9 / 2k8 + 3k3 / 235
Ring 8×1	7k3 / 2k1 + 1k0 / 328	7k3 / 2k1 + 1k0 / 726	7k3 / 2k1 + 1k0 / 440	7k3 / 1k6 + 1k9 / 397	7k3 / 1k7 + 1k9 / 272	7k4 / 3k6 + 4k2 / 235
Ring 11×1	9k5 / 2k8 + 1k4 / 333	9k5 / 2k8 + 1k4 / 753	9k5 / 2k8 + 1k4 / 446	9k5 / 2k1 + 2k4 / 434	9k6 / 2k2 + 2k4 / 241	9k7 / 4k9 + 5k8 / 235
Ring 15×1	12k5/ 3k7 + 1k8 / 325	12k5/ 3k7 + 1k8 / 763	12k5/ 3k7 + 1k8 / 443	12k5/ 2k7 + 3k1 / 403	12k5/ 2k8 + 3k0 / 222	12k8/ 6k5 + 7k6 / 235
Switch from [3]		--- / 1678 / 470				
Switch from [4]	--- / 430 / 241 (results for Virtex-6 LX760 speed grade -2)					
DSP	/ 392	/ 687	/ 463	/ 400	/ 200	/ 185
BRAM no RW check	/ 388	/ 660	/ 458	/ 650	/ 285	/ 272
BRAM RW check	/ 339	/ 575	/ 400	/ 455	/ 240	/ 214

obtained maximum clock frequencies are estimations, and actual frequency values should be generated by implementation tools provided by a device vendor like Vivado from Xilinx (see the next section).

Resource utilization for basic RingNet modules and rings is reported in Table IV. The utilization of flip-flops (FFs) and look up tables (LUTs) is presented.

The utilization of FFs in different devices is almost the same with the maximum deviation from the average module size of 2%.

In Table IV, two types of utilized LUTs are distinguished: LUTs utilized for logic implementation and LUTs utilized as RAM. The number of LUTs utilized as logic is 1.8 times higher in the Lattice device than in the devices from Xilinx and Intel, on average. The reason for the disproportion is because Xilinx and Intel use 6-input LUTs in their products and Lattice uses smaller, 4-input LUTs. Among Xilinx and Intel devices, the maximum deviation in the number of LUTs utilized as logic is 19% among all synthesized RingNet rings. Differences in LUT utilization between Xilinx and Intel FPGAs are the result of minor differences in the architectures used by each vendor, such as different multiplexing hardware at inputs and outputs of the LUTs.

The number of LUTs utilized as RAM differs significantly between vendors. It is a result of different numbers of LUTRAM bits available per utilized LUT; on average, 32 bits per LUT for Intel devices, 48 for Xilinx and 10.7 for devices from Lattice (see Section I). Hence, the expected ratio between the LUTs utilized as RAM in Xilinx, Intel, and Lattice FPGAs is 1 : 1.5 : 4.5. The obtained average ratio for RingNet rings is 1 : 1.8 : 4, so it is close to the expected ratio.

The average ratio between the number of utilized FFs and the number of utilized LUTs is equal to 2.2 : 1 for RingNet rings synthesized for Intel and Xilinx devices and 1 : 1 for Lattice. Those ratios are close to the ratios of FFs and LUTs available in the devices (see Table I) that we find desirable for FPGA design.

For most NoCs, a significant part of FPGA resources is utilized for buffers. It is also true for RingNet, where buffers are implemented using LUTRAM, whereas RAM-capable LUTs are a limited resource (see Table I). As for any limited resource, using less LUTRAM can make routing of synthesized design easier and provide shorter links and higher clock frequency. We checked what share of the LUTs utilized for RingNet rings is utilized for RAM. This share happens to be constant for a given FPGA vendor: about 33% for rings synthesized for FPGAs from Xilinx, 55% for Intel and 54% for ECP5 from Lattice. The share of utilized RAM-capable LUTs exceeds the share of available RAM-capable LUTs, but only by 4% and 5% in the case of Lattice and Intel, respectively. This shows reasonable LUTRAM utilization, suitable for FPGA implementation.

Through the presented synthesis results, it is shown that RingNet modules can be efficiently implemented in devices of various types from different vendors, with comparable resource utilization and a high value of the maximum clock frequency.

In Table IV, the results for the state-of-the-art 64-bit flits switches are presented for reference [3], [4]. The 5-port switch proposed in [3] is reported to utilize 1678 LUTs per PE when implemented in the Xilinx Kintex UltraScale device. The CONNECT network using 3-port switches and configured in ring topology [4] is reported to utilize 430 LUTs per PE when implemented in Xilinx Virtex-6 LX760. On the other hand, RingNet utilizes only 367 LUTs per PE if 15×1 ring is used. It is clear that 3-port switches used in RingNet and CONNECT utilize fewer LUTs than the 5-port switch from [3]. In RingNet ring, unlike in [3] and [4], some network elements, i.e., Slot Generators and L2R Manager are common to the number of connected PEs. The cost of the elements is shared among the number of the connected PEs, which decreases the resources utilized per PE. It is the reason for fewer LUTs utilized per PE

in RingNet when compared with CONNECT. It needs to be pointed out that the shared L2R Manager used in RingNet not only lowers resource utilization but also increases latency, which is discussed in Section VII. On the other hand, the L2R Manager has knowledge about the required load, therefore it can provide a fair access for all its associated PEs.

When implemented in the devices of the same speed class, the RingNet rings can be clocked with higher clock frequency than switches proposed in [3], as well as the CONNECT network. This feature is achieved by the use of small and optimized 3-port switches in RingNet. For example, for the Kintex UltraScale Xilinx devices, the maximum clock frequency is above 700 MHz and about 470 MHz for RingNet and the network of switches [3], respectively (cf. Table IV). Another example shows that for a Virtex6 device, the maximum frequency for CONNECT network is about 31% – 59% of the maximum DSP frequency, whereas it is 84% to 107% of the maximum DSP frequency for a RingNet ring.

The RingNet network utilizes less resources than the NoCs from [3], and [4], but it comes at the expense of increased latency (see Section VII). On the other hand, the high latency can be mitigated with high maximum clock frequency of RingNet.

IX. Comparison between RingNet Ring and AXI4 Interconnect

We compare RingNet with AXI4, a widely used communication infrastructure for FPGAs. Both have common features. According to [32], AXI4 is designed for high-performance memory-mapped requirements, just like RingNet. Both use packets with a single address flit and separate write and read channels. In contrast to RingNet, AXI4 supports packets of different sizes up to 256 flits and various data widths from 32 up to 1024 bits.

AXI4 connects memory-mapped devices using AXI Interconnect. It is built of a crossbar, a data fifo used for buffering, pipeline flip-flops used to break a critical timing path, and an address range decoder. AXI Interconnect is available as a standalone IP in the Xilinx IP Catalog.

We configured a RingNet ring and AXI4 Interconnect to have similar features and compared the implementations in terms of resource utilization and the maximum clock frequency. The configurations are described in Table V.

The implementation results were obtained just for a sample FPGA device, namely Artix7 (xc7a100tcsg324-1). The synthesis discussed in Section VIII already showed comparable results of RingNet for different types of FPGA. Therefore, the conclusions from the implementation can also be extended to the other FPGAs.

AXI4 Interconnect was configured using the AXI Interconnect RTL 1.7 generator from Vivado 2016.4. The implementations were performed using Vivado 2016.4.

The goal of the implementations was to estimate the maximum clock frequency. In order to find it, multiple implementations were performed for the requested clock frequencies changed with the resolution of 1MHz. The results obtained for the maximum obtained frequencies are presented in Table VI.

TABLE V
RingNet Ring and AXI4 Interconnect Configuration

Parameter	RingNet ring	AXI4 Interconnect
Address width	37 bits	37 bits
Data width	64 bits (single ring)	64 bits
Number and types of available interfaces	*N*×1: - 1 RI for connecting another ring or a memory device - *N* (2 ÷ 15) of NIs for connecting PEs	*N*×1: - 1 slave interface for connecting memory device, - *N* (2 ÷ 15) of master interfaces for connecting PEs
Arbiter	L2R Manager	Round-Robin arbiter
Buffering	LUTRAM-based fifos	
Performance optimization	---	Crossbar in performance optimized version named Shared-Address, Multiple-Data (SAMD); Use of pipelining FFs called AXI Register Slice
Data enable	Enable bit for each transmitted data byte.	

TABLE VI
Utilized Resources and The Maximum Frequency of RingNet Ring and AXI4 Interconnect Implementations.

		LUTs utilized	FFs utilized	Max freq. [MHz]
AXI4 Interc.	2×1	1370 (22% as RAM)	2801	268
	4×1	2205 (23% as RAM)	4480	226
	6×1	3175 (22% as RAM)	6151	192
	15×1	7181 (23% as RAM)	13650	151
RingNet ring	2×1	1185 (44% as RAM)	2047	396
	4×1	1846 (43% as RAM)	3188	392
	6×1	2557 (44% as RAM)	4343	382
	15×1	5650 (43% as RAM)	9595	365

One can compare the results obtained for the implementation and synthesis of RingNet rings (see Table IV for Artix7 and Table VI). The differences in the number of utilized LUTs reported in both tables are minor. The maximum clock frequency and the number of utilized FFs are higher for the implementation by 12% and 34% on average, respectively. Those increases are expected and are the result of register replication enabled during the implementation.

From Table VI, one can conclude that AXI4 Interconnect does not scale well and the maximum clock frequency decreases rapidly for a growing size of the AXI4 crossbar. On the other hand, RingNet, due to its optimized 3-port switches and ring topology, provides high maximum clock frequency across a wide range of network sizes. For the corresponding configurations, RingNet supports higher frequencies than AXI4 Interconnect. The increase in the maximum clock frequency is from 48% (for the 2×1 configuration) up to 142% (for the 15×1 configuration).

For all the cases RingNet ring requires less resources than AXI4 Interconnect. The FF utilization is about 40% lower for all configurations. The reduction in the number of used LUTs ranges from 16% (for the 2×1 configuration) to 27% (for the 15×1 configuration). Despite utilizing less LUTs, a RingNet

ring utilizes more LUTs as RAM than a corresponding AXI4 Interconnect, providing more buffer space.

Both RingNet rings and AXI4 Interconnect utilize a part of the resources for core elements that are shared between a number of attached PEs. Those core elements for RingNet ring are: L2R Manager, Slot Generators, and Root Interface, whereas the core elements of AXI4 Interconnect are the Round-Robin arbiter and a slave interface. For both architectures the numbers of utilized LUTs and FFs follow the equation:

$$R_{Total} = p \cdot R_{PE} + R_{core} \quad (3),$$

where p is the number of connected PEs, R_{PE} is the number of LUTs or FFs utilized per PE and R_{core} is a constant number of LUTs or FFs utilized for the core elements. The parameters R_{PE} and R_{core} can be calculated for LUTs and FFs based on the results presented in Table VI by using linear regression. Results of the regression are presented in Table VII.

TABLE VII
RESOURCES UTILIZED FOR RINGNET RING AND AXI4 INTERCONNECT IMPLEMENTATIONS.

	Resources added per PE (R_{PE})		Resources utilized for core elements (R_{core})	
	LUTs	FFs	LUTs	FFs
AXI4 Interc.	449	834	473	1133
RingNet ring	344	581	497	884

For other state-of-the-art NoCs from [3], [4], one switch is added per processing element, therefore the utilization of resources is proportional to the number of PEs. For interconnects like RingNet or AXI4, core elements are used, and their cost is shared between all the connected PEs. As already stated in the previous section, this approach can reduce the overall resource utilization of a network.

X. CONCLUSIONS

In the paper, the problems related to the design of NoCs on FPGAs are considered in the context of the typical features of FPGAs manufactured by leading vendors. As a result of these considerations, the RingNet architecture and communication protocol are proposed with the aim to exploit the specific potential of FPGA devices as much as possible. Although few FPGA-oriented NoCs have been proposed so far ([3], [4], [19], [31], [36]–[43]), the RingNet design is likely the most determinedly adopted to the typical FPGA resources and their architectures. The specific features of RingNet include: communication exclusively through system memory (large SDRAM or block RAM), control over the traffic load executed by the processing elements, FPGA-optimized 3-port switches organized into the tree-of-rings topology, distributed memory (LUTRAM) used as small buffers in the switches, and virtual cut-through switching. In the paper, it is demonstrated that RingNet features guaranteed throughput, predictable latency, and fair network access. In the paper, the simulation results demonstrate that the RingNet implementations are efficient in terms of the maximum clock frequency and resource consumption for flagship FPGA devices from major manufacturers. As compared to the widely-accepted state-of-the-art interconnection architecture AXI4, the implementations demonstrate higher maximum clock frequency and lower resource consumption in RingNet. Therefore, the authors believe that the RingNet NoC architecture and protocol may be widely adopted in FPGA-based SoC designs, especially in high-volume data processing applications, e.g., in video processing.

In complex SoCs, it may be desired to establish clock domains with individual clock frequencies, e.g., with the aim of optimizing power dissipation. We think that the RingNet ring is a natural candidate for grouping PEs with similar maximum clock frequency, or more generally, in sections with common power management. Nevertheless, the usage of rings in separate clock domains requires methods for inter-domain transitions. Considerations on such methods for RingNet would be an interesting direction for future research, in particular, when searching for techniques that would be efficient for several families of FPGA devices.

REFERENCES

[1] K. A. Helal, S. Attia, T. Ismail, and H. Mostafa, "Comparative review of NoCs in the context of ASICs and FPGAs," in *Proc. 2015 IEEE Int. Symp. Circuits and Syst. (ISCAS)*, Lisbon, pp. 1866-1869.

[2] A. Łuczak, M. Stępniewska, J. Siast, M. Domański, O. Stankiewicz, M. Kurc, and J. Konieczny, "Network-on-multi-chip (NoMC) with monitoring and debugging support," *J. Telecommun. and Inform. Techno.*, no. 3, pp. 81, 2011.

[3] P. Maidee and A. Kaviani, "Improving FPGA NoC performance using virtual cut-through switching technique," in *Proc. Int. Conf. ReConFigurable Computing and FPGAs (ReConFig)*, Mexico City, 2015, pp. 1-6.

[4] M. K. Papamichael and J. C. Hoe, "The CONNECT network-on-chip generator," *Computer*, vol. 48, no. 12, pp. 72-79, Dec. 2015.

[5] *Stratix 10 GX/SX device overview*, S10-overview, Intel Co., Oct. 2016.

[6] *Arria 10 device overview*, A10-overwiev, Altera Co., Oct. 2016.

[7] *Intel Cyclone 10 GX device overview*, C10GX51001, Intel Co., Nov. 2017.

[8] *Stratix V device overview*, SV51001, Altera Co., Oct. 2015.

[9] *Devices: 28nm device portfolio*, Altera Product Catalog, Altera Co., 2015.

[10] *Zynq UltraScale+ MPSoC*, XMP104, v2.1, Xilinx Inc., 2016.

[11] *UltraScale+ FPGAs*, XMP103, v1.10, Xilinx Inc., 2017.

[12] *UltraScale FPGA*, XMP102, v1.7, Xilinx Inc., 2016.

[13] *All programmable 7 Series*, XMP101, v1.2, Xilinx Inc., 2016.

[14] *Cost-optimized portfolio*, XMP100, v1.6, Xilinx Inc., 2016.

[15] *ECP5 and ECP5-5G family*, DS1044, v1.6, Lattice Semiconductor Co., Feb. 2016.

[16] *LatticeECP3 family data sheet*, DS1021, v02.8EA, Lattice Semiconductor Co., Mar. 2015.

[17] *LatticeECP2/M family data sheet*, DS1006, v04.1, Lattice Semiconductor Co., Sept. 2013.

[18] *Stratix V device handbook*, SV51002, v1.3, vol. 2, Altera Co., Nov. 2011.

[19] J. Rettkowski and D. Göhringer, "RAR-NoC: A reconfigurable and adaptive routable network-on-chip for FPGA-based multiprocessor systems," in *Proc. Int. Conf. ReConFigurable Computing and FPGAs (ReConFig)*, Cancun, 2014, pp. 1-6.

[20] D. Bertozzi and L. Benini, "Xpipes: A network-on-chip architecture for gigascale systems-on-chip," *IEEE Circuits and Syst. Magazine*, vol. 4, no. 2, pp. 18-31, 2004.

[21] S. Murali *et al.*, "Synthesis of predictable networks-on-chip-based interconnect architectures for chip multiprocessors," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 15, no. 8, pp. 869-880, Aug. 2007.

[22] S. Liu, A. Jantsch, and Z. Lu, "A fair and maximal allocator for single-cycle on-chip homogeneous resource allocation," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 22, no. 10, pp. 2230-2234, Oct. 2014.

[23] T. Wehbe and X. Wang, "Secure and dependable NoC-connected systems on an FPGA chip," *IEEE Trans. on Reliability*, vol. 65, no. 4, pp. 1852-1863, Dec. 2016.
[24] E. Todorovich, M. Leonetti, and R. Brinks, "An advanced NoC with debug services on FPGA," in *Proc. Southern Conf. Programmable Logic (SPL)*, Buenos Aires, 2014, pp. 1-6.
[25] M. Zhu, Y. Jiang, M. Yang and L. De Luna, "A scalable parameterized NoC emulator built upon Xilinx Virtex-7 FPGA," in *Proc. Int. Conf. Syst. Engineering (ICSEng)*, Las Vegas, 2017, pp. 287-290.
[26] H. M. Kamali, K. Z. Azar and S. Hessabi, "DuCNoC: A high-throughput FPGA-based NoC simulator using dual-clock lightweight router micro-architecture," *IEEE Trans. Comput.*, vol. 67, no. 2, pp. 208-221, Feb. 2018.
[27] Y. J. Yoon, N. Concer, M. Petracca, and L. P. Carloni, "Virtual channels and multiple physical networks: Two alternatives to improve NoC performance," *IEEE Trans. Computer-Aided Design of Integr. Circuits and Syst.*, vol. 32, no. 12, pp. 1906-1919, Dec. 2013.
[28] Yong Ho Song and T. M. Pinkston, "A progressive approach to handling message-dependent deadlock in parallel computer systems," *IEEE Trans. Parallel and Distrib. Syst.*, vol. 14, no. 3, pp. 259-275, Mar. 2003.
[29] M. A. Al Faruque, T. Ebi, and J. Henkel, "AdNoC: Runtime adaptive network-on-chip architecture," *IEEE Trans. Very Large Integr. (VLSI) Syst.*, vol. 20, no. 2, pp. 257-269, Feb. 2012.
[30] J. Postman, T. Krishna, C. Edmonds, L. S. Peh, and P. Chiang, "SWIFT: A low-power network-on-chip implementing the token flow control router architecture with swing-reduced interconnects," *IEEE Trans. on Very Large Scale Integr. (VLSI) Syst.*, vol. 21, no. 8, pp. 1432-1446, Aug. 2013.
[31] M. K. Papamichael and J. C. Hoe, "CONNECT: Re-examining conventional wisdom for designing NoCs in the context of FPGAs," in *Proc. ACM/SIGDA Int. Symp. Field Programmable Gate Arrays*, New York, 2012, pp. 37-46.
[32] *Vivado AXI reference guide*, UG1037, v3.0, Xilinx Inc., June, 2015.
[33] M. S. Abdelfattah and V. Betz, "Power analysis of embedded NoCs on FPGAs and comparison with custom buses," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 24, no. 1, pp. 165-177, Jan. 2016.
[34] M. S. Abdelfattah and V. Betz, "The case for embedded networks on chip on field-programmable gate arrays," *IEEE Micro*, vol. 34, no. 1, pp. 80-89, Jan.-Feb. 2014.
[35] M. S. Abdelfattah, A. Bitar and V. Betz, "Design and applications for embedded networks-on-chip on FPGAs," *IEEE Trans. Comput.*, vol. 66, no. 6, pp. 1008-1021, June 1 2017.
[36] N. Kapre and J. Gray, "Hoplite: Building austere overlay NoCs for FPGAs," in *Proc. Int. Conf. Field Programmable Logic and Applications (FPL)*, London, 2015, pp. 1-8.
[37] S. Wasly, R. Pellizzoni and N. Kapre, "HopliteRT: An efficient FPGA NoC for real-time applications," in *Proc. Int. Conf. Field Programmable Technology (ICFPT)*, Melbourne, 2017, pp. 64-71.
[38] K. Vipin, J. Gray and N. Kapre, "Enabling partial reconfiguration and low latency routing using segmented FPGA NoCs," in *Proc. Int. Conf. Field Programmable Logic and Applications (FPL)*, Ghent, 2017, pp. 1-8.
[39] N. Kapre, "On bit-serial NoCs for FPGAs," in *Proc. IEEE Int. Symp. Field-Programmable Custom Computing Machines (FCCM)*, Napa, CA, 2017, pp. 32-39.
[40] N. Kapre, "Implementing FPGA overlay NoCs using the Xilinx UltraScale memory cascades," in *Proc. IEEE Annu. Int. Symp. Field-Programmable Custom Computing Machines (FCCM)*, Napa, CA, 2017, pp. 40-47.
[41] N. Kapre, "Deflection-routed butterfly fat trees on FPGAs," in *Proc. Int. Conf. Field Programmable Logic and Applications (FPL)*, Ghent, 2017, pp. 1-8.
[42] S. N. Shelke and P. B. Patil, "Low-latency, low-area overhead and high throughput NoC architecture for FPGA based computing system," in *Proc. Int. Conf. Electronic Syst., Signal Process. and Computing Technol.*, Nagpur, 2014, pp. 53-57.
[43] P. Maidee, A. Kaviani and K. Zeng, "LinkBlaze: Efficient global data movement for FPGAs," in *Proc. Int. Conf. on ReConFigurable Computing and FPGAs (ReConFig)*, Cancun, 2017, pp. 1-8.
[44] W. Snyder, "Introduction to Verilator," [Online], Available: https://www.veripool.org/wiki/verilator, Accessed on: Dec. 19, 2017.
[45] R. Kumar and A. Gordon-Ross, "MACS: A highly customizable low-latency communication architecture," *IEEE Trans. Parallel and Distrib. Syst.*, vol. 27, no. 1, pp. 237-249, Jan. 1 2016.
[46] M. Domański, J. Konieczny, M. Kurc, A. Łuczak, J. Siast, O. Stankiewicz, and K. Wegner, "Fast depth estimation on mobile platforms and FPGA devices," in *Proc.* 2015 3DTV-Conf.: The True Vision - Capture, Transmission and Display of 3D Video (3DTV-CON), Lisbon, 2015.
[47] *Synopsys FPGA synthesis user guide*, Synopsys Inc., Apr. 2017.
[48] M. S. Abdelfattah and V. Betz, "LYNX: CAD for FPGA-based networks-on-chip," in *Proc. 26th Int. Conf. on Field Programmable Logic and Applications (FPL)*, Lausanne, 2016, pp. 1-10.
[49] S. Abba and Jeong-A Lee, "A parametric-based performance evaluation and design trade-offs for interconnect architectures using FPGAs for networks-on-chip," Microprocessors and Microsystems, vol. 38, no. 5, pp. 375-398, July 2014.

Jakub Siast received the M.Sc. degree (2009) from the Faculty of Electronics and Telecommunications, Poznan University of Technology, where he is currently pursuing the Ph.D. degree. He is the co-author of several papers related to hardware implementation in FPGA, 3D video coding and depth estimation. His current research interests include FPGA and microprocessor architecture design.

Adam Łuczak received his M.Sc. and Ph.D. degrees from Poznan University of Technology in 1997 and 2001, respectively. Member of Polish Society Theoretical and Applied Electrical Engineering (PTETiS). His research activities include video coders control, and hardware implementations (FPGA) of digital signal processing algorithms. He received Annual Fellowship for Young Scientist from the Foundation for Polish Science (FNP) and also a Group Award of the Minister of the National Education for research in the National Education for research in image compression.

Marek Domański received M.S. (1978), Ph.D. (1983) and Habilitation (1990) degrees from Poznań University of Technology, Poland, where he is a Professor (since 1993) and leads the Chair (Department) of Multimedia Telecommunications and Microelectronics. He co-authored highly ranked technology proposals submitted in response to MPEG calls for scalable video compression (2004) and 3D video coding (2011). He also leaded the team that developed one of the very first AVC decoders for TV set-top boxes (2004) and various AVC, HEVC and AAC HE codec implementations and improvements. He is an author of 3 books, 14 patents and over 300 research papers, mostly on image, video and audio compression, virtual navigation, free-viewpoint television, image processing, multimedia systems, 3D video and color image technology, digital filters and multidimensional signal processing. He was General Chairman/Co-Chairman and host of several international conferences: Picture Coding Symposium, PCS 2012; IEEE Int. Conf. on Advanced and Signal-based Surveillance, AVSS 2013, European Signal Processing Conference, EUSIPCO 2007; 73rd and 112nd Meetings of MPEG, 69th JPEG Meeting and others. He served as a member of various steering, program and editorial committees of international journals and international conferences.

Appendix I. Simulation results of the average latency test

In Appendix I, the average latency for the read and write channel is presented for various loads and network sizes. The parameters of the test are:

- *R*: Multiplication degree of a root level, i.e., number of parallel rings used at the root level, set in the range of 1–4.
- *F*: The number of 1st level rings, set in the range of 1–5.
- *G*: The number of packet generators (PGs) connected to a single 1st level ring, set in the range of 1–15.
- Logical channel load. The aggregated load generated by all PGs is set in the range of 0%–100% of the theoretical throughput T_{RW_MAX} (1). The logical channel load is separately set for the read and write channel.

TABLE I.1
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR READ CHANNEL FOR 1 RING USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS EQUAL 100%.

		Number of PGs connected to a 1st level ring (*G*)					
		1	2	3	4	7	15
Number of 1st level rings (*F*)	1	88	166	209	252	394	764
	2	210	297	390	478	749	1528
	3	318	459	594	729	1127	2300
	4	417	602	780	958	1516	3058
	5	514	750	980	1193	1877	3672

TABLE I.2
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR READ CHANNEL FOR 2 PARALLEL RINGS USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS EQUAL 100%.

		Number of PGs connected to a 1st level ring (*G*)					
		1	2	3	4	7	15
Number of 1st level rings (*F*)	1	---	---	---	---	---	---
	2	124	193	236	280	421	791
	3	272	341	407	476	667	1257
	4	353	449	541	632	892	1672
	5	436	550	666	780	1108	2002

TABLE I.3
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR READ CHANNEL FOR 3 PARALLEL RINGS USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS EQUAL 100%.

		Number of PGs connected to a 1st level ring (G)					
		1	2	3	4	7	15
Number of 1st level rings (F)	1	---	---	---	---	---	---
	2	---	---	---	---	---	---
	3	152	220	264	308	449	819
	4	311	365	432	494	674	1188
	5	374	452	533	610	830	1441

TABLE I.4
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR READ CHANNEL FOR 4 PARALLEL RINGS USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS EQUAL 100%.

		Number of PGs connected to a 1st level ring (G)					
		1	2	3	4	7	15
Number of 1st level rings (F)	1	---	---	---	---	---	---
	2	---	---	---	---	---	---
	3	---	---	---	---	---	---
	4	156	223	272	316	454	824
	5	349	401	459	515	689	1157

TABLE I.5
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR READ CHANNEL FOR 1 RING USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS ARE IN THE RANGE 92%–97%.

		Number of PGs connected to a 1st level ring (G)					
		1	2	3	4	7	15
Number of 1st level rings (F)	1	95	118	120	122	147	194
	2	113	134	142	145	182	225
	3	129	148	155	161	185	241
	4	130	151	157	163	189	245
	5	140	163	169	175	201	258

TABLE I.6

AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR READ CHANNEL FOR 2 PARALLEL RINGS USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS ARE IN THE RANGE 92%–97%.

		Number of PGs connected to a 1st level ring (G)					
		1	2	3	4	7	15
Number of 1st level rings (F)	1	---	---	---	---	---	---
	2	132	144	146	154	179	226
	3	139	156	158	163	186	235
	4	141	158	163	167	190	238
	5	155	170	175	180	202	253

TABLE I.7

AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR READ CHANNEL FOR 3 PARALLEL RINGS USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS ARE IN THE RANGE 92%–97%.

		Number of PGs connected to a 1st level ring (G)					
		1	2	3	4	7	15
Number of 1st level rings (F)	1	---	---	---	---	---	---
	2	---	---	---	---	---	---
	3	152	169	171	177	202	249
	4	149	165	170	174	196	244
	5	159	174	178	186	207	254

TABLE I.8

AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR READ CHANNEL FOR 4 PARALLEL RINGS USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS ARE IN THE RANGE 92%–97%.

		Number of PGs connected to a 1st level ring (G)					
		1	2	3	4	7	15
Number of 1st level rings (F)	1	---	---	---	---	---	---
	2	---	---	---	---	---	---
	3	---	---	---	---	---	---
	4	155	170	176	181	204	254
	5	162	179	184	188	211	259

TABLE I.9

AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR READ CHANNEL FOR 1 RING USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS ARE IN THE RANGE 85%–92%.

		Number of PGs connected to a 1st level ring (*G*)					
		1	2	3	4	7	15
Number of 1st level rings (*F*)	1	94	118	119	120	144	189
	2	113	132	139	141	169	214
	3	127	145	151	157	178	230
	4	128	147	153	157	179	231
	5	138	158	164	168	192	242

TABLE I.10

AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR READ CHANNEL FOR 2 PARALLEL RINGS USED AT ROOT LEVEL, AND WRITE CHANNEL LOAD IS IN THE RANGE 85%–91%, READ CHANNEL LOAD IS IN THE RANGE 85%–92%.

		Number of PGs connected to a 1st level ring (*G*)					
		1	2	3	4	7	15
Number of 1st level rings (*F*)	1	---	---	---	---	---	---
	2	130	143	145	151	175	221
	3	138	155	156	161	183	229
	4	140	156	160	164	186	232
	5	153	168	172	177	197	244

TABLE I.11

AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR READ CHANNEL FOR 3 PARALLEL RINGS USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS ARE IN THE RANGE 85%–91%.

		Number of PGs connected to a 1st level ring (*G*)					
		1	2	3	4	7	15
Number of 1st level rings (*F*)	1	---	---	---	---	---	---
	2	---	---	---	---	---	---
	3	150	167	169	174	197	243
	4	147	163	168	171	193	239
	5	158	173	177	182	203	249

TABLE I.12

AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR READ CHANNEL FOR 4 PARALLEL RINGS USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS ARE IN THE RANGE 85%–91%.

		Number of PGs connected to a 1st level ring (G)					
		1	2	3	4	7	15
Number of 1st level rings (F)	1	---	---	---	---	---	---
	2	---	---	---	---	---	---
	3	---	---	---	---	---	---
	4	152	169	174	177	199	247
	5	162	176	181	186	208	254

TABLE I.13

AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR READ CHANNEL FOR 1 RING USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS ARE IN THE RANGE 69%–73%.

		Number of PGs connected to a 1st level ring (G)					
		1	2	3	4	7	15
Number of 1st level rings (F)	1	93	116	117	118	140	183
	2	111	130	136	137	160	203
	3	126	142	146	151	170	216
	4	125	143	147	150	170	214
	5	136	153	158	160	181	224

TABLE I.14

AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR READ CHANNEL FOR 2 PARALLEL RINGS USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS ARE IN THE RANGE 69%–73%.

		Number of PGs connected to a 1st level ring (G)					
		1	2	3	4	7	15
Number of 1st level rings (F)	1	---	---	---	---	---	---
	2	129	141	142	148	170	213
	3	137	153	153	158	178	221
	4	138	153	156	160	180	223
	5	150	165	168	172	190	235

TABLE I.15
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR READ CHANNEL FOR 3 PARALLEL RINGS USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS ARE IN THE RANGE 69%–73%.

		Number of PGs connected to a 1st level ring (G)					
		1	2	3	4	7	15
Number of 1st level rings (F)	1	---	---	---	---	---	---
	2	---	---	---	---	---	---
	3	147	163	164	169	190	235
	4	145	160	164	168	188	232
	5	156	170	174	178	197	241

TABLE I.16
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR READ CHANNEL FOR 4 PARALLEL RINGS USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS ARE IN THE RANGE 69%–73%.

		Number of PGs connected to a 1st level ring (G)					
		1	2	3	4	7	15
Number of 1st level rings (F)	1	---	---	---	---	---	---
	2	---	---	---	---	---	---
	3	---	---	---	---	---	---
	4	149	164	169	172	193	238
	5	158	173	178	182	202	246

TABLE I.17
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR READ CHANNEL FOR 1 RING USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS ARE IN THE RANGE 27%–28%.

		Number of PGs connected to a 1st level ring (G)					
		1	2	3	4	7	15
Number of 1st level rings (F)	1	93	115	115	115	136	176
	2	110	127	132	132	152	192
	3	123	138	141	144	162	205
	4	122	138	141	143	161	204
	5	132	147	151	153	171	213

TABLE I.18
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR READ CHANNEL FOR 2 PARALLEL RINGS USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS ARE IN THE RANGE 27%–28%.

		Number of PGs connected to a 1st level ring (G)					
		1	2	3	4	7	15
Number of 1st level rings (F)	1	---	---	---	---	---	---
	2	127	138	138	143	162	204
	3	134	149	149	152	171	213
	4	136	149	152	154	173	215
	5	147	160	162	166	184	227

TABLE I.19
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR READ CHANNEL FOR 3 PARALLEL RINGS USED AT ROOT LEVEL, AND WRITE CHANNEL LOAD IS IN THE RANGE 27%–28%, READ CHANNEL LOAD EQUALS 27%.

		Number of PGs connected to a 1st level ring (G)					
		1	2	3	4	7	15
Number of 1st level rings (F)	1	---	---	---	---	---	---
	2	---	---	---	---	---	---
	3	142	157	158	162	181	223
	4	141	155	159	161	180	222
	5	152	165	168	172	190	233

TABLE I.20
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR READ CHANNEL FOR 4 PARALLEL RINGS USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS EQUAL 27%.

		Number of PGs connected to a 1st level ring (G)					
		1	2	3	4	7	15
Number of 1st level rings (F)	1	---	---	---	---	---	---
	2	---	---	---	---	---	---
	3	---	---	---	---	---	---
	4	144	159	162	165	184	226
	5	154	168	171	175	193	236

TABLE I.21
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR WRITE CHANNEL FOR 1 RING USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS EQUAL 100%.

		Number of PGs connected to a 1st level ring (G)					
		1	2	3	4	7	15
Number of 1st level rings (F)	1	95	172	216	260	401	771
	2	217	304	398	485	756	1535
	3	325	468	601	736	1134	2307
	4	425	611	790	969	1523	3064
	5	521	762	978	1201	1879	3712

TABLE I.22
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR WRITE CHANNEL FOR 2 PARALLEL RINGS USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS EQUAL 100%.

		Number of PGs connected to a 1st level ring (G)					
		1	2	3	4	7	15
Number of 1st level rings (F)	1	---	---	---	---	---	---
	2	132	200	238	287	428	798
	3	296	365	430	500	691	1281
	4	379	478	571	661	916	1700
	5	469	589	699	812	1153	2037

TABLE I.23
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR WRITE CHANNEL FOR 3 PARALLEL RINGS USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS EQUAL 100%.

		Number of PGs connected to a 1st level ring (G)					
		1	2	3	4	7	15
Number of 1st level rings (F)	1	---	---	---	---	---	---
	2	---	---	---	---	---	---
	3	163	224	275	318	456	830
	4	334	393	454	515	696	1210
	5	402	476	558	633	856	1466

TABLE I.24
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR WRITE CHANNEL FOR 4 PARALLEL RINGS USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS EQUAL 100%.

		Number of PGs connected to a 1st level ring (G)					
		1	2	3	4	7	15
Number of 1st level rings (F)	1	---	---	---	---	---	---
	2	---	---	---	---	---	---
	3	---	---	---	---	---	---
	4	168	230	279	323	462	834
	5	368	423	481	538	708	1175

TABLE I.25
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR WRITE CHANNEL FOR 1 RING USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS ARE IN THE RANGE 92%–97%.

		Number of PGs connected to a 1st level ring (G)					
		1	2	3	4	7	15
Number of 1st level rings (F)	1	100	125	127	129	154	202
	2	120	141	149	152	187	232
	3	135	155	162	167	193	249
	4	136	158	163	170	196	252
	5	147	169	177	183	210	266

TABLE I.26
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR WRITE CHANNEL FOR 2 PARALLEL RINGS USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS ARE IN THE RANGE 92%–97%.

		Number of PGs connected to a 1st level ring (G)					
		1	2	3	4	7	15
Number of 1st level rings (F)	1	---	---	---	---	---	---
	2	140	152	154	162	187	234
	3	146	163	166	170	194	242
	4	148	165	170	174	198	246
	5	162	176	182	187	209	261

TABLE I.27
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR WRITE CHANNEL FOR 3 PARALLEL RINGS USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS ARE IN THE RANGE 92%–97%.

		Number of PGs connected to a 1st level ring (G)					
		1	2	3	4	7	15
Number of 1st level rings (F)	1	---	---	---	---	---	---
	2	---	---	---	---	---	---
	3	159	177	181	187	211	260
	4	157	173	178	183	205	254
	5	167	183	187	194	215	263

TABLE I.28
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR WRITE CHANNEL FOR 4 PARALLEL RINGS USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS ARE IN THE RANGE 92%–97%.

		Number of PGs connected to a 1st level ring (G)					
		1	2	3	4	7	15
Number of 1st level rings (F)	1	---	---	---	---	---	---
	2	---	---	---	---	---	---
	3	---	---	---	---	---	---
	4	162	179	185	191	213	264
	5	170	185	191	195	218	267

TABLE I.29
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR WRITE CHANNEL FOR 1 RING USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS ARE IN THE RANGE 85%–92%.

		Number of PGs connected to a 1st level ring (G)					
		1	2	3	4	7	15
Number of 1st level rings (F)	1	99	124	126	127	151	196
	2	119	139	146	148	175	222
	3	134	152	158	164	185	237
	4	134	154	159	164	186	237
	5	145	164	172	176	200	249

TABLE I.30
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR WRITE CHANNEL FOR 2 PARALLEL RINGS USED AT ROOT LEVEL, AND WRITE CHANNEL LOAD IS IN THE RANGE 85%–91%, READ CHANNEL LOAD IS IN THE RANGE 85%–92%.

		Number of PGs connected to a 1st level ring (G)					
		1	2	3	4	7	15
Number of 1st level rings (F)	1	---	---	---	---	---	---
	2	138	151	152	159	182	228
	3	145	162	163	168	190	236
	4	147	163	167	171	193	238
	5	160	175	179	184	204	252

TABLE I.31
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR WRITE CHANNEL FOR 3 PARALLEL RINGS USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS ARE IN THE RANGE 85%–91%.

		Number of PGs connected to a 1st level ring (G)					
		1	2	3	4	7	15
Number of 1st level rings (F)	1	---	---	---	---	---	---
	2	---	---	---	---	---	---
	3	157	175	178	183	206	254
	4	155	171	176	180	202	248
	5	166	181	185	191	211	258

TABLE I.32
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR WRITE CHANNEL FOR 4 PARALLEL RINGS USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS ARE IN THE RANGE 85%–91%.

		Number of PGs connected to a 1st level ring (G)					
		1	2	3	4	7	15
Number of 1st level rings (F)	1	---	---	---	---	---	---
	2	---	---	---	---	---	---
	3	---	---	---	---	---	---
	4	159	177	182	186	207	257
	5	168	184	188	194	215	262

TABLE I.33
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR WRITE CHANNEL FOR 1 RING USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS ARE IN THE RANGE 69%–73%.

		Number of PGs connected to a 1st level ring (*G*)					
		1	2	3	4	7	15
Number of 1st level rings (*F*)	1	97	122	123	124	146	189
	2	118	137	143	144	166	210
	3	131	148	153	157	175	221
	4	130	149	153	156	175	220
	5	142	160	165	168	188	231

TABLE I.34
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR WRITE CHANNEL FOR 2 PARALLEL RINGS USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS ARE IN THE RANGE 69%–73%.

		Number of PGs connected to a 1st level ring (*G*)					
		1	2	3	4	7	15
Number of 1st level rings (*F*)	1	---	---	---	---	---	---
	2	136	148	149	155	177	221
	3	142	159	160	164	184	227
	4	144	159	163	166	186	229
	5	158	172	175	179	197	242

TABLE I.35
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR WRITE CHANNEL FOR 3 PARALLEL RINGS USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS ARE IN THE RANGE 69%–73%.

		Number of PGs connected to a 1st level ring (*G*)					
		1	2	3	4	7	15
Number of 1st level rings (*F*)	1	---	---	---	---	---	---
	2	---	---	---	---	---	---
	3	153	170	172	177	198	243
	4	151	167	172	175	196	240
	5	163	178	182	186	205	250

TABLE I.36
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR WRITE CHANNEL FOR 4 PARALLEL RINGS USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS ARE IN THE RANGE 69%–73%.

		Number of PGs connected to a 1st level ring (G)					
		1	2	3	4	7	15
Number of 1st level rings (F)	1	---	---	---	---	---	---
	2	---	---	---	---	---	---
	3	---	---	---	---	---	---
	4	154	171	175	179	200	246
	5	166	180	185	189	209	254

TABLE I.37
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR WRITE CHANNEL FOR 1 RING USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS ARE IN THE RANGE 27%–28%.

		Number of PGs connected to a 1st level ring (G)					
		1	2	3	4	7	15
Number of 1st level rings (F)	1	91	114	114	114	135	176
	2	117	134	139	138	159	200
	3	123	138	141	144	162	205
	4	121	138	140	143	161	204
	5	139	154	158	160	178	220

TABLE I.38
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR WRITE CHANNEL FOR 2 PARALLEL RINGS USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS ARE IN THE RANGE 27%–28%.

		Number of PGs connected to a 1st level ring (G)					
		1	2	3	4	7	15
Number of 1st level rings (F)	1	---	---	---	---	---	---
	2	134	145	145	150	169	211
	3	134	149	149	152	171	213
	4	135	149	152	154	174	215
	5	154	167	169	173	191	234

TABLE I.39
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR WRITE CHANNEL FOR 3 PARALLEL RINGS USED AT ROOT LEVEL, AND WRITE CHANNEL LOAD IS IN THE RANGE 27%–28%, READ CHANNEL LOAD EQUALS 27%.

		Number of PGs connected to a 1st level ring (*G*)					
		1	2	3	4	7	15
Number of 1st level rings (*F*)	1	---	---	---	---	---	---
	2	---	---	---	---	---	---
	3	142	158	159	162	181	224
	4	141	155	159	162	181	223
	5	159	172	175	179	197	240

TABLE I.40
AVERAGE LATENCY (EXPRESSED IN CLOCK CYCLES) FOR WRITE CHANNEL FOR 4 PARALLEL RINGS USED AT ROOT LEVEL, AND WRITE AND READ CHANNEL LOADS EQUAL 27%.

		Number of PGs connected to a 1st level ring (*G*)					
		1	2	3	4	7	15
Number of 1st level rings (*F*)	1	---	---	---	---	---	---
	2	---	---	---	---	---	---
	3	---	---	---	---	---	---
	4	144	158	162	165	184	226
	5	161	175	178	182	200	243

Appendix II. Load-latency curves for RingNet

In Appendix II, load-latency curves for RingNet of various sizes are presented. The parameters of the test are:

- *R*: Multiplication degree of a root level, i.e., number of parallel rings used at the root level, set in the range of 1–4.
- *F*: The number of 1st level rings, set in the range of 1–5.
- *G*: The number of packet generators (PGs) connected to a single 1st level ring, set in the range of 1–15.
- Logical channel load. The aggregated load generated by all PGs is set in the range of 0%–100% of the theoretical throughput T_{RW_MAX} (1). The logical channel load is separately set for the read and write channel.

Error bars represent the minimum and maximum average latency of the channel at given load when the second channel load changes in range 0% to 100%. Dots denote the results for the load of the second channel is in the range 85% to 92% of its maximum throughput.

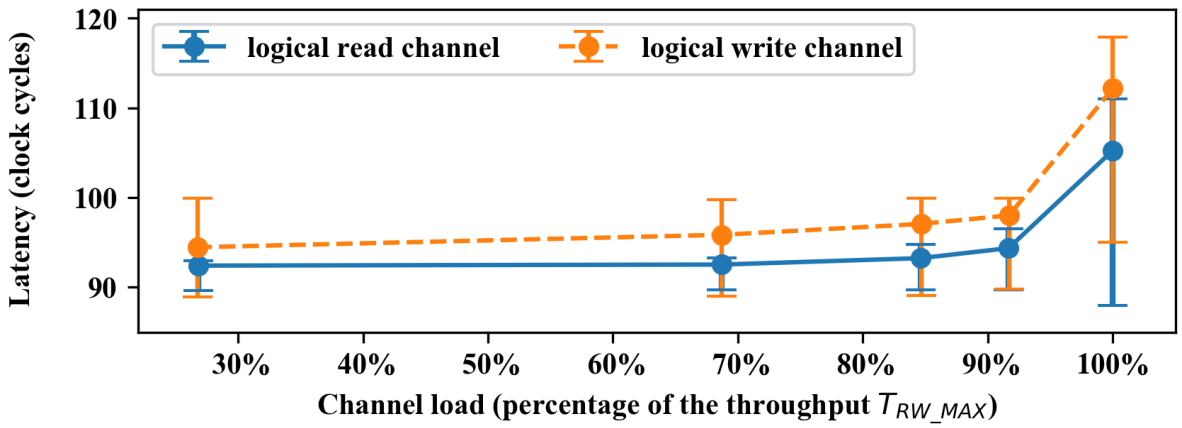

Fig.II.1. Load-latency curve for RingNet with one 1st level ring and 1 PG connected at the ring, and one ring used at the network root (F=1 × G=1, R=1).

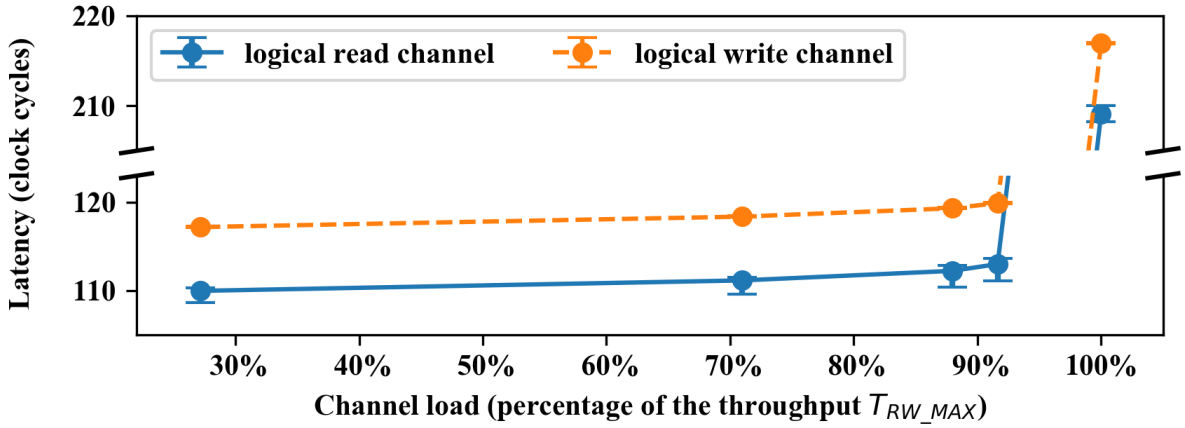

Fig.II.2. Load-latency curve for RingNet with 2 1st level rings and 1 PG connected at each ring, and one ring used at the network root (F=2 × G=1, R=1).

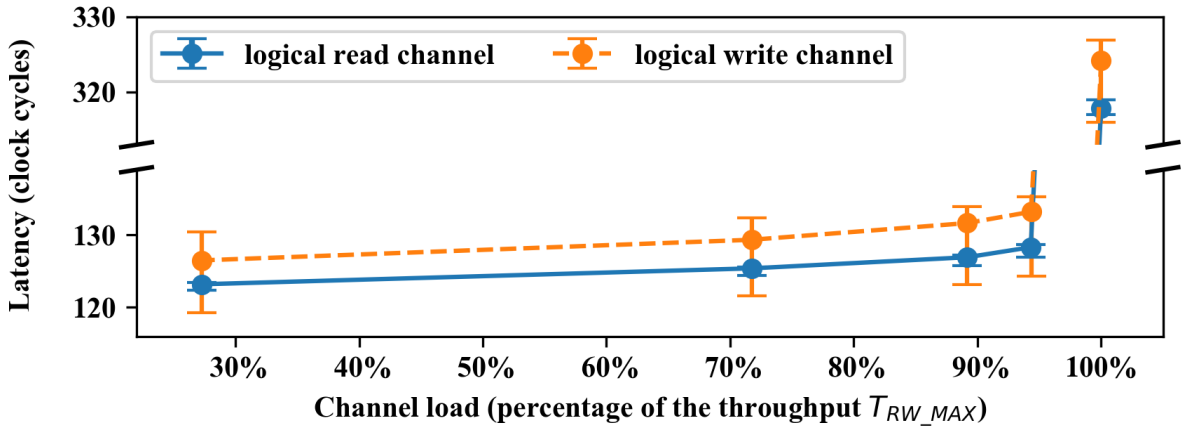

Fig.II.3. Load-latency curve for RingNet with 3 1st level rings and 1 PG connected at each ring, and one ring used at the network root (F=3 × G=1, R=1).

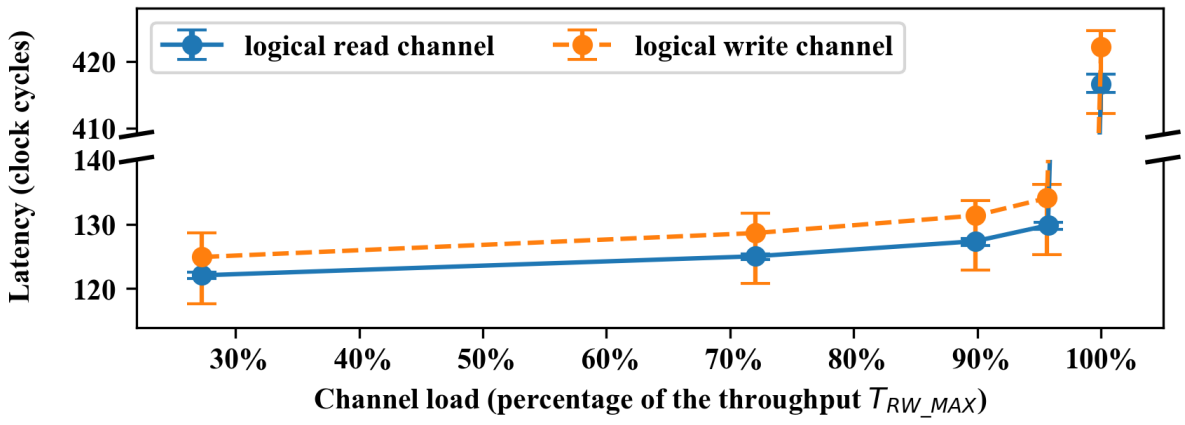

Fig.II.4. Load-latency curve for RingNet with 4 1st level rings and 1 PG connected at each ring, and one ring used at the network root (F=4 × G=1, R=1).

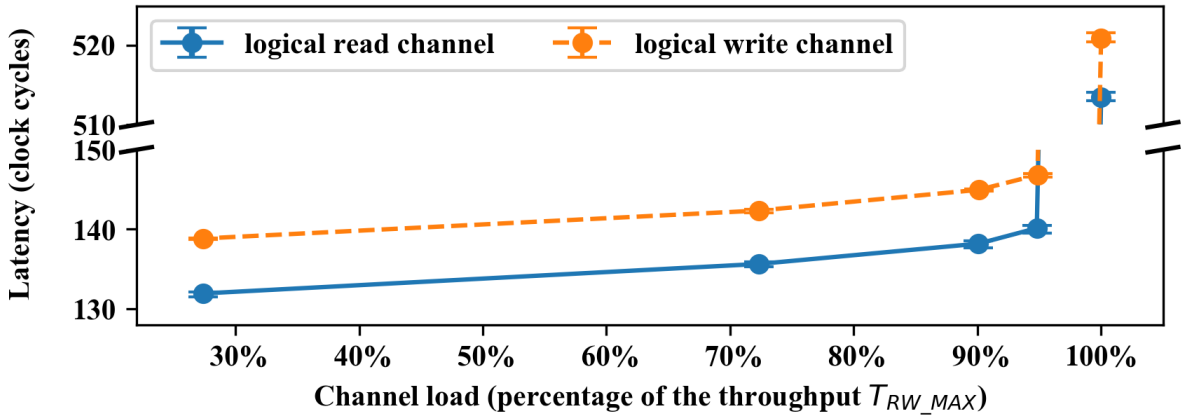

Fig.II.5. Load-latency curve for RingNet with 5 1st level rings and 1 PG connected at each ring, and one ring used at the network root (F=5 × G=1, R=1).

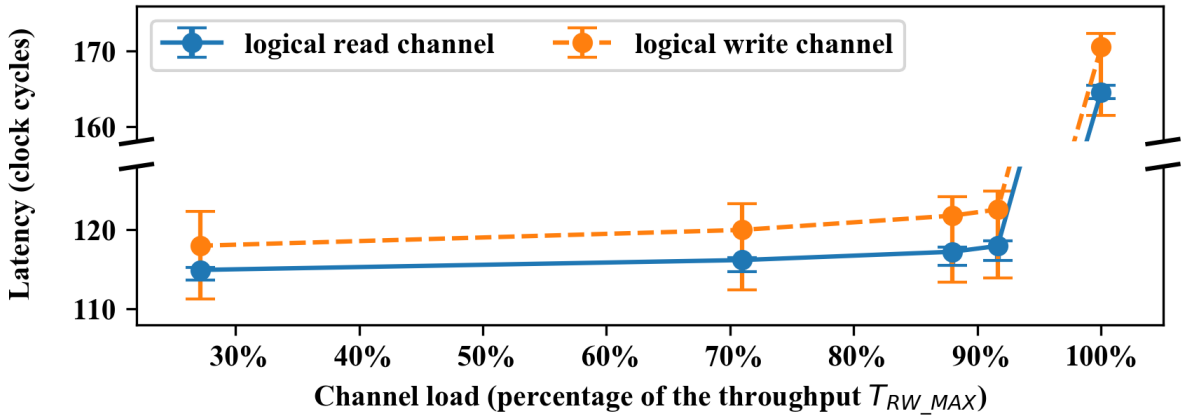

Fig.II.6. Load-latency curve for RingNet with one 1st level ring and 2 PGs connected at the ring, and one ring used at the network root (F=1 × G=2, R=1).

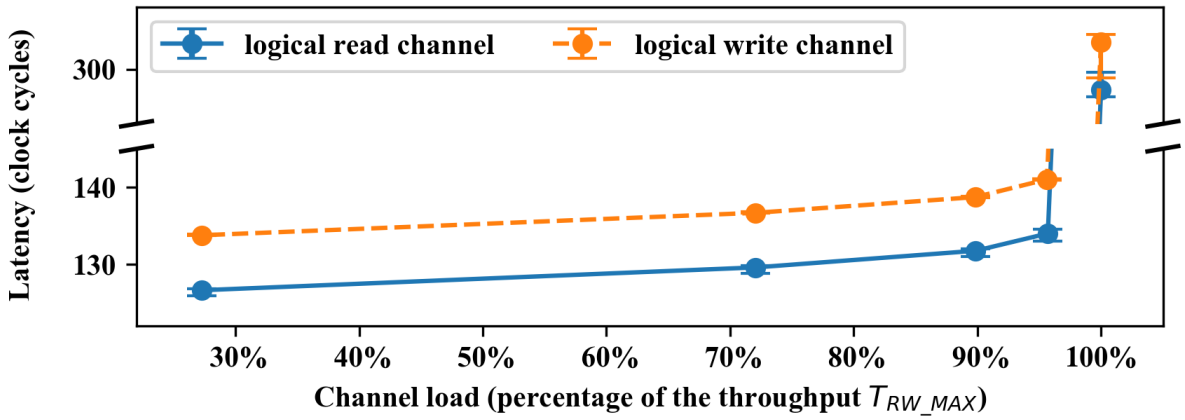

Fig.II.7. Load-latency curve for RingNet with 2 1st level rings and 2 PGs connected at each ring, and one ring used at the network root (F=2 × G=2, R=1).

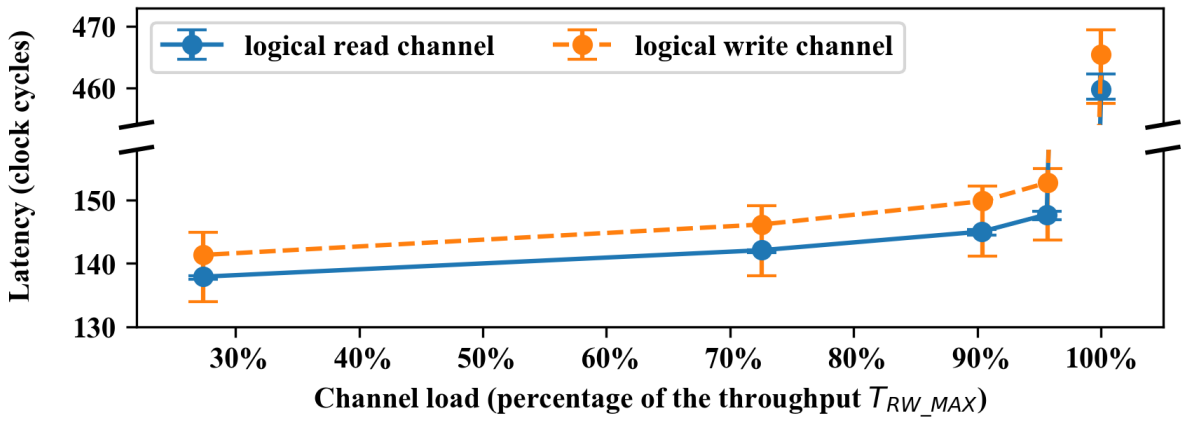

Fig.II.8. Load-latency curve for RingNet with 3 1st level rings and 2 PGs connected at each ring, and one ring used at the network root (F=3 × G=2, R=1).

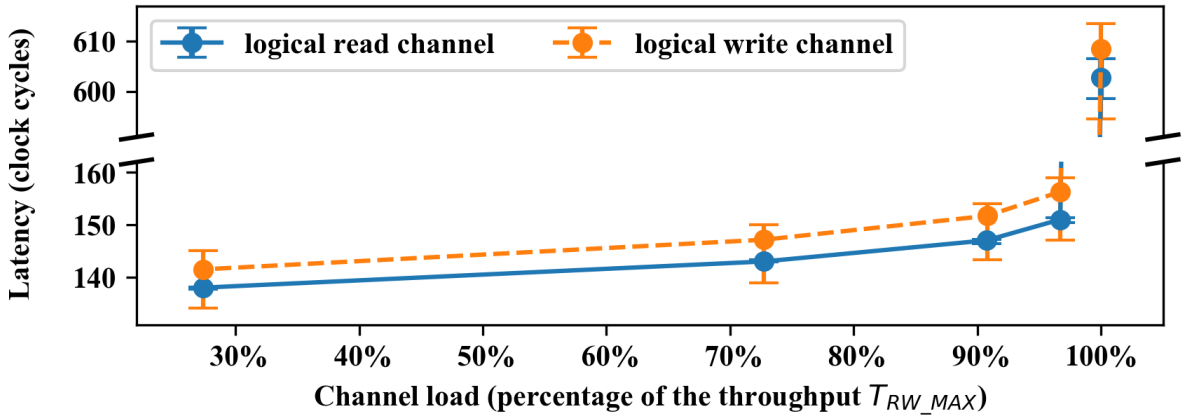

Fig.II.9. Load-latency curve for RingNet with 4 1st level rings and 2 PGs connected at each ring, and one ring used at the network root (F=4 × G=2, R=1).

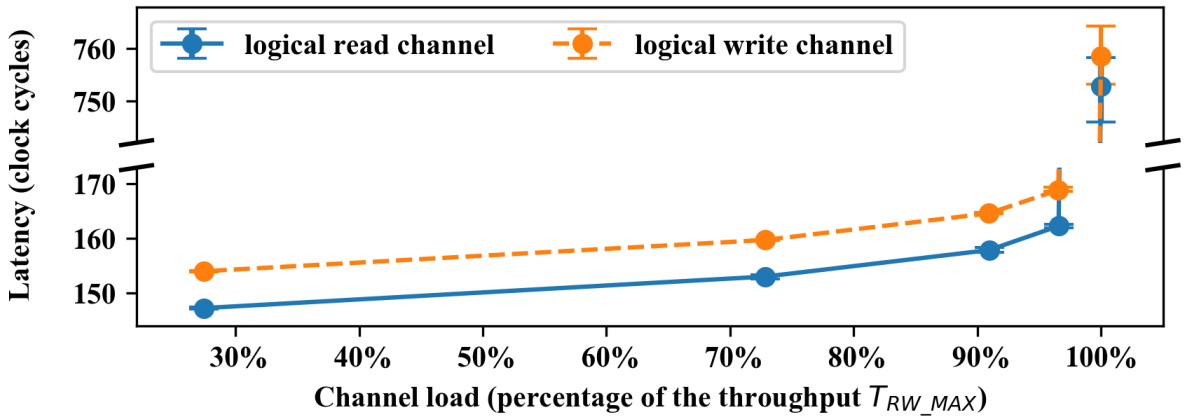

Fig.II.10. Load-latency curve for RingNet with 5 1st level rings and 2 PGs connected at each ring, and one ring used at the network root (F=5 × G=2, R=1).

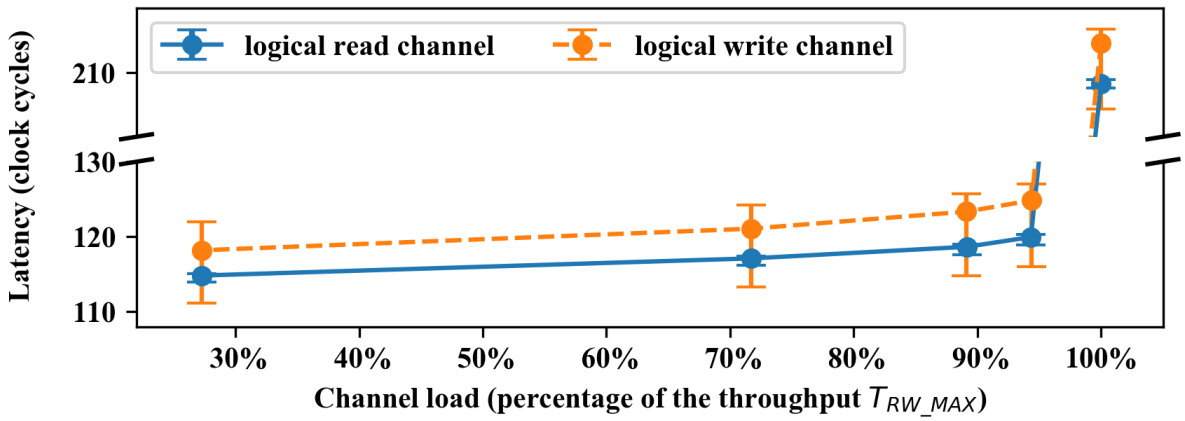

Fig.II.11. Load-latency curve for RingNet with one 1st level ring and 3 PGs connected at the ring, and one ring used at the network root (F=1 × G=3, R=1).

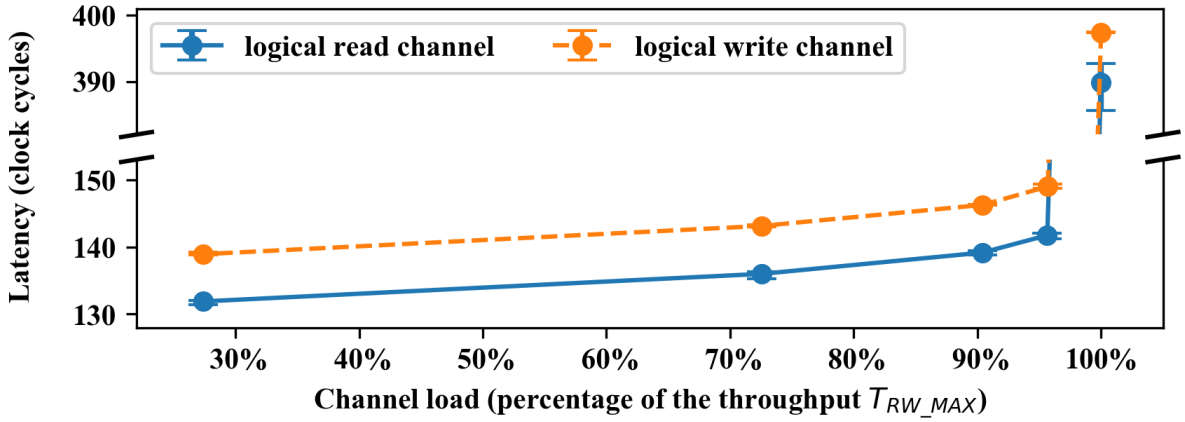

Fig.II.12. Load-latency curve for RingNet with 2 1st level rings and 3 PGs connected at each ring, and one ring used at the network root (F=2 × G=3, R=1).

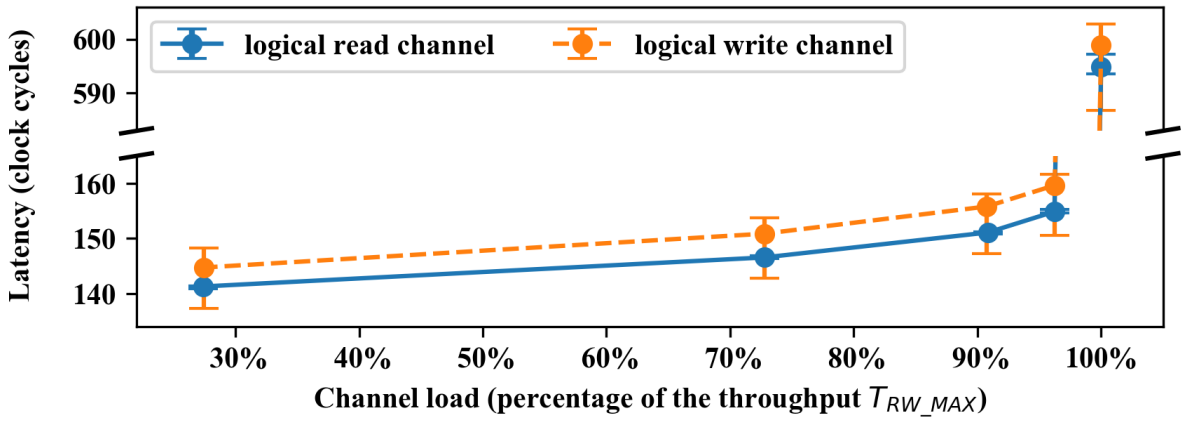

Fig.II.13. Load-latency curve for RingNet with 3 1st level rings and 3 PGs connected at each ring, and one ring used at the network root (F=3 × G=3, R=1).

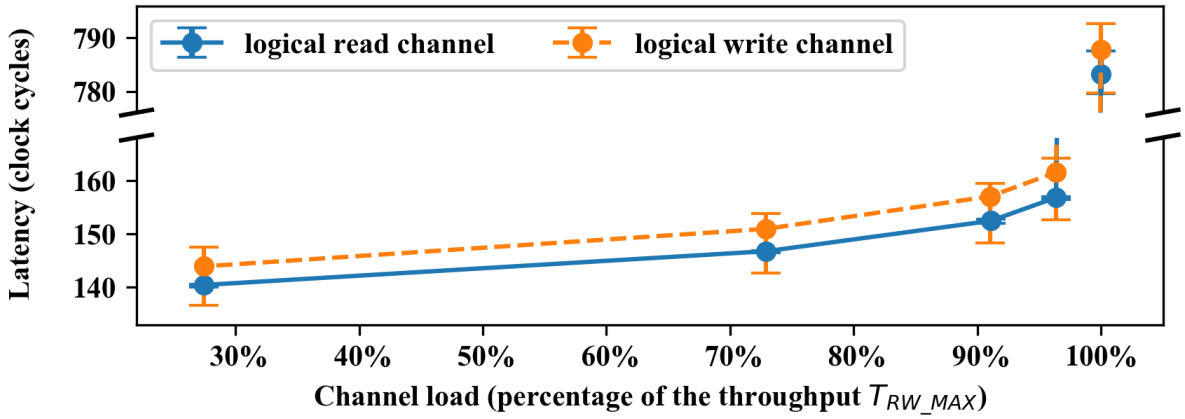

Fig.II.14. Load-latency curve for RingNet with 4 1st level rings and 3 PGs connected at each ring, and one ring used at the network root (F=4 × G=3, R=1).

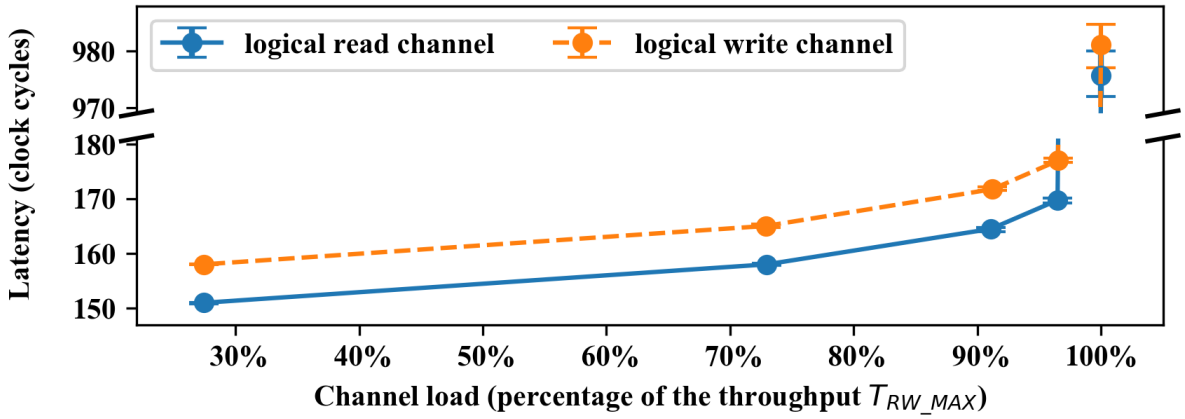

Fig.II.15. Load-latency curve for RingNet with 5 1st level rings and 3 PGs connected at each ring, and one ring used at the network root (F=5 × G=3, R=1).

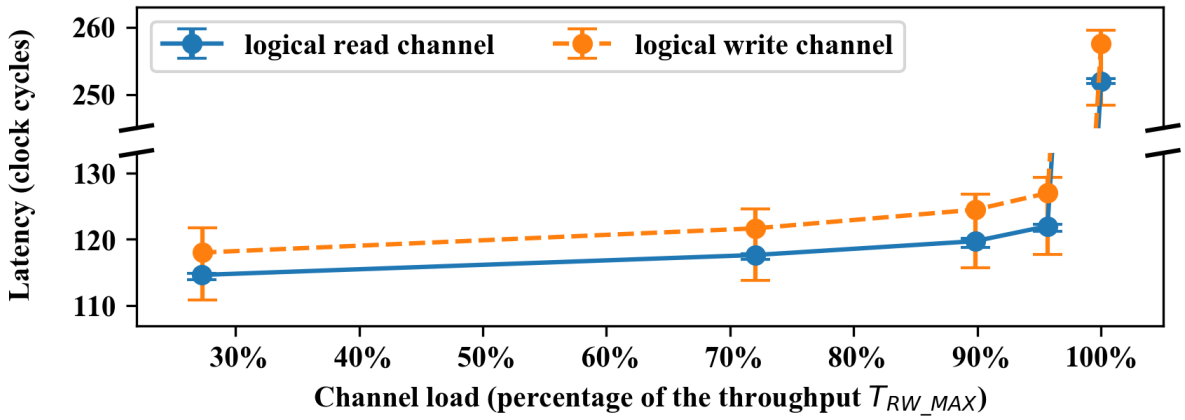

Fig.II.16. Load-latency curve for RingNet with one 1st level ring and 4 PGs connected at the ring, and one ring used at the network root (F=1 × G=4, R=1).

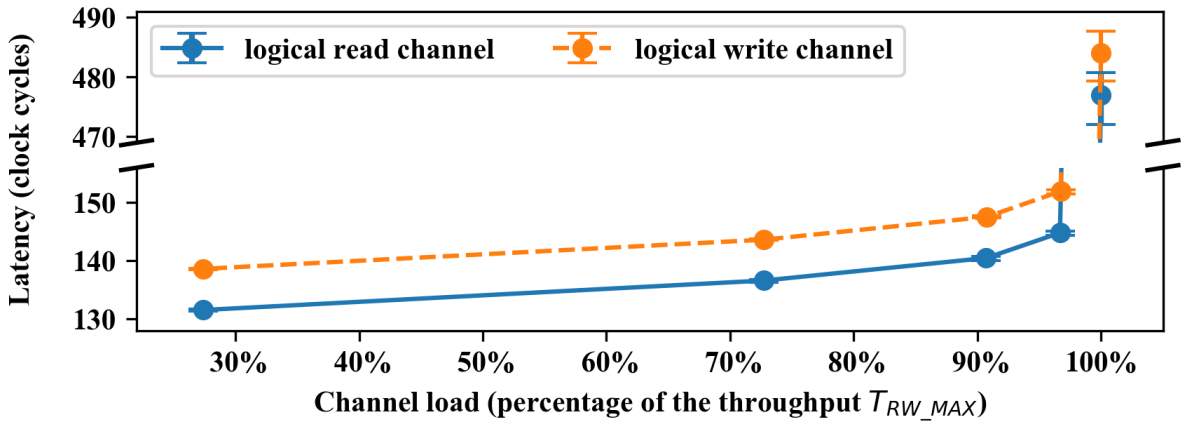

Fig.II.17. Load-latency curve for RingNet with 2 1st level rings and 4 PGs connected at each ring, and one ring used at the network root (F=2 × G=4, R=1).

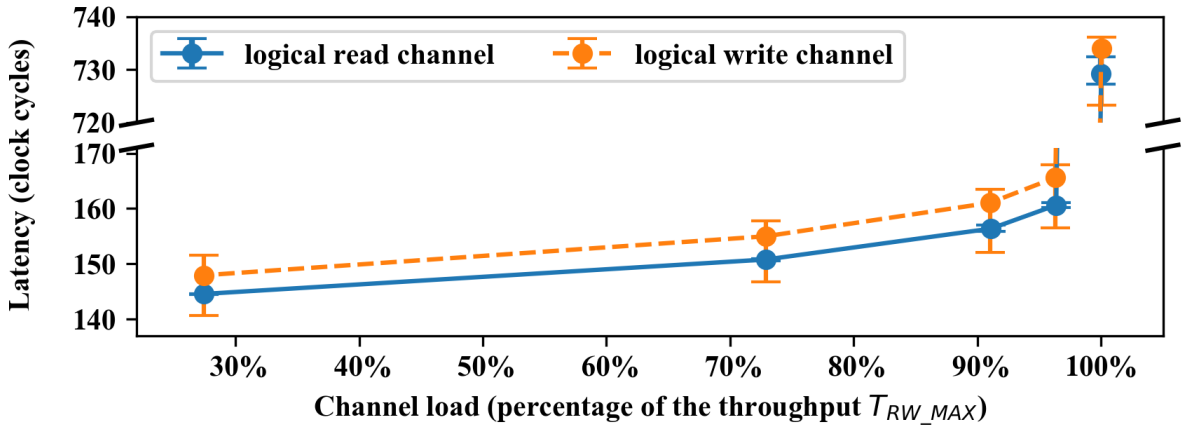

Fig.II.18. Load-latency curve for RingNet with 3 1st level rings and 4 PGs connected at each ring, and one ring used at the network root (F=3 × G=4, R=1).

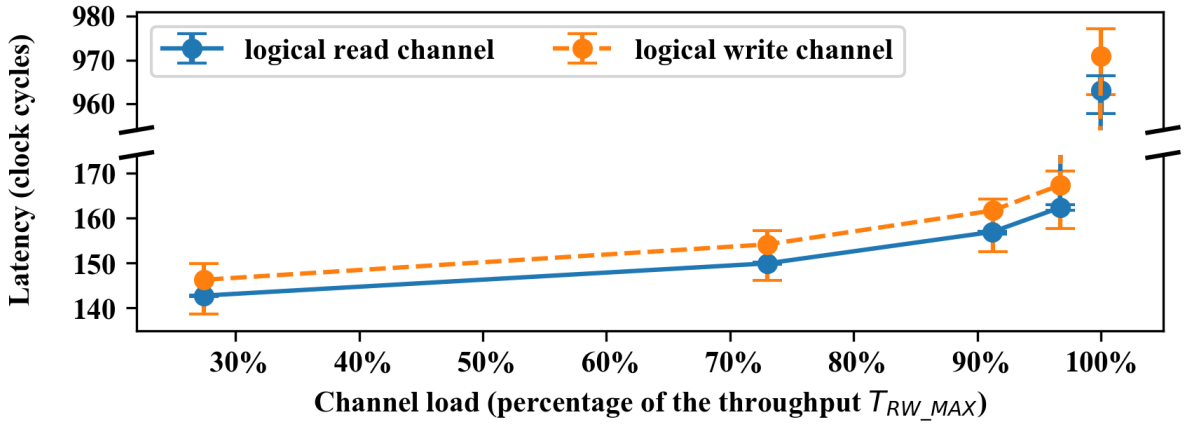

Fig.II.19. Load-latency curve for RingNet with 4 1st level rings and 4 PGs connected at each ring, and one ring used at the network root (F=4 × G=4, R=1).

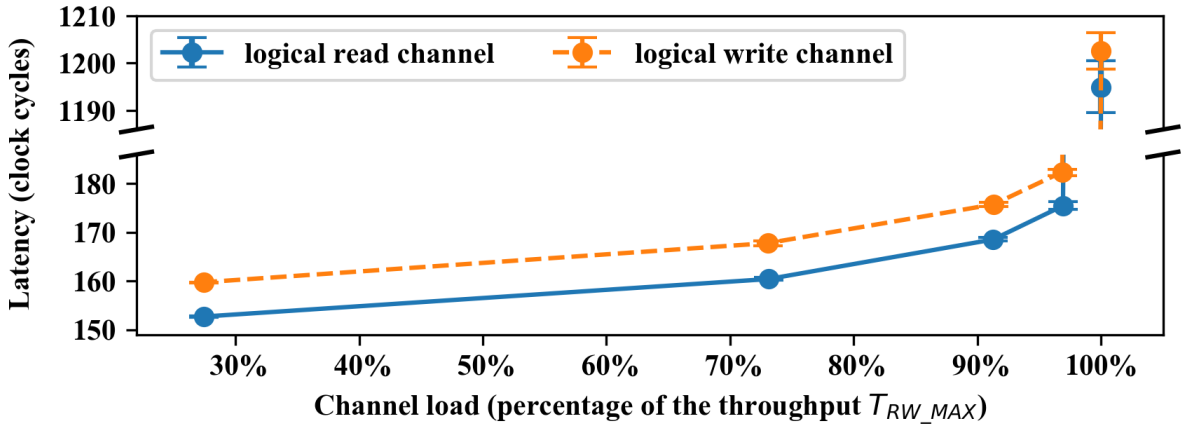

Fig.II.20. Load-latency curve for RingNet with 5 1st level rings and 4 PGs connected at each ring, and one ring used at the network root (F=5 × G=4, R=1).

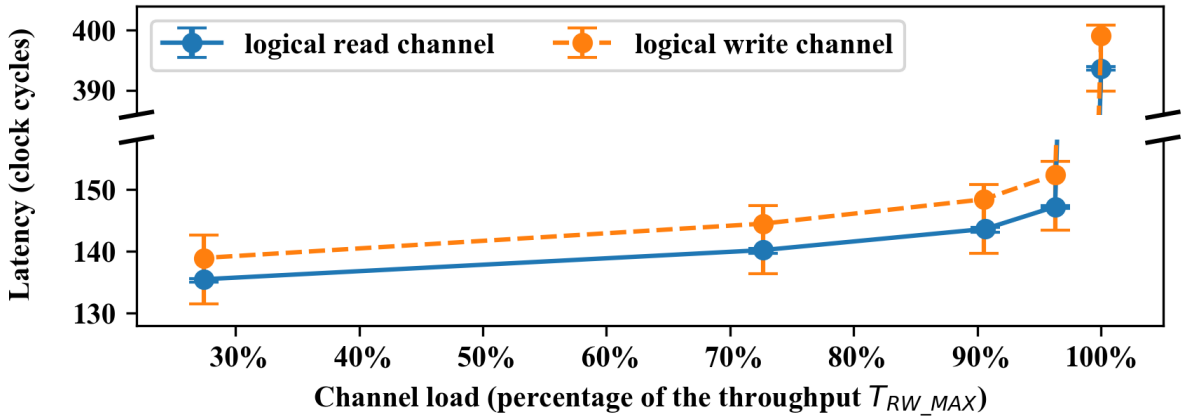

Fig.II.21. Load-latency curve for RingNet with one 1st level ring and 7 PGs connected at the ring, and one ring used at the network root (F=1 × G=7, R=1).

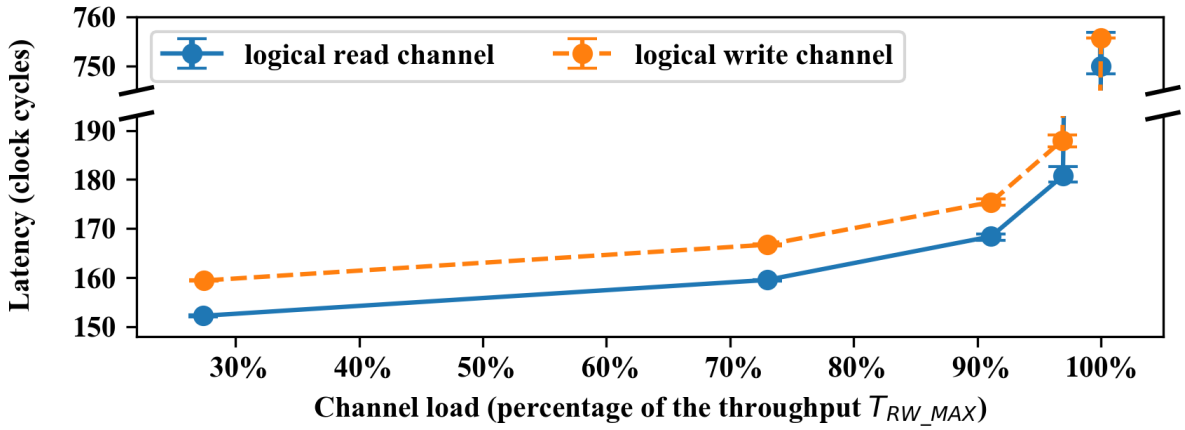

Fig.II.22. Load-latency curve for RingNet with 2 1st level rings and 7 PGs connected at each ring, and one ring used at the network root (F=2 × G=7, R=1).

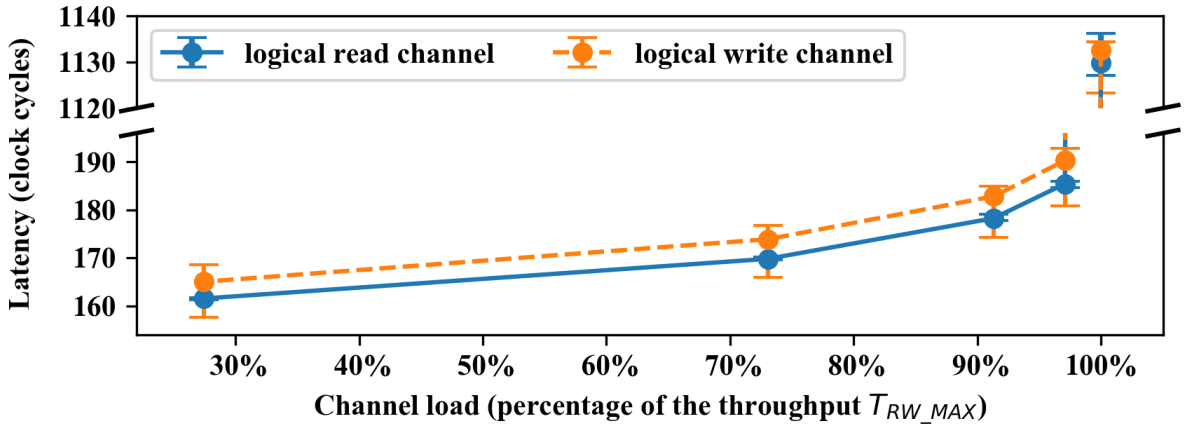

Fig.II.23. Load-latency curve for RingNet with 3 1st level rings and 7 PGs connected at each ring, and one ring used at the network root (F=3 × G=7, R=1).

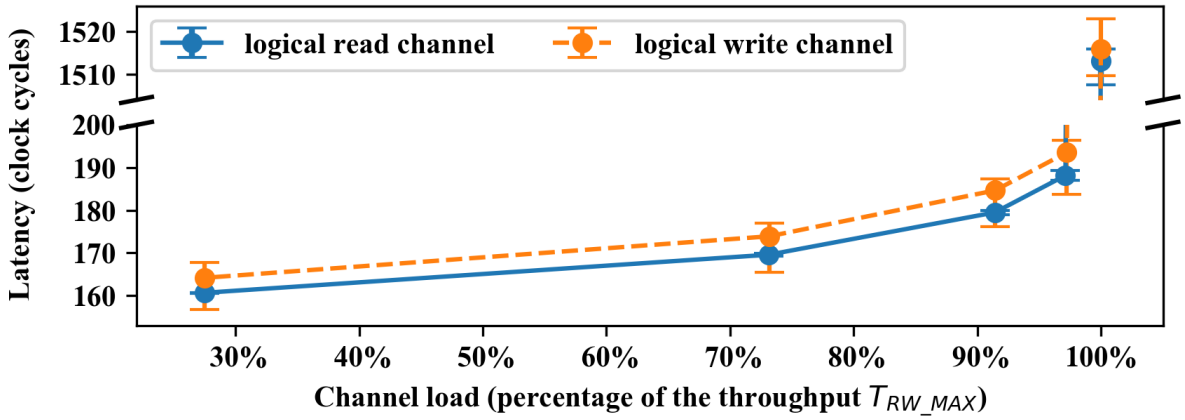

Fig.II.24. Load-latency curve for RingNet with 4 1st level rings and 7 PGs connected at each ring, and one ring used at the network root (F=4 × G=7, R=1).

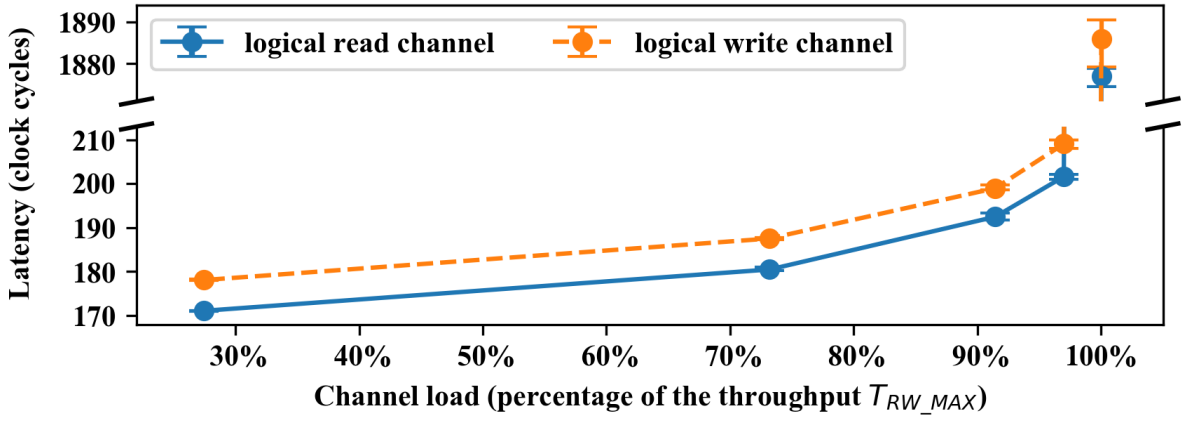

Fig.II.25. Load-latency curve for RingNet with 5 1st level rings and 7 PGs connected at each ring, and one ring used at the network root (F=5 × G=7, R=1).

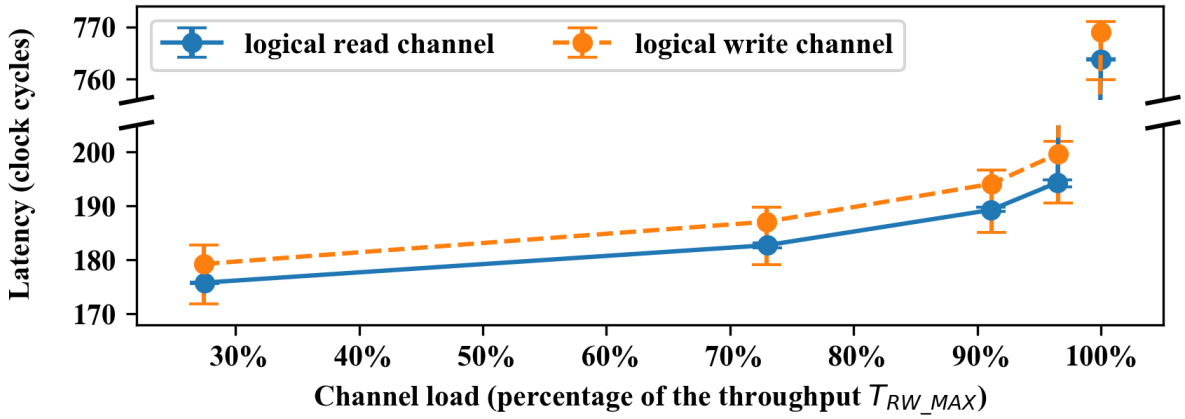

Fig.II.26. Load-latency curve for RingNet with one 1st level ring and 15 PGs connected at the ring, and one ring used at the network root (F=1 × G=15, R=1).

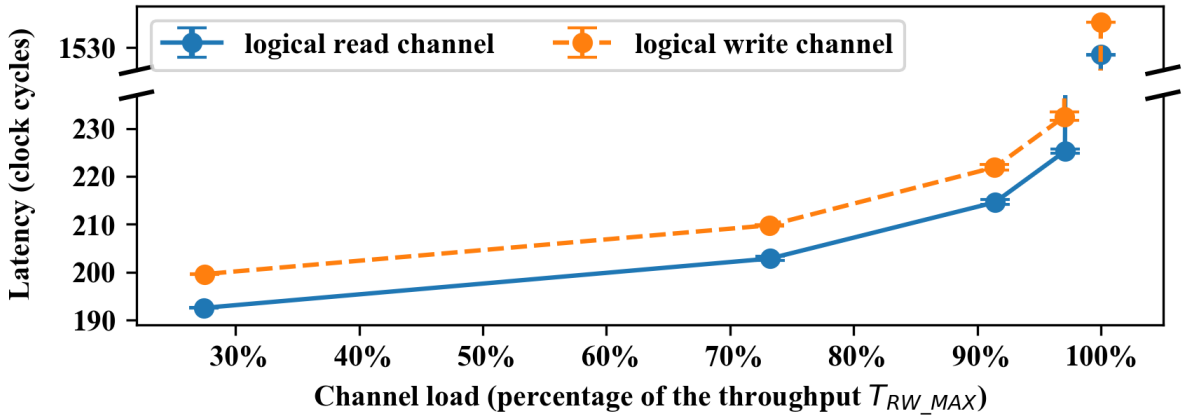

Fig.II.27. Load-latency curve for RingNet with 2 1st level rings and 15 PGs connected at each ring, and one ring used at the network root (F=2 × G=15, R=1).

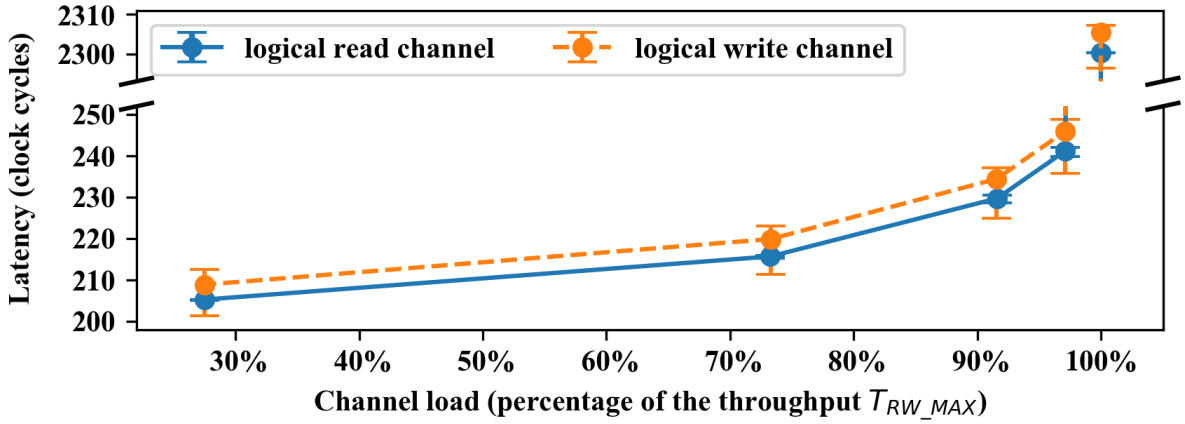

Fig.II.28. Load-latency curve for RingNet with 3 1st level rings and 15 PGs connected at each ring, and one ring used at the network root (F=3 × G=15, R=1).

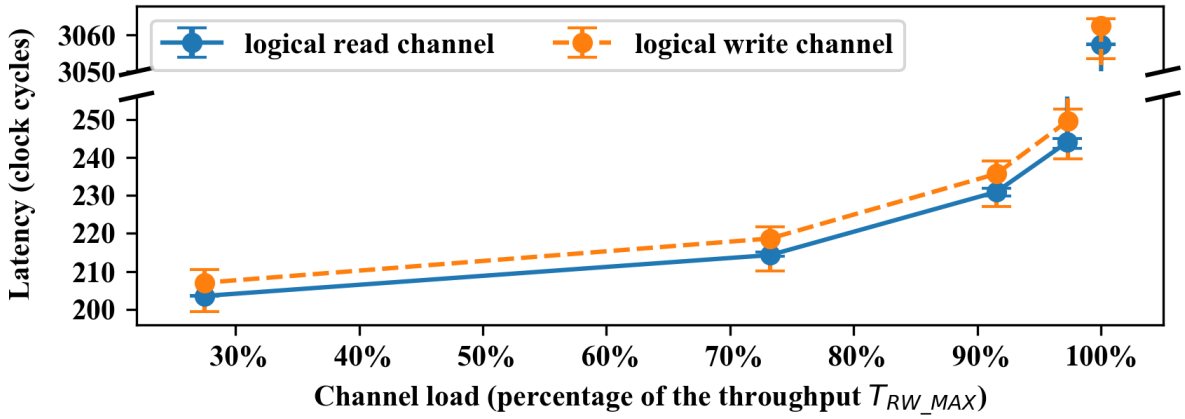

Fig.II.29. Load-latency curve for RingNet with 4 1st level rings and 15 PGs connected at each ring, and one ring used at the network root (F=4 × G=15, R=1).

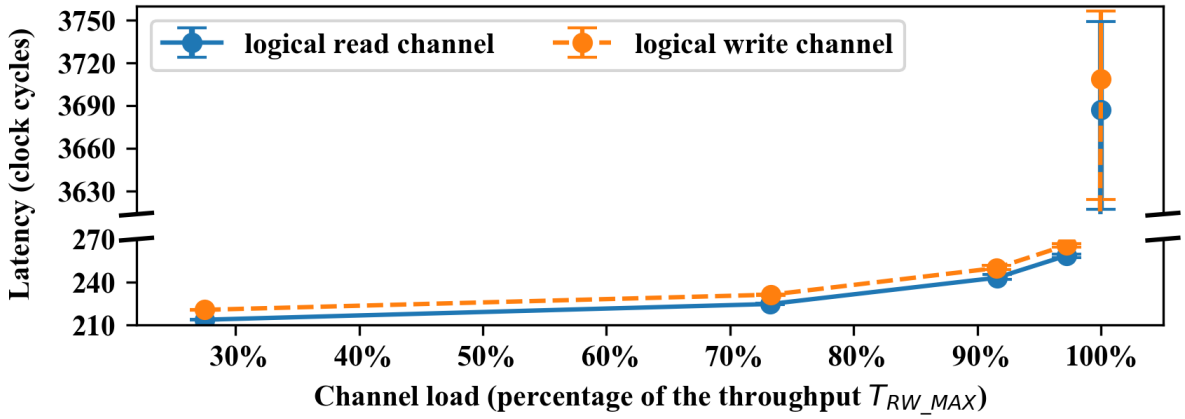

Fig.II.30. Load-latency curve for RingNet with 5 1st level rings and 15 PGs connected at each ring, and one ring used at the network root (F=5 × G=15, R=1).

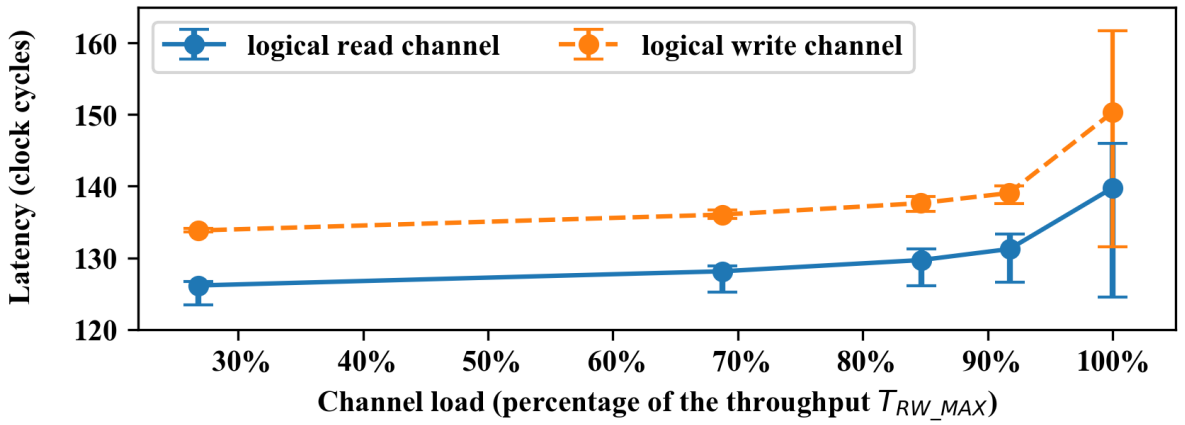

Fig.II.31. Load-latency curve for RingNet with 2 1st level rings and 1 PG connected at each ring, and 2 parallel rings used at the network root (F=2 × G=1, R=2).

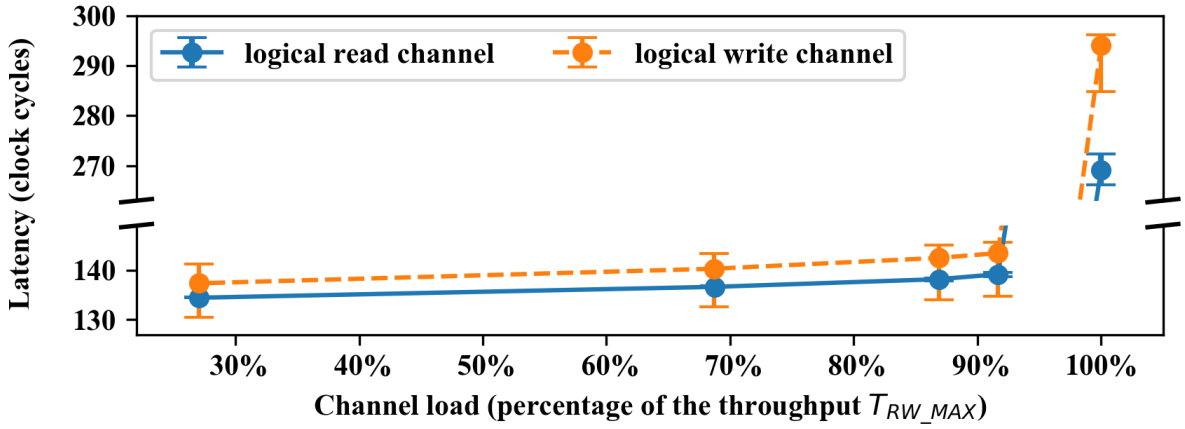

Fig.II.32. Load-latency curve for RingNet with 3 1st level rings and 1 PG connected at each ring, and 2 parallel rings used at the network root (F=3 × G=1, R=2).

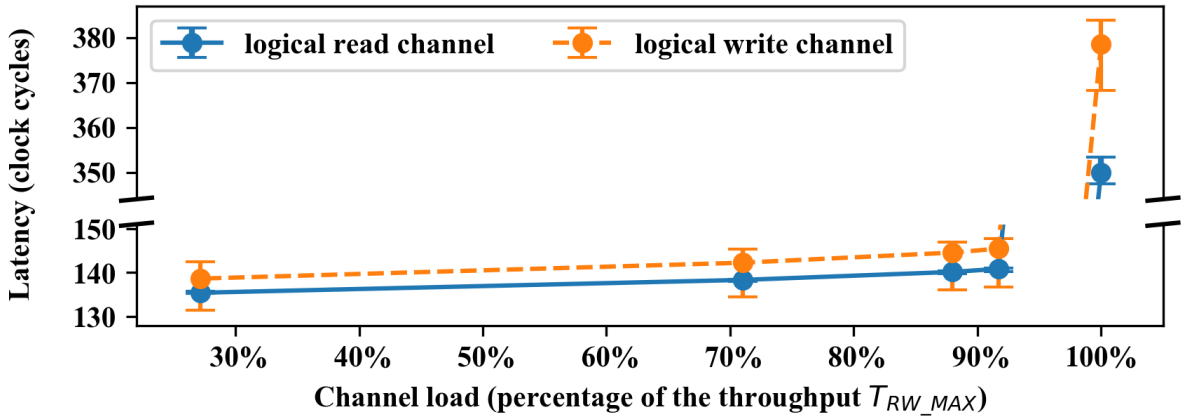

Fig.II.33. Load-latency curve for RingNet with 4 1st level rings and 1 PG connected at each ring, and 2 parallel rings used at the network root (F=4 × G=1, R=2).

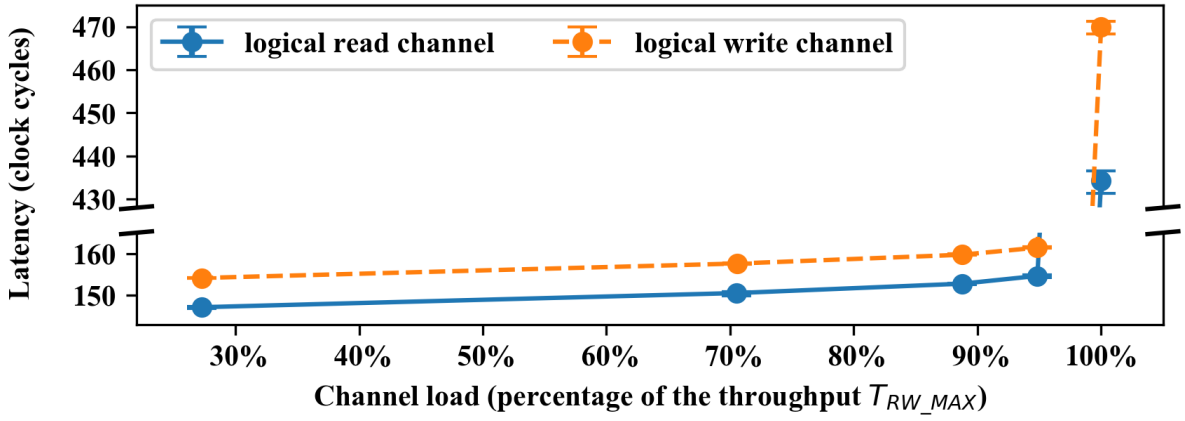

Fig.II.34. Load-latency curve for RingNet with 5 1st level rings and 1 PG connected at each ring, and 2 parallel rings used at the network root (F=5 × G=1, R=2).

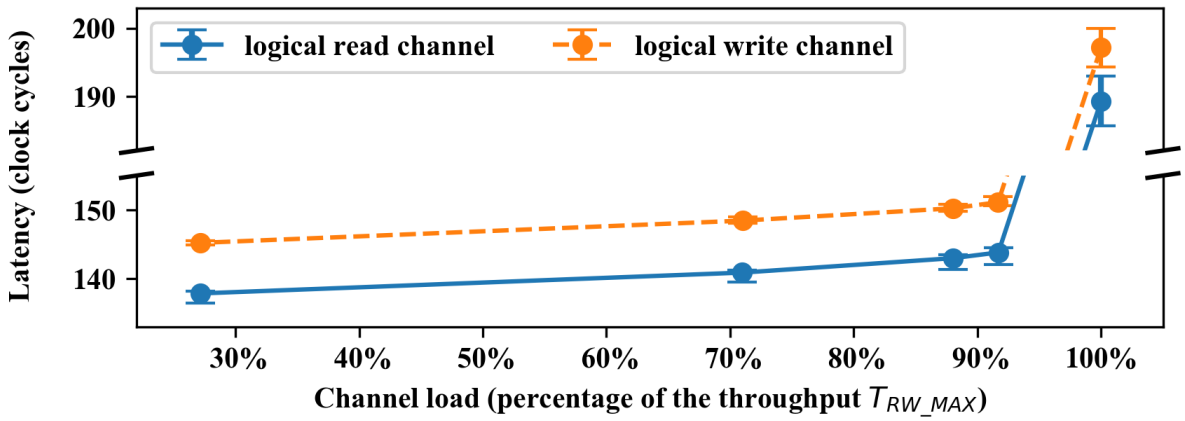

Fig.II.35. Load-latency curve for RingNet with 2 1st level rings and 2 PGs connected at each ring, and 2 parallel rings used at the network root (F=2 × G=2, R=2).

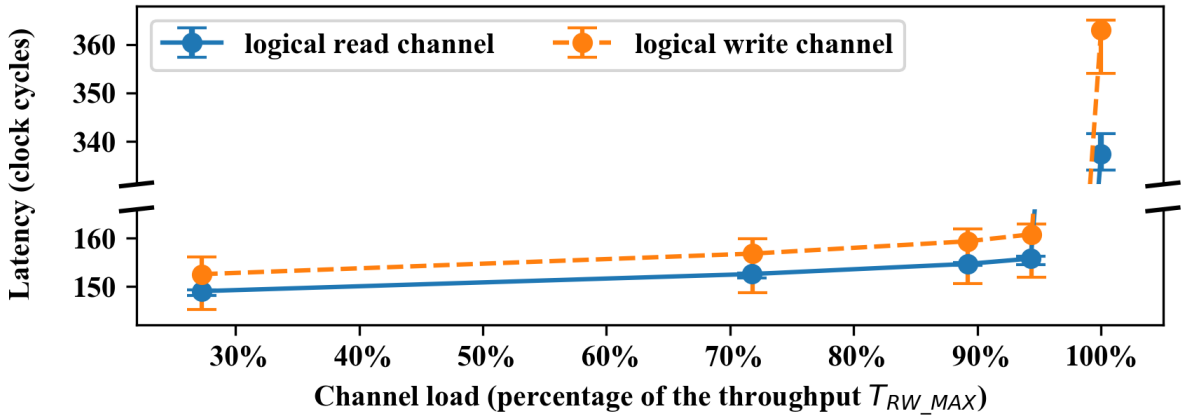

Fig.II.36. Load-latency curve for RingNet with 3 1st level rings and 2 PGs connected at each ring, and 2 parallel rings used at the network root (F=3 × G=2, R=2).

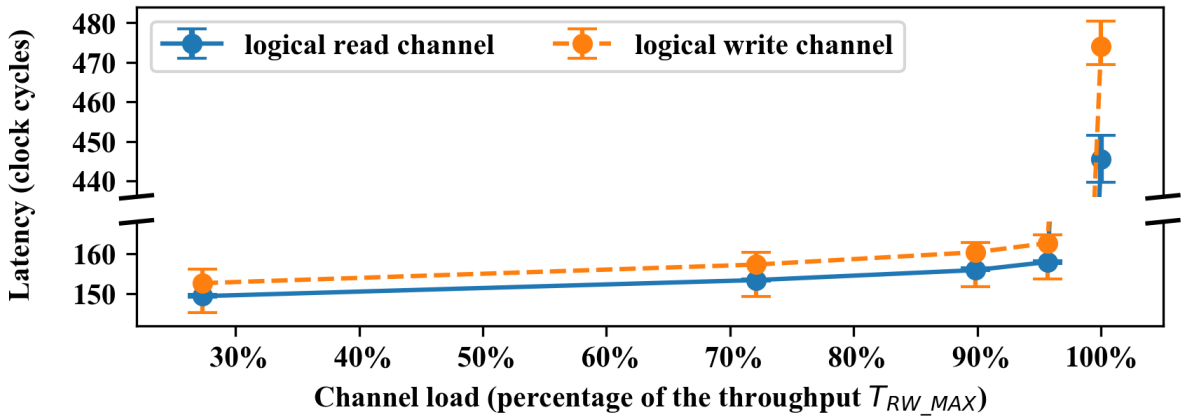

Fig.II.37. Load-latency curve for RingNet with 4 1st level rings and 2 PGs connected at each ring, and 2 parallel rings used at the network root (F=4 × G=2, R=2).

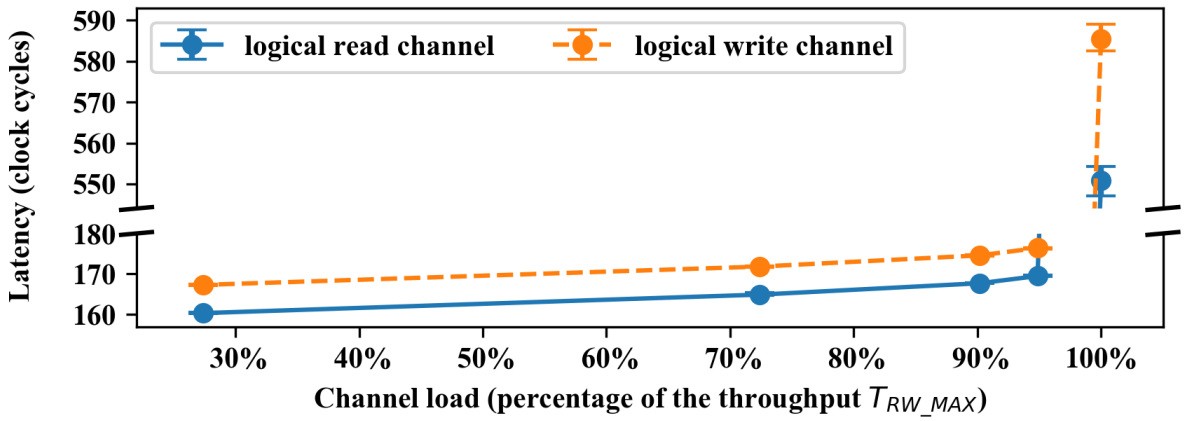

Fig.II.38. Load-latency curve for RingNet with 5 1st level rings and 2 PGs connected at each ring, and 2 parallel rings used at the network root (F=5 × G=2, R=2).

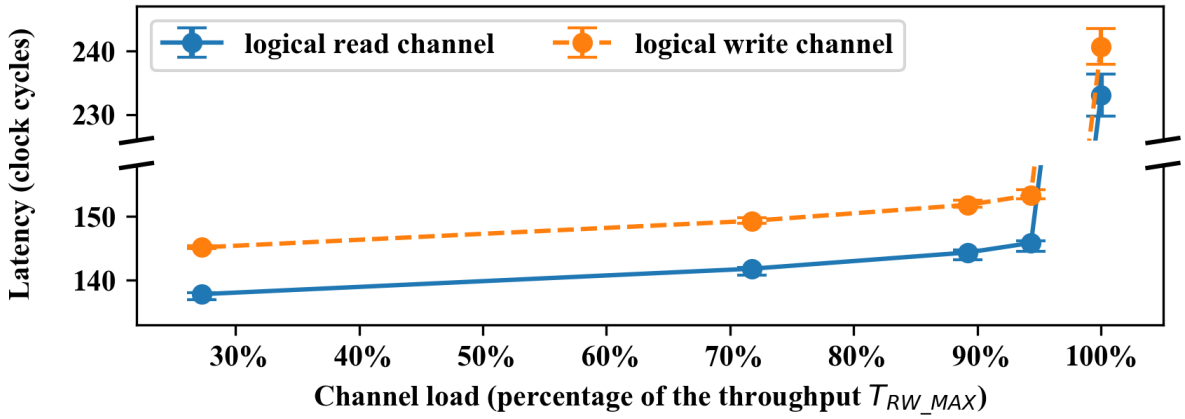

Fig.II.39. Load-latency curve for RingNet with 2 1st level rings and 3 PGs connected at each ring, and 2 parallel rings used at the network root (F=2 × G=3, R=2).

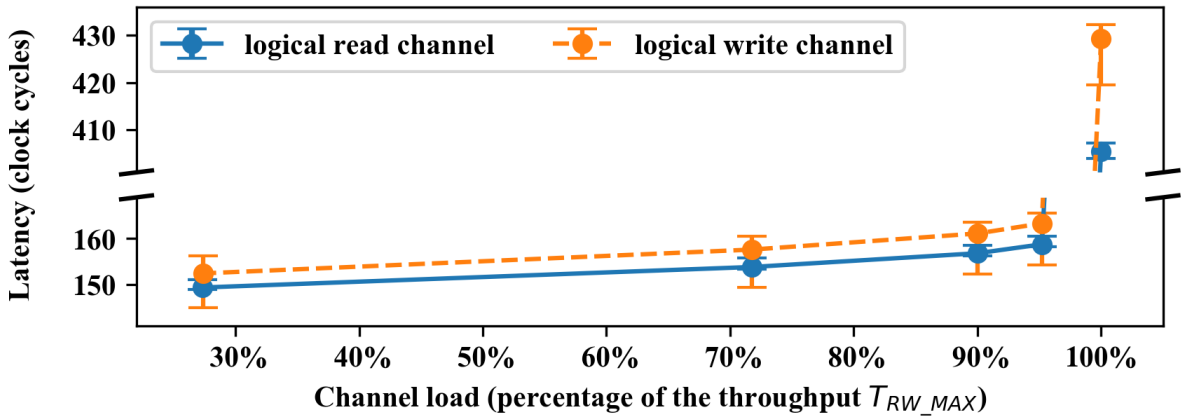

Fig.II.40. Load-latency curve for RingNet with 3 1st level rings and 3 PGs connected at each ring, and 2 parallel rings used at the network root (F=3 × G=3, R=2).

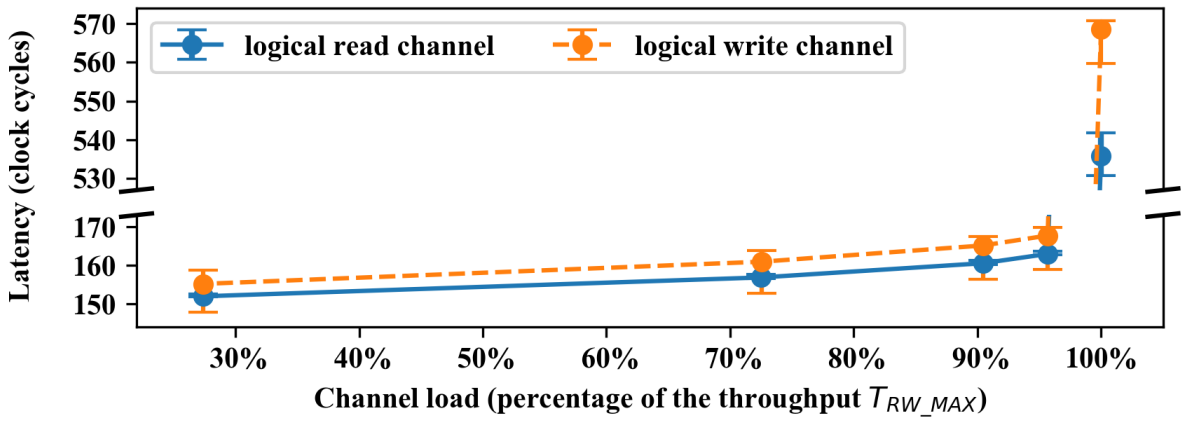

Fig.II.41. Load-latency curve for RingNet with 4 1st level rings and 3 PGs connected at each ring, and 2 parallel rings used at the network root (F=4 × G=3, R=2).

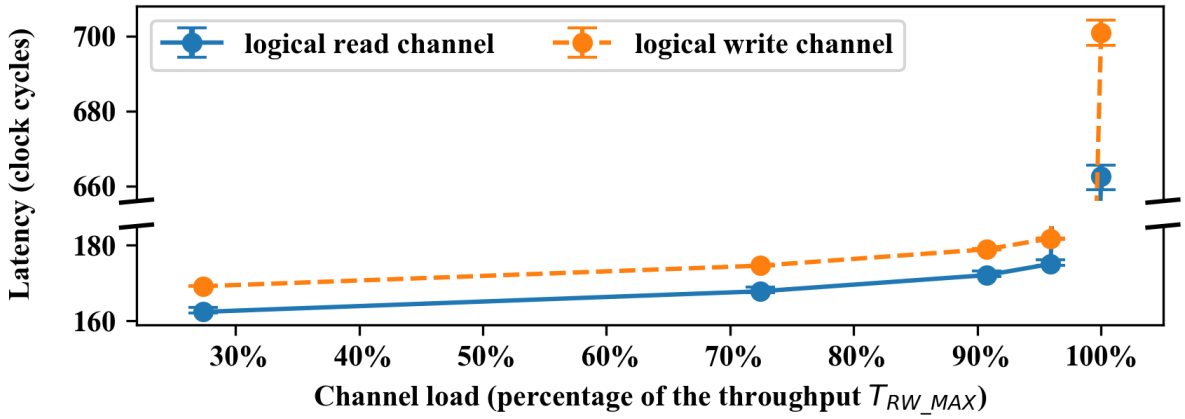

Fig.II.42. Load-latency curve for RingNet with 5 1st level rings and 3 PGs connected at each ring, and 2 parallel rings used at the network root (F=5 × G=3, R=2).

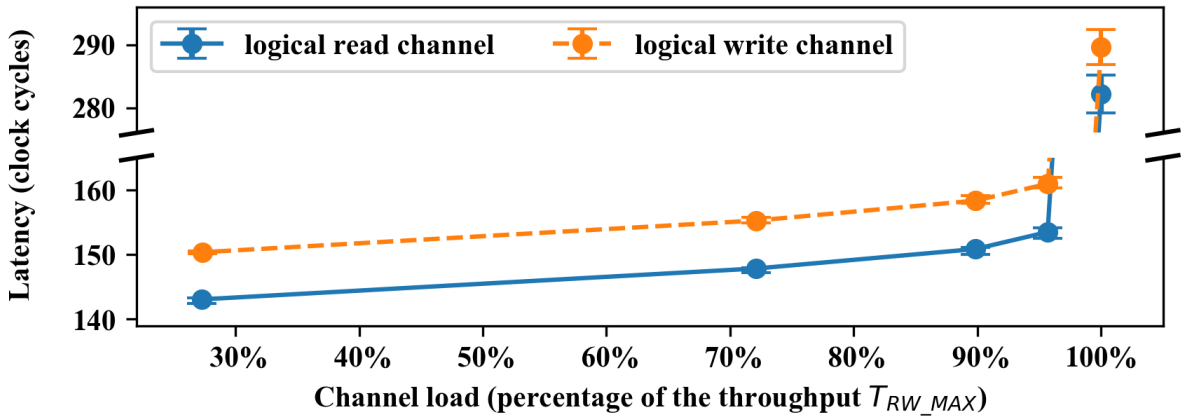

Fig.II.43. Load-latency curve for RingNet with 2 1st level rings and 4 PGs connected at each ring, and 2 parallel rings used at the network root (F=2 × G=4, R=2).

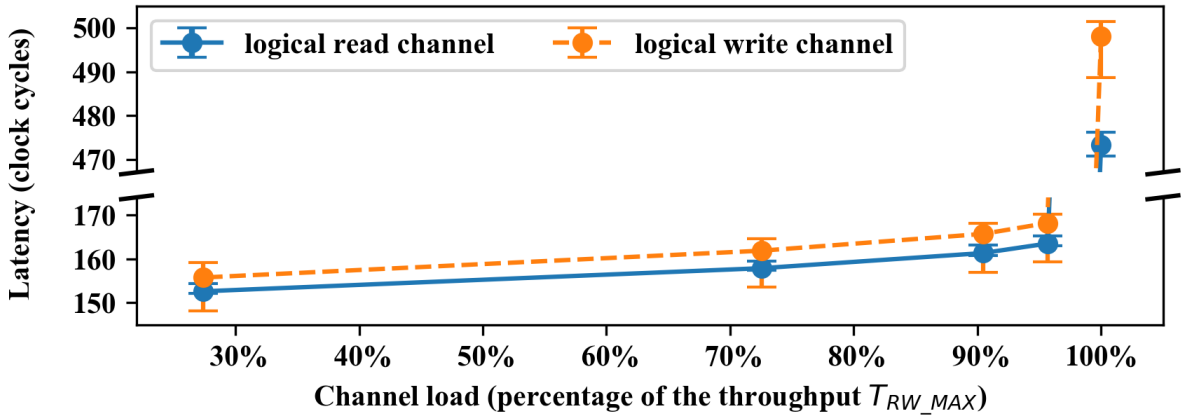

Fig.II.44. Load-latency curve for RingNet with 3 1st level rings and 4 PGs connected at each ring, and 2 parallel rings used at the network root (F=3 × G=4, R=2).

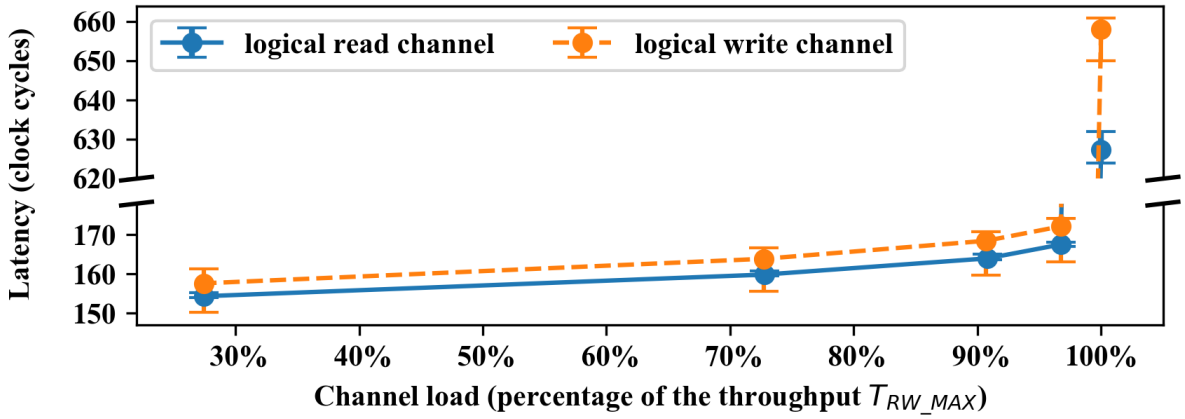

Fig.II.45. Load-latency curve for RingNet with 4 1st level rings and 4 PGs connected at each ring, and 2 parallel rings used at the network root (F=4 × G=4, R=2).

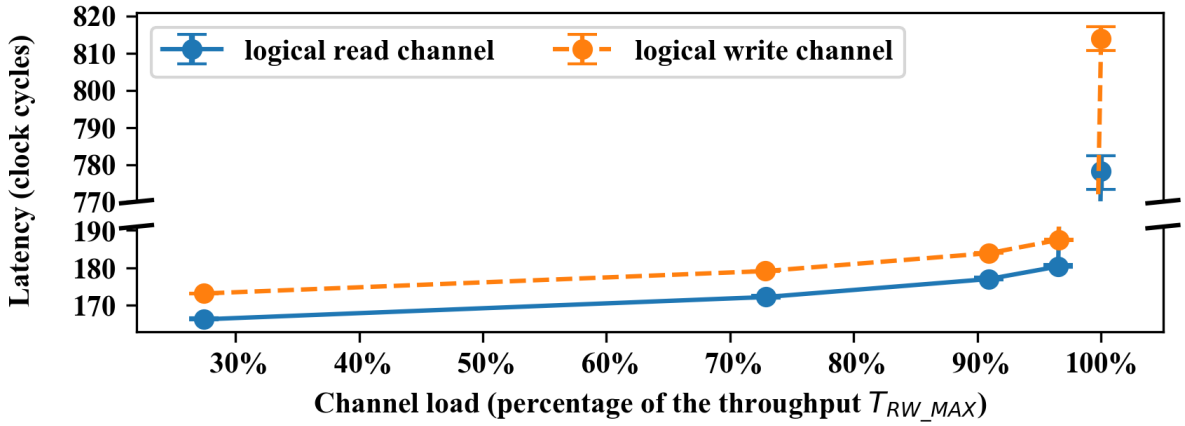

Fig.II.46. Load-latency curve for RingNet with 5 1st level rings and 4 PGs connected at each ring, and 2 parallel rings used at the network root (F=5 × G=4, R=2).

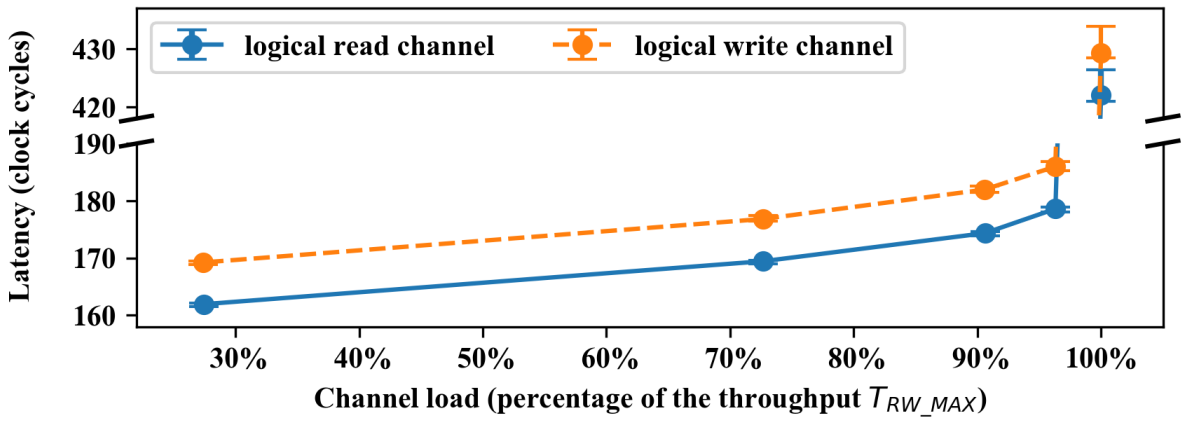

Fig.II.47. Load-latency curve for RingNet with 2 1st level rings and 7 PGs connected at each ring, and 2 parallel rings used at the network root (F=2 × G=7, R=2).

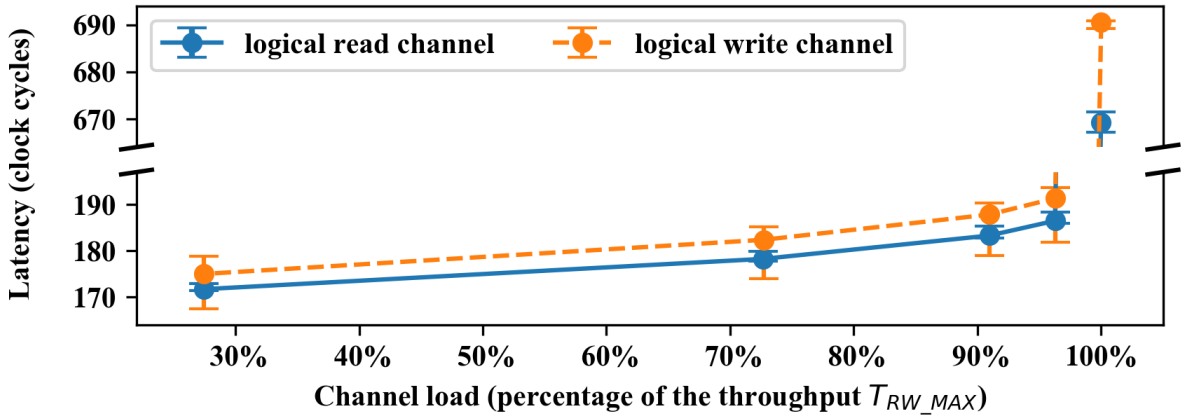

Fig.II.48. Load-latency curve for RingNet with 3 1st level rings and 7 PGs connected at each ring, and 2 parallel rings used at the network root (F=3 × G=7, R=2).

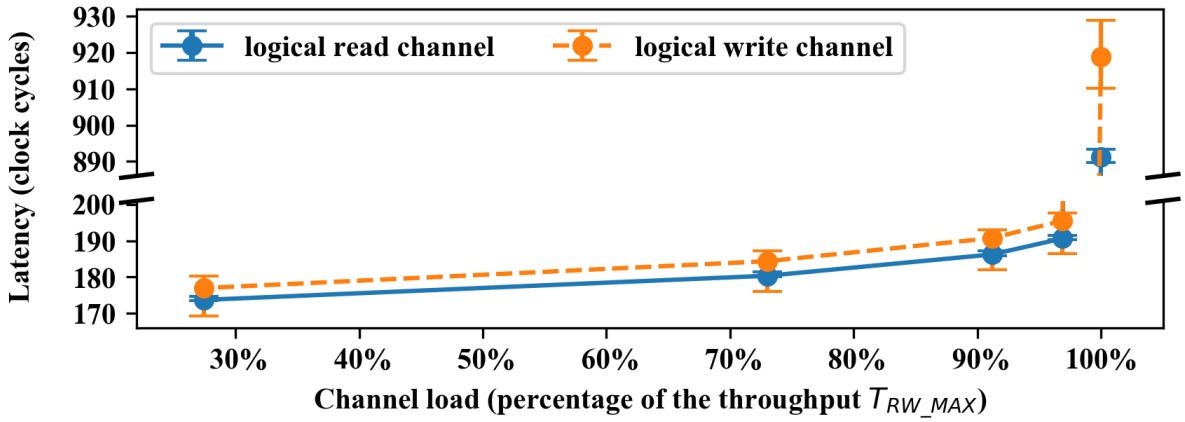

Fig.II.49. Load-latency curve for RingNet with 4 1st level rings and 7 PGs connected at each ring, and 2 parallel rings used at the network root (F=4 × G=7, R=2).

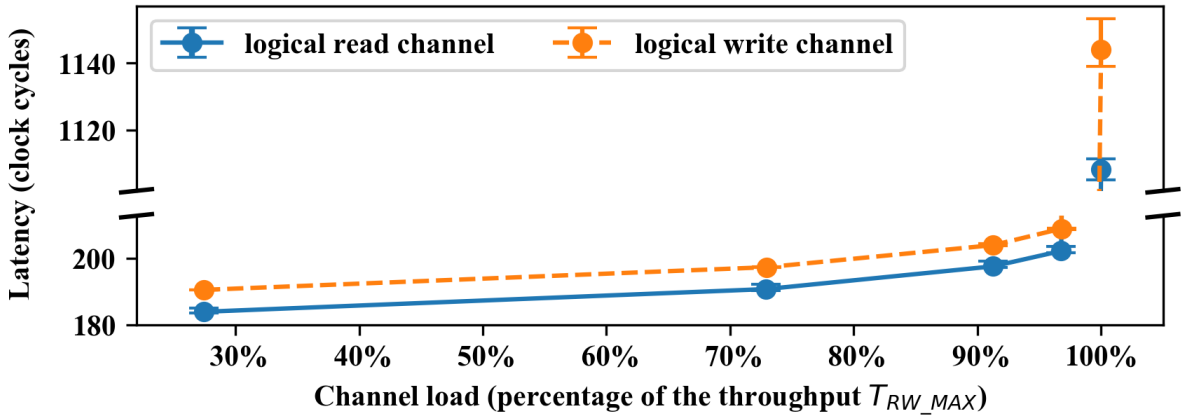

Fig.II.50. Load-latency curve for RingNet with 5 1st level rings and 7 PGs connected at each ring, and 2 parallel rings used at the network root (F=5 × G=7, R=2).

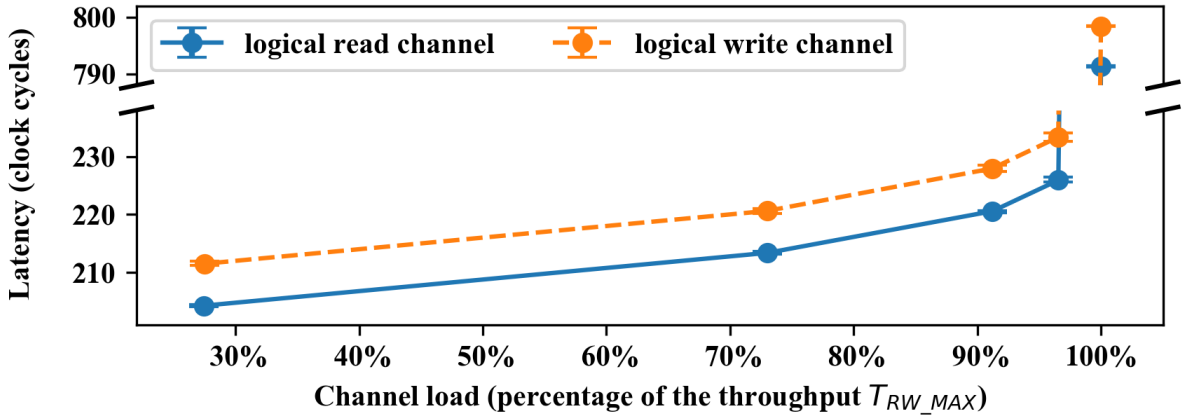

Fig.II.51. Load-latency curve for RingNet with 2 1st level rings and 15 PGs connected at each ring, and 2 parallel rings used at the network root (F=2 × G=15, R=2).

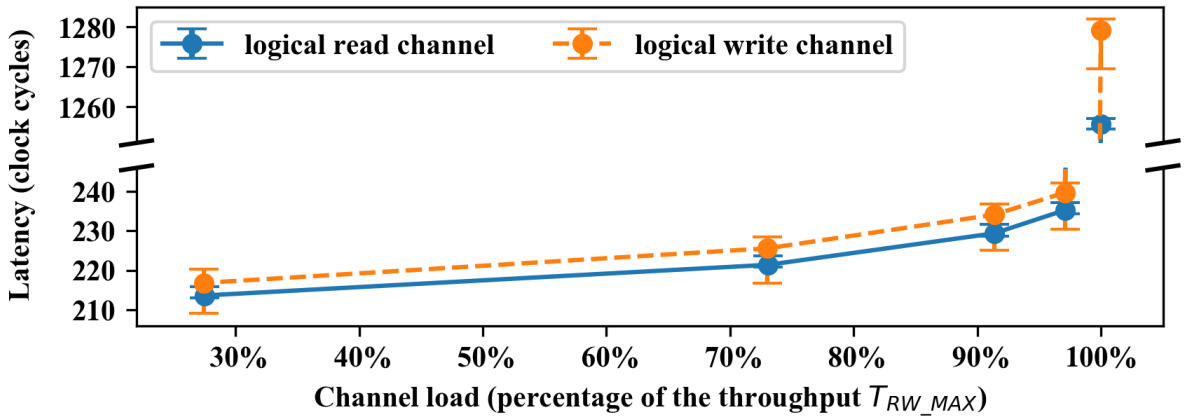

Fig.II.52. Load-latency curve for RingNet with 3 1st level rings and 15 PGs connected at each ring, and 2 parallel rings used at the network root (F=3 × G=15, R=2).

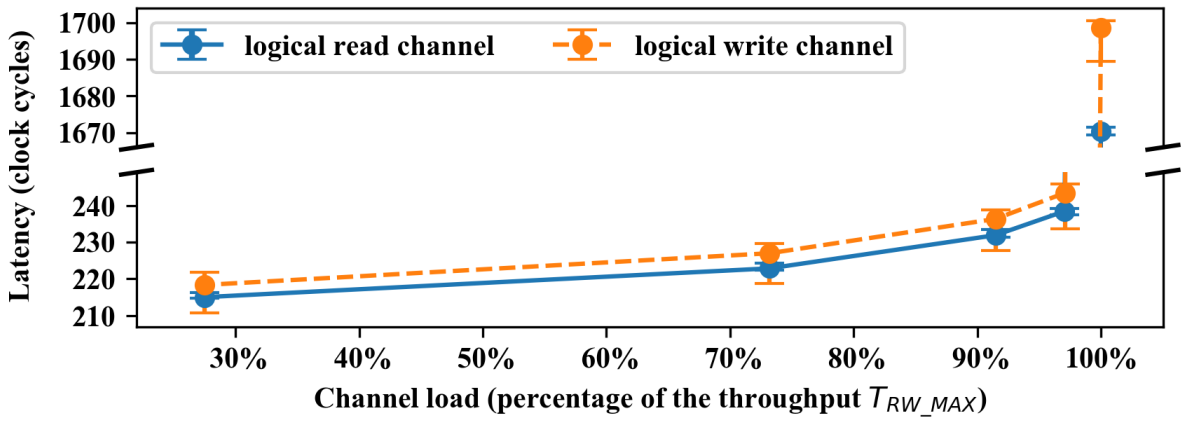

Fig.II.53. Load-latency curve for RingNet with 4 1st level rings and 15 PGs connected at each ring, and 2 parallel rings used at the network root (F=4 × G=15, R=2).

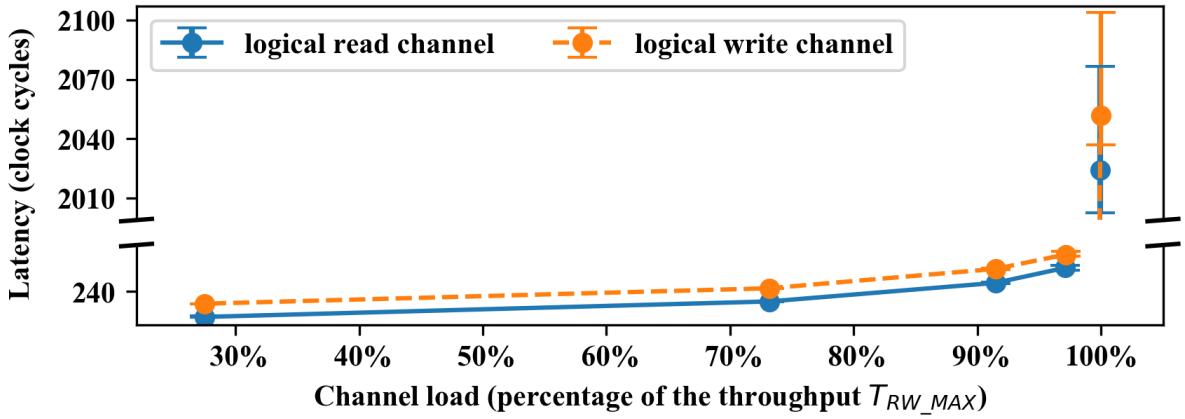

Fig.II.54. Load-latency curve for RingNet with 5 1st level rings and 15 PGs connected at each ring, and 2 parallel rings used at the network root (F=5 × G=15, R=2).

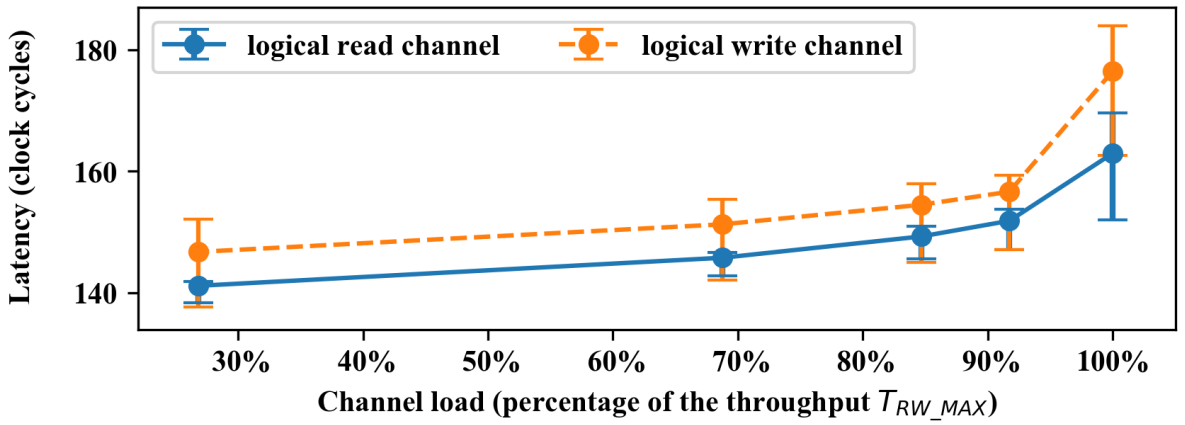

Fig.II.55. Load-latency curve for RingNet with 3 1st level rings and 1 PG connected at each ring, and 3 parallel rings used at the network root (F=3 × G=1, R=3).

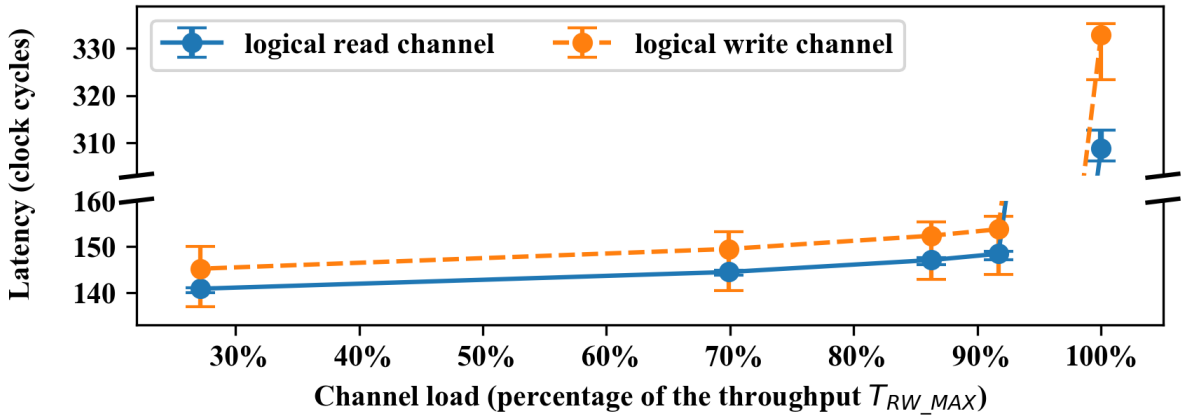

Fig.II.56. Load-latency curve for RingNet with 4 1st level rings and 1 PG connected at each ring, and 3 parallel rings used at the network root (F=4 × G=1, R=3).

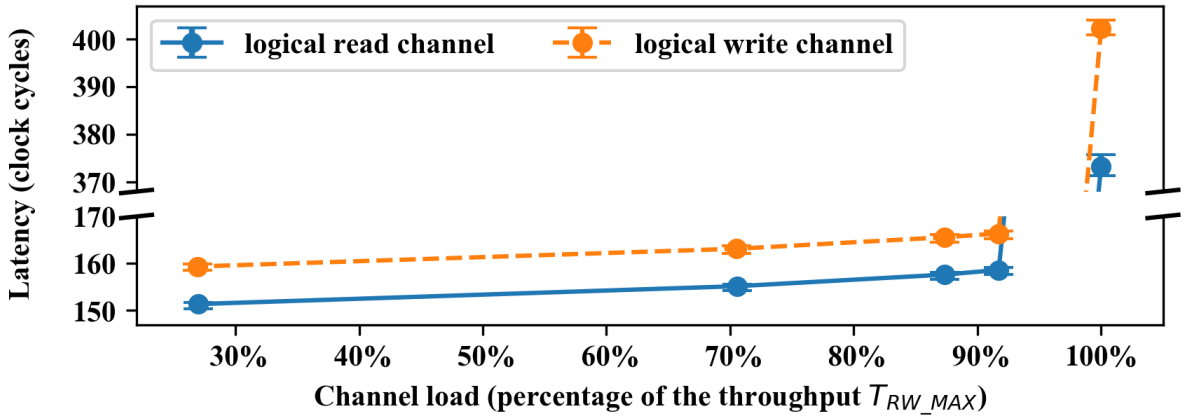

Fig.II.57. Load-latency curve for RingNet with 5 1st level rings and 1 PG connected at each ring, and 3 parallel rings used at the network root (F=5 × G=1, R=3).

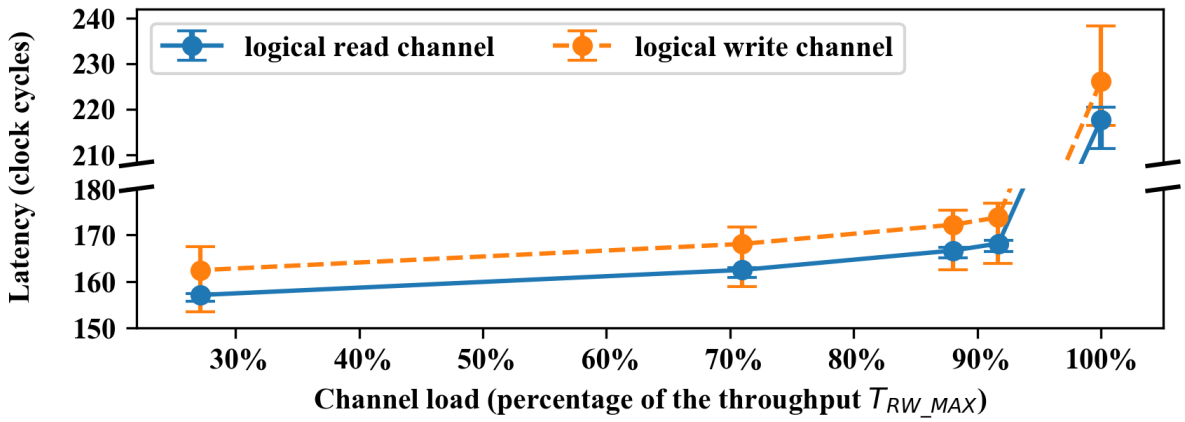

Fig.II.58. Load-latency curve for RingNet with 3 1st level rings and 2 PGs connected at each ring, and 3 parallel rings used at the network root (F=3 × G=2, R=3).

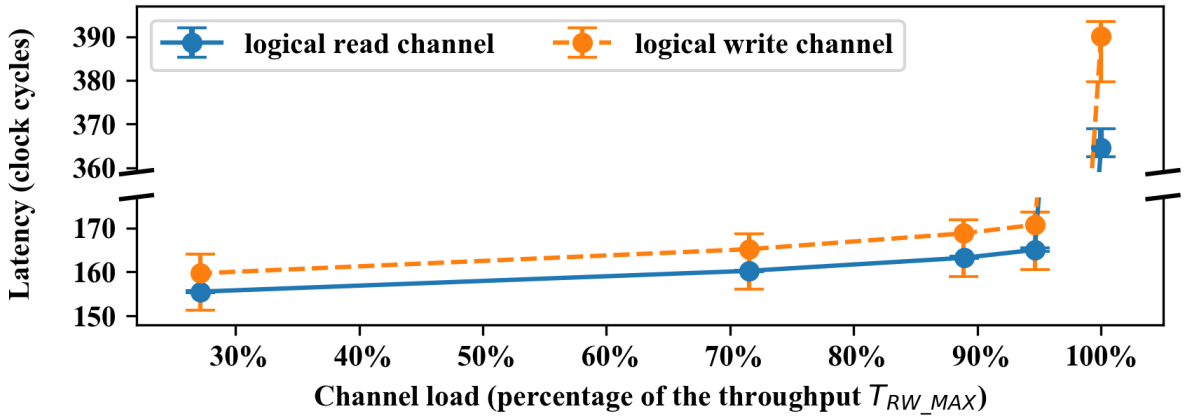

Fig.II.59. Load-latency curve for RingNet with 4 1st level rings and 2 PGs connected at each ring, and 3 parallel rings used at the network root (F=4 × G=2, R=3).

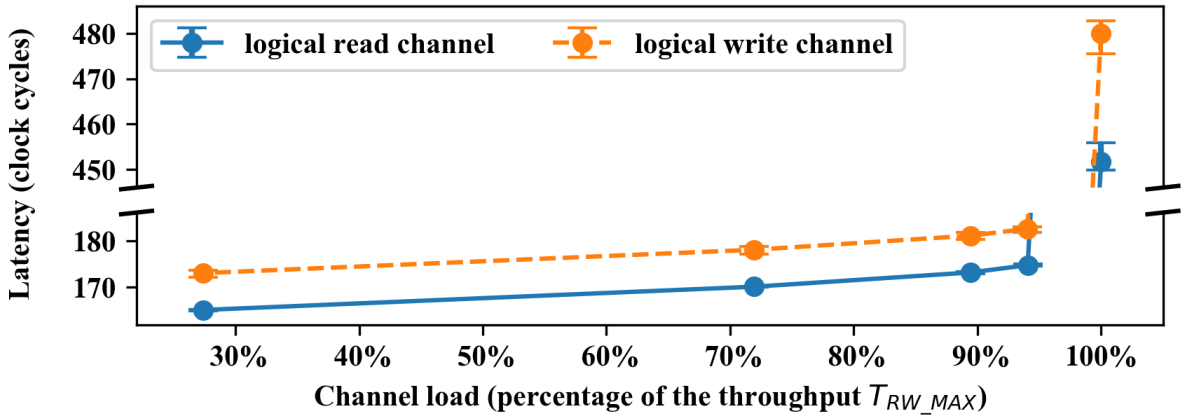

Fig.II.60. Load-latency curve for RingNet with 5 1st level rings and 2 PGs connected at each ring, and 3 parallel rings used at the network root (F=5 × G=2, R=3).

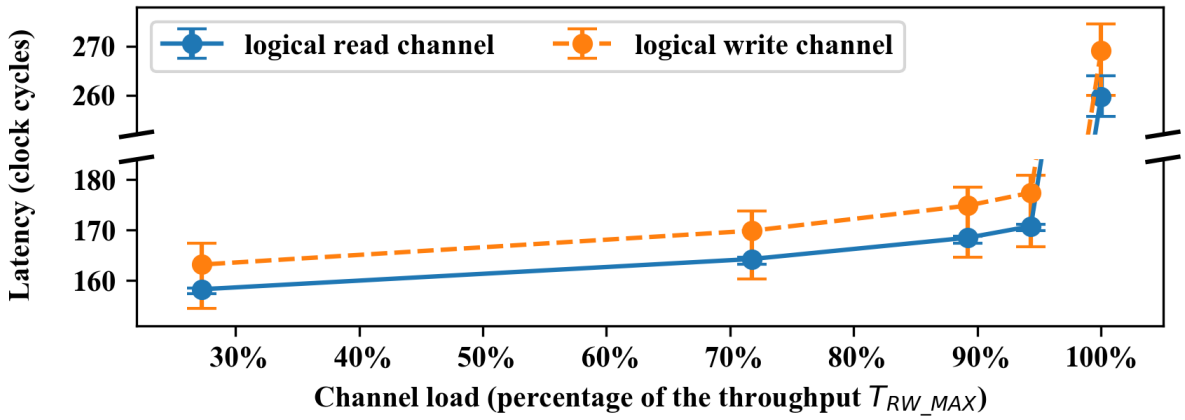

Fig.II.61. Load-latency curve for RingNet with 3 1st level rings and 3 PGs connected at each ring, and 3 parallel rings used at the network root (F=3 × G=3, R=3).

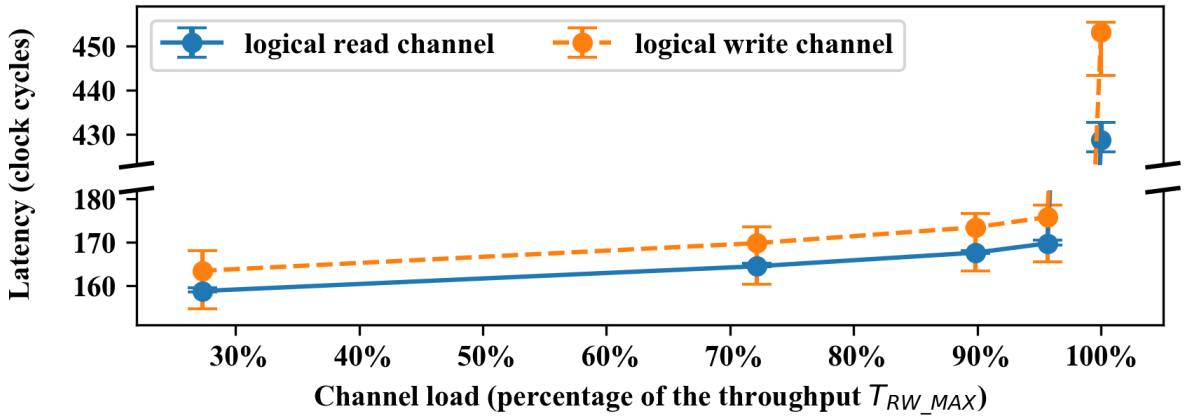

Fig.II.62. Load-latency curve for RingNet with 4 1st level rings and 3 PGs connected at each ring, and 3 parallel rings used at the network root (F=4 × G=3, R=3).

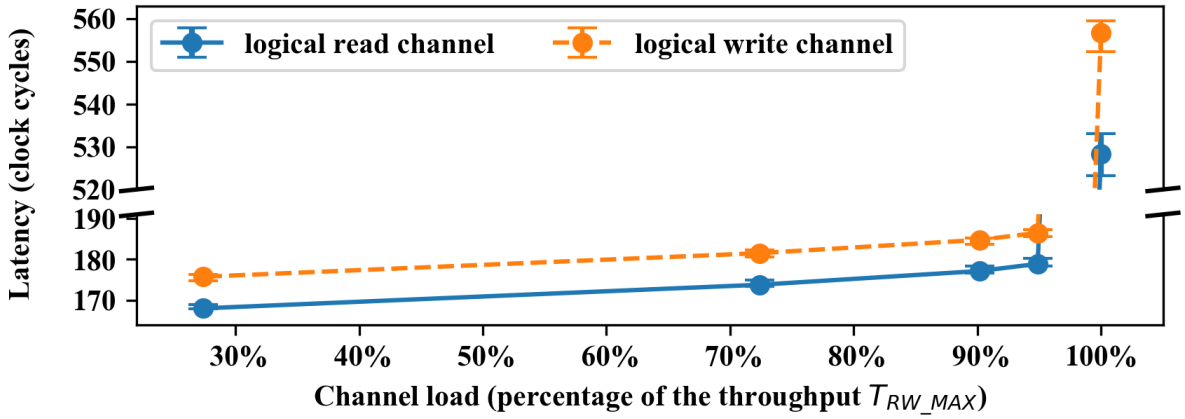

Fig.II.63. Load-latency curve for RingNet with 5 1st level rings and 3 PGs connected at each ring, and 3 parallel rings used at the network root (F=5 × G=3, R=3).

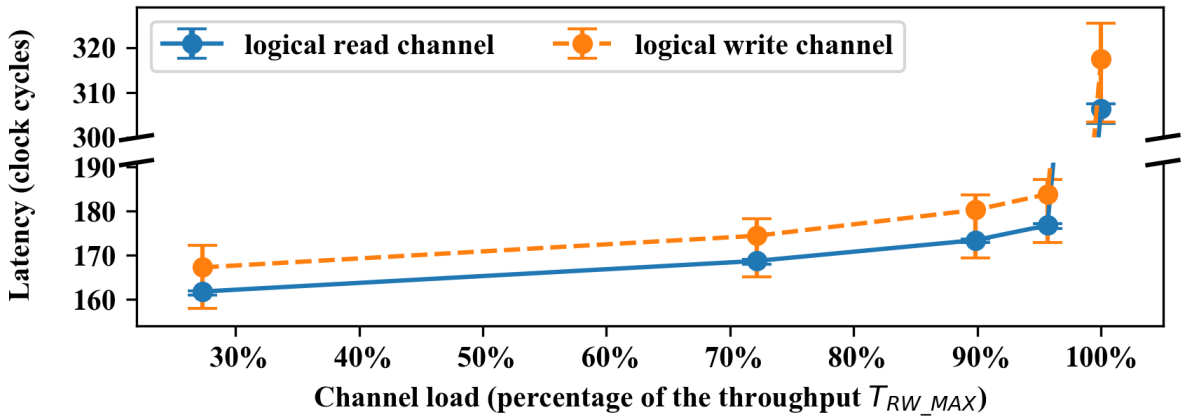

Fig.II.64. Load-latency curve for RingNet with 3 1st level rings and 4 PGs connected at each ring, and 3 parallel rings used at the network root (F=3 × G=4, R=3).

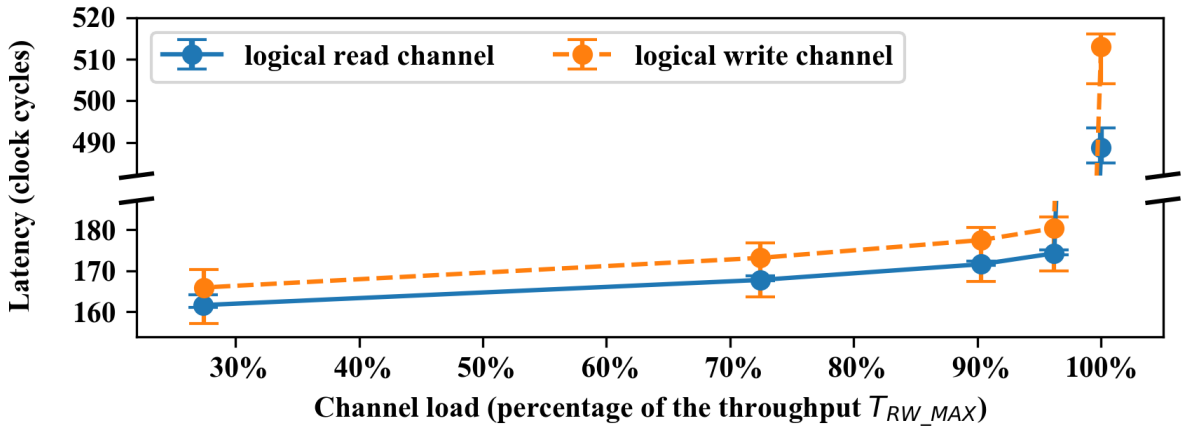

Fig.II.65. Load-latency curve for RingNet with 4 1st level rings and 4 PGs connected at each ring, and 3 parallel rings used at the network root (F=4 × G=4, R=3).

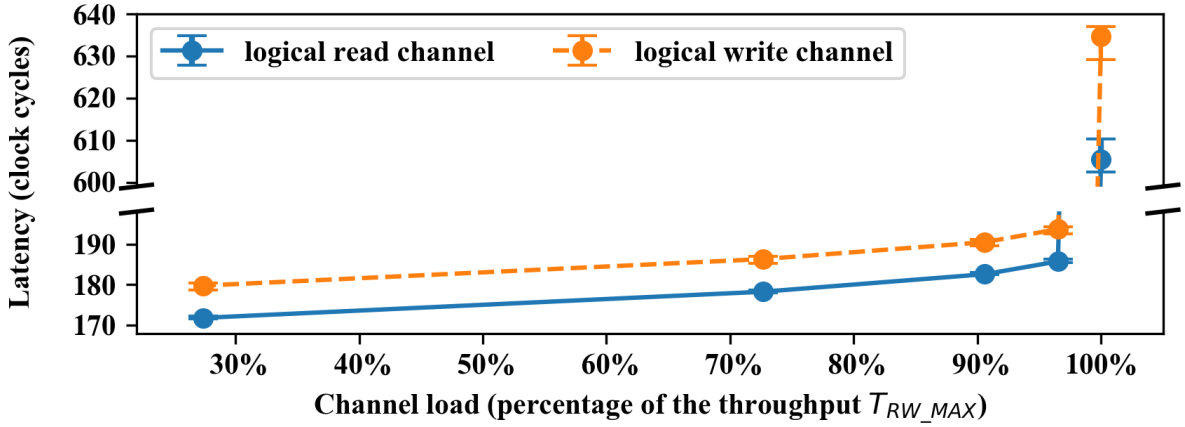

Fig.II.66. Load-latency curve for RingNet with 5 1st level rings and 4 PGs connected at each ring, and 3 parallel rings used at the network root (F=5 × G=4, R=3).

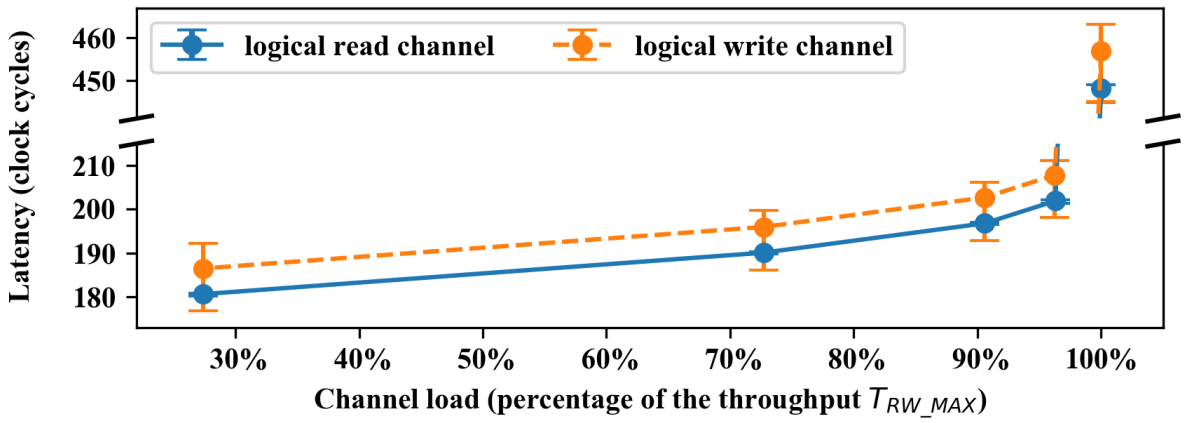

Fig.II.67. Load-latency curve for RingNet with 3 1st level rings and 7 PGs connected at each ring, and 3 parallel rings used at the network root (F=3 × G=7, R=3).

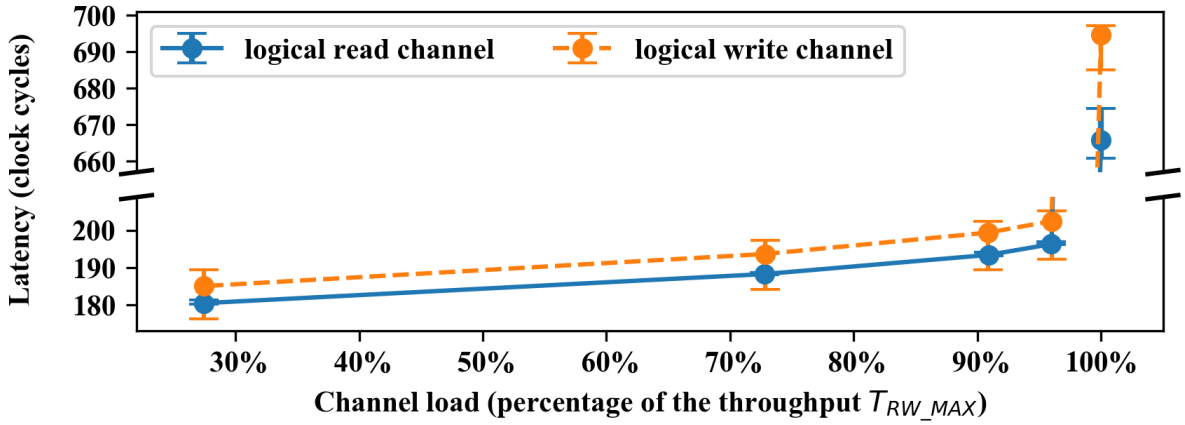

Fig.II.68. Load-latency curve for RingNet with 4 1st level rings and 7 PGs connected at each ring, and 3 parallel rings used at the network root (F=4 × G=7, R=3).

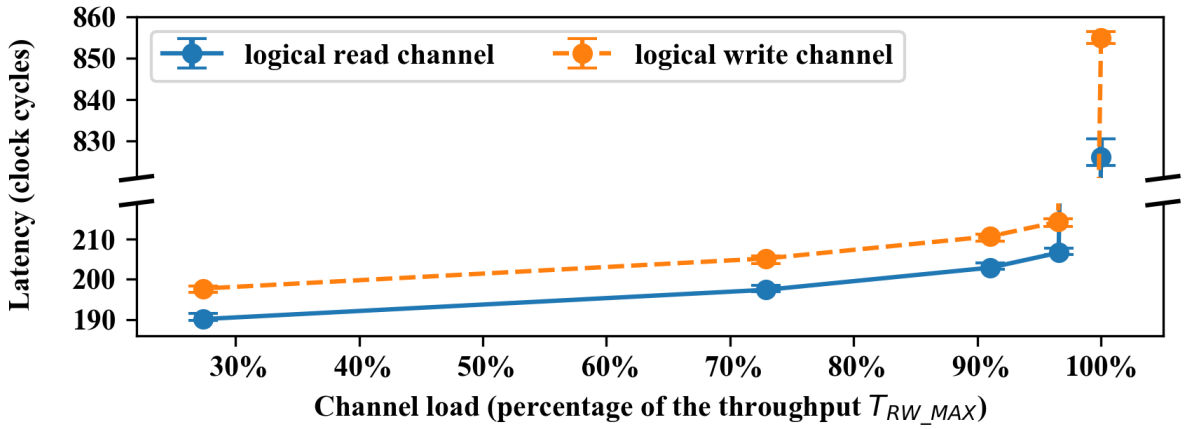

Fig.II.69. Load-latency curve for RingNet with 5 1st level rings and 7 PGs connected at each ring, and 3 parallel rings used at the network root (F=5 × G=7, R=3).

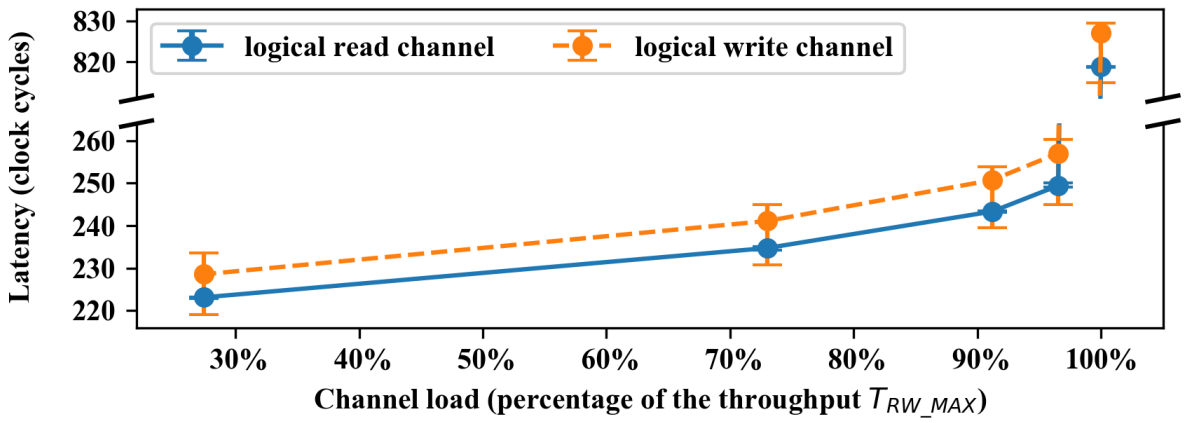

Fig.II.70. Load-latency curve for RingNet with 3 1st level rings and 15 PGs connected at each ring, and 3 parallel rings used at the network root (F=3 × G=15, R=3).

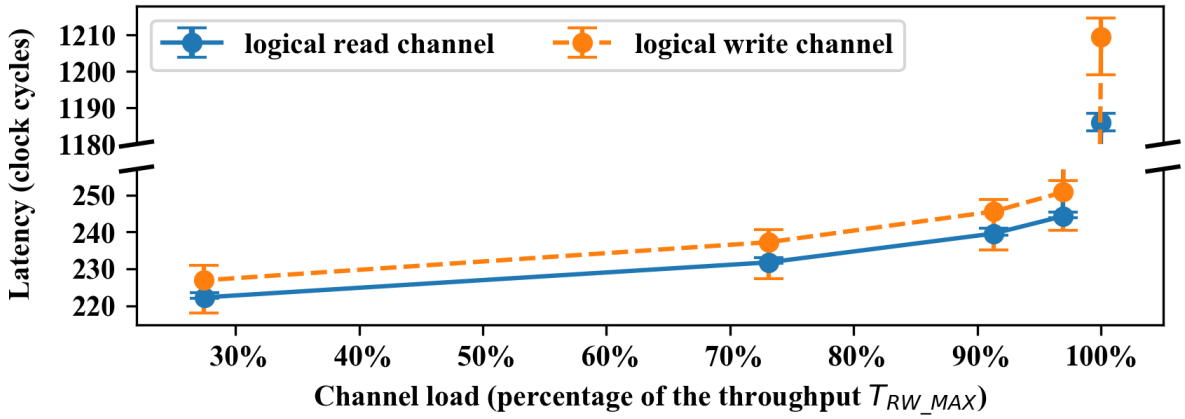

Fig.II.71. Load-latency curve for RingNet with 4 1st level rings and 15 PGs connected at each ring, and 3 parallel rings used at the network root (F=4 × G=15, R=3).

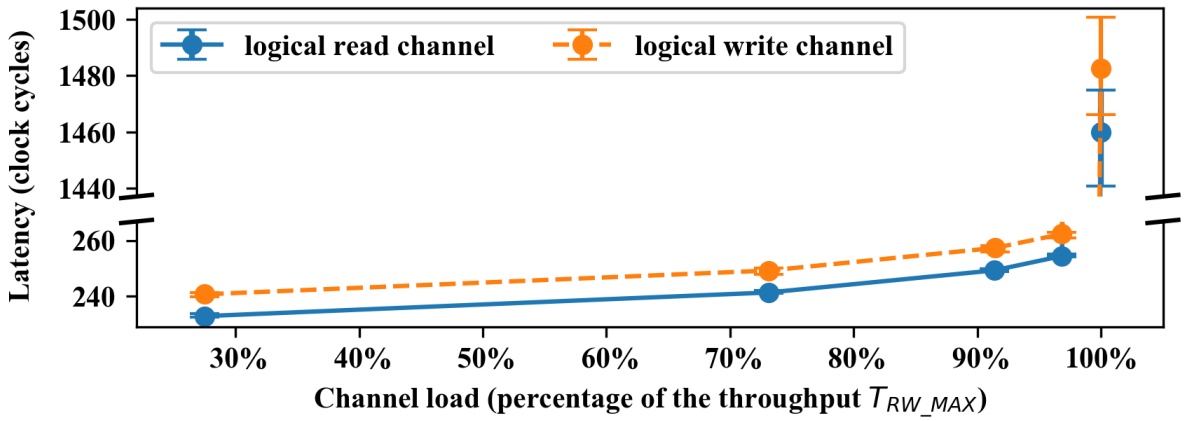

Fig.II.72. Load-latency curve for RingNet with 5 1st level rings and 15 PGs connected at each ring, and 3 parallel rings used at the network root (F=5 × G=15, R=3).

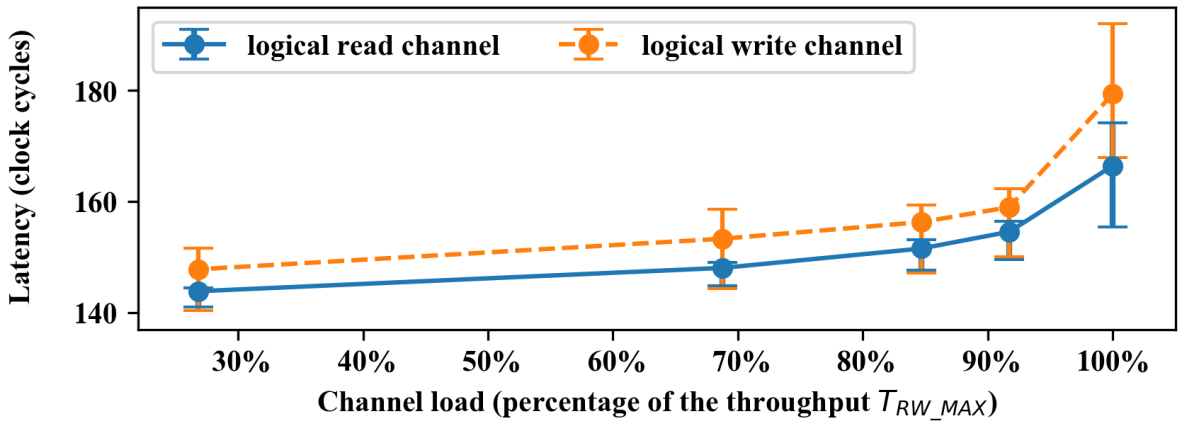

Fig.II.73. Load-latency curve for RingNet with 4 1st level rings and 1 PG connected at each ring, and 4 parallel rings used at the network root (F=4 × G=1, R=4).

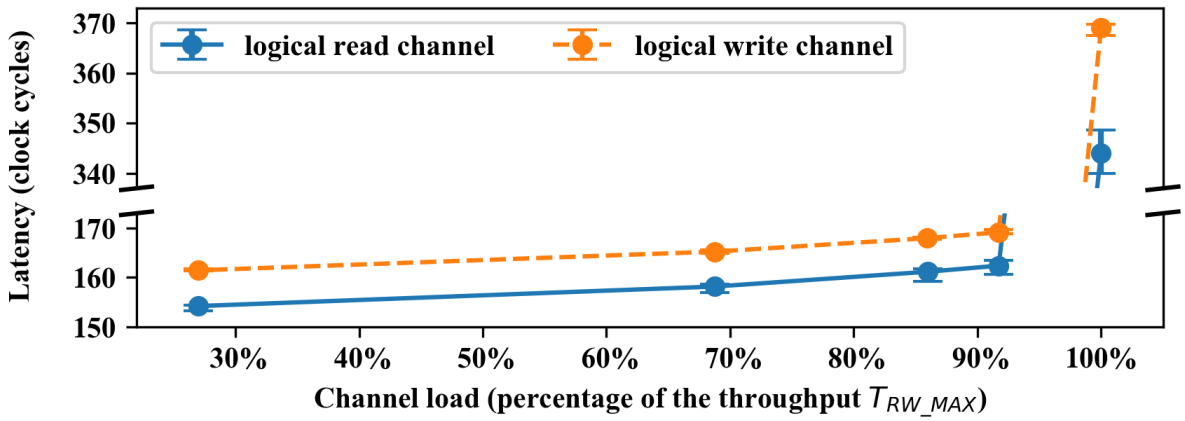

Fig.II.74. Load-latency curve for RingNet with 5 1st level rings and 1 PG connected at each ring, and 4 parallel rings used at the network root (F=5 × G=1, R=4).

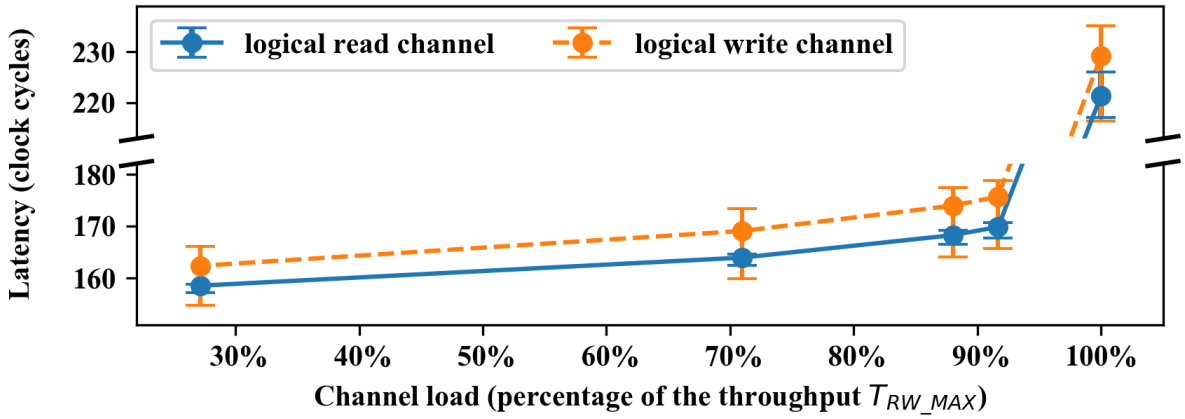

Fig.II.75. Load-latency curve for RingNet with 4 1st level rings and 2 PGs connected at each ring, and 4 parallel rings used at the network root (F=4 × G=2, R=4).

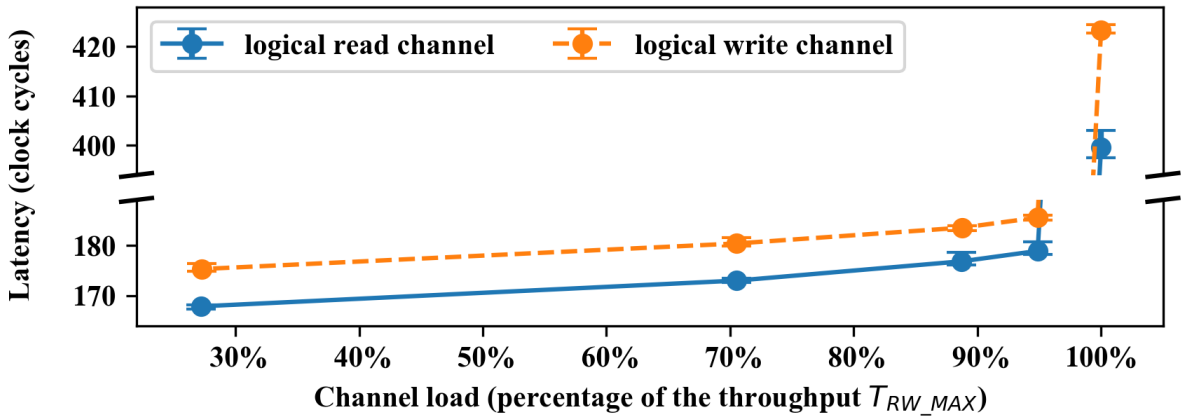

Fig.II.76. Load-latency curve for RingNet with 5 1st level rings and 2 PGs connected at each ring, and 4 parallel rings used at the network root (F=5 × G=2, R=4).

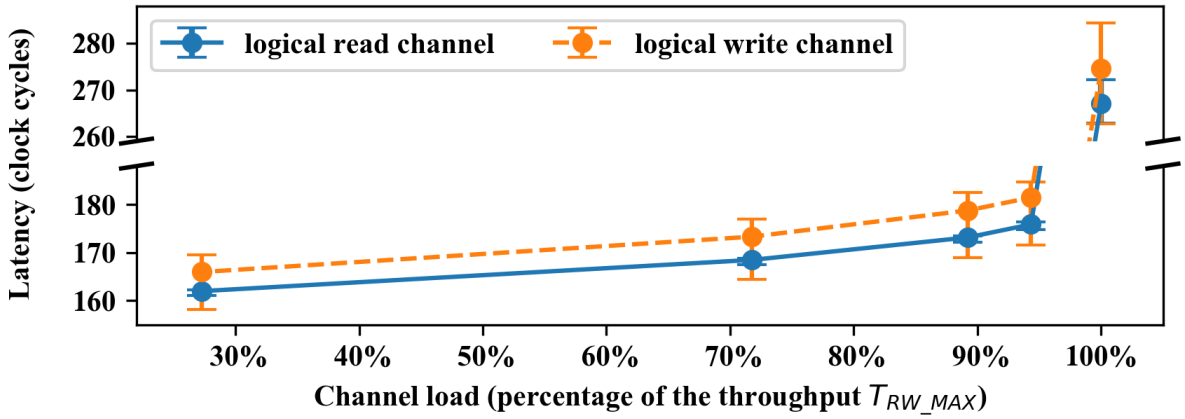

Fig.II.77. Load-latency curve for RingNet with 4 1st level rings and 3 PGs connected at each ring, and 4 parallel rings used at the network root (F=4 × G=3, R=4).

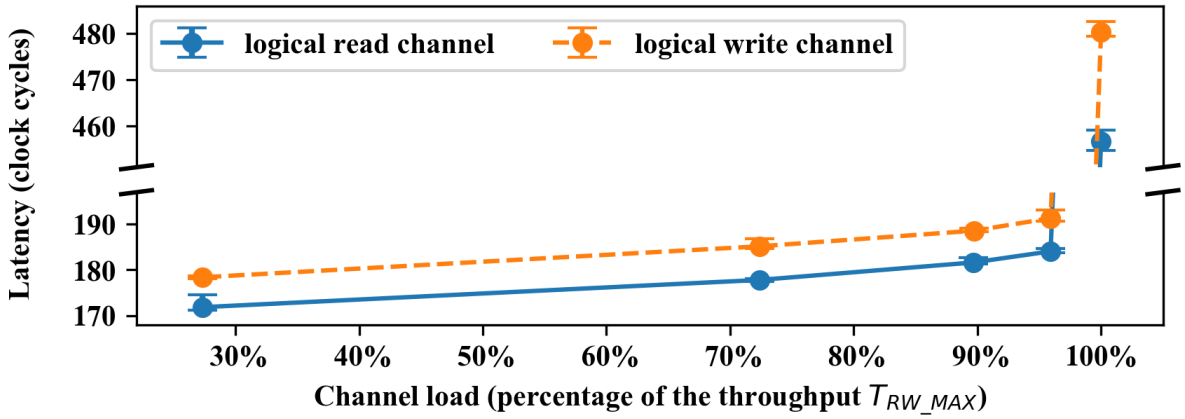

Fig.II.78. Load-latency curve for RingNet with 5 1st level rings and 3 PGs connected at each ring, and 4 parallel rings used at the network root (F=5 × G=3, R=4).

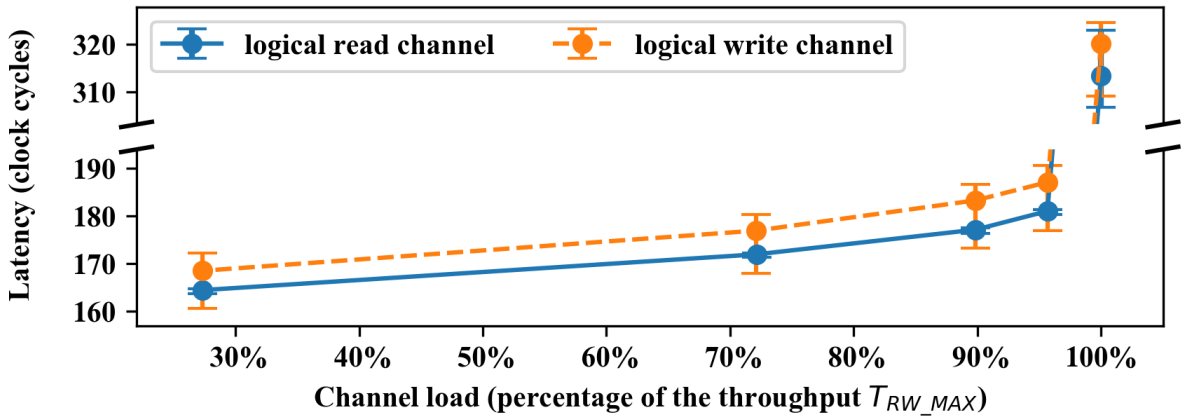

Fig.II.79. Load-latency curve for RingNet with 4 1st level rings and 4 PGs connected at each ring, and 4 parallel rings used at the network root (F=4 × G=4, R=4).

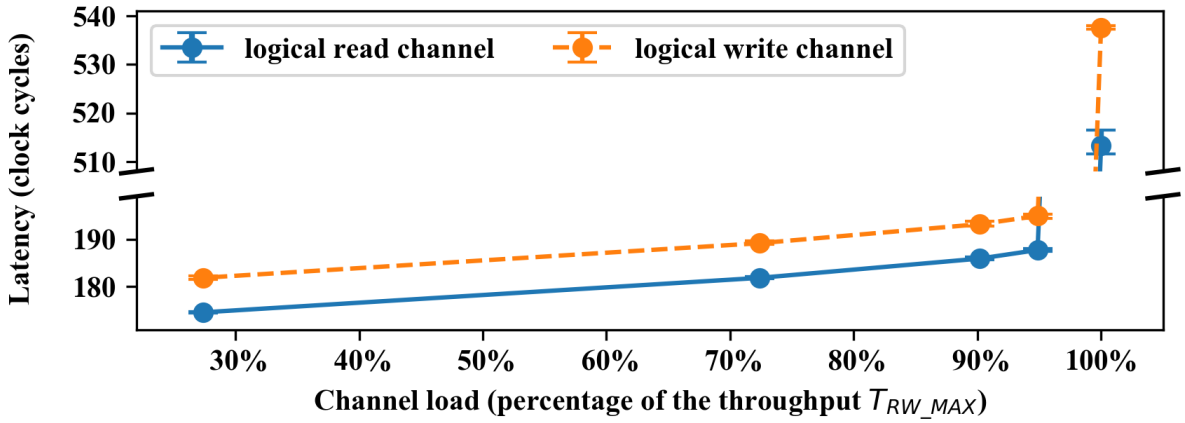

Fig.II.80. Load-latency curve for RingNet with 5 1st level rings and 4 PGs connected at each ring, and 4 parallel rings used at the network root (F=5 × G=4, R=4).

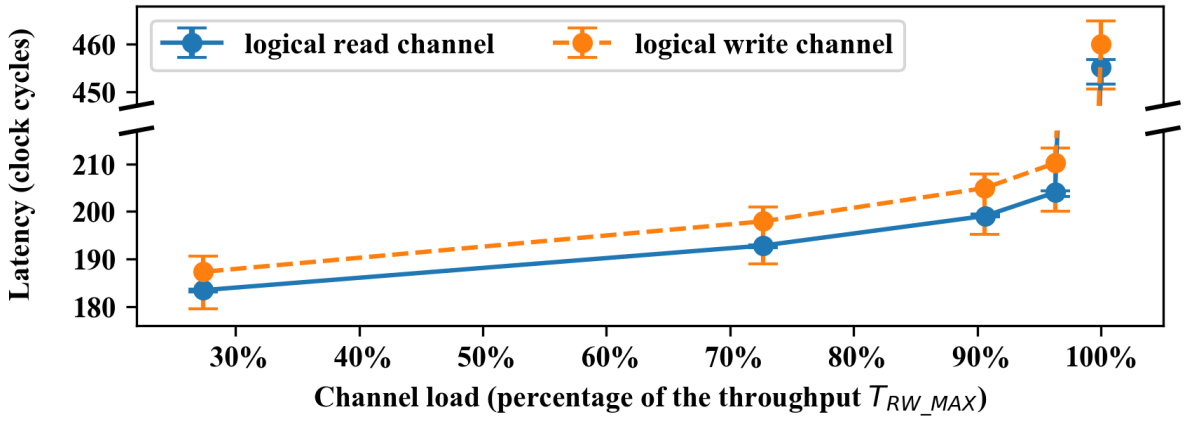

Fig.II.81. Load-latency curve for RingNet with 4 1st level rings and 7 PGs connected at each ring, and 4 parallel rings used at the network root (F=4 × G=7, R=4).

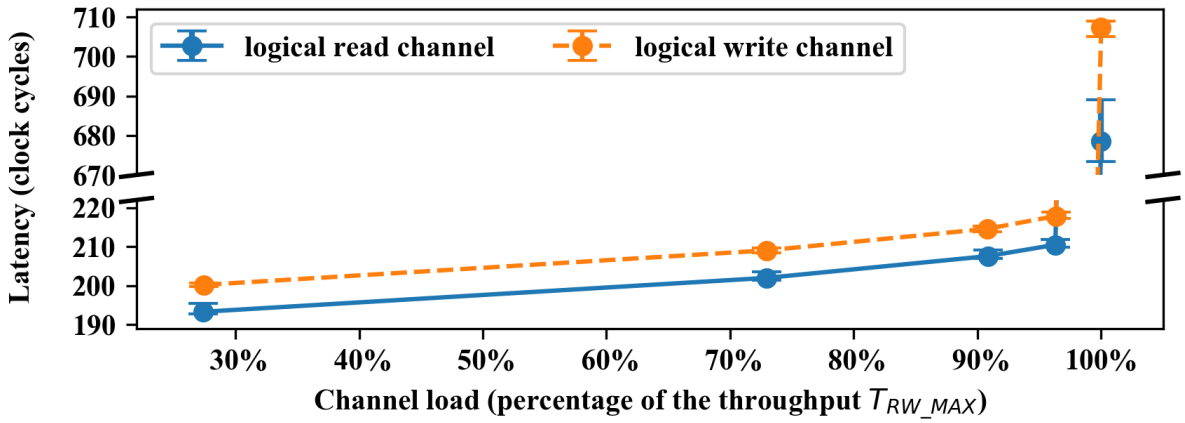

Fig.II.82. Load-latency curve for RingNet with 5 1st level rings and 7 PGs connected at each ring, and 4 parallel rings used at the network root (F=5 × G=7, R=4).

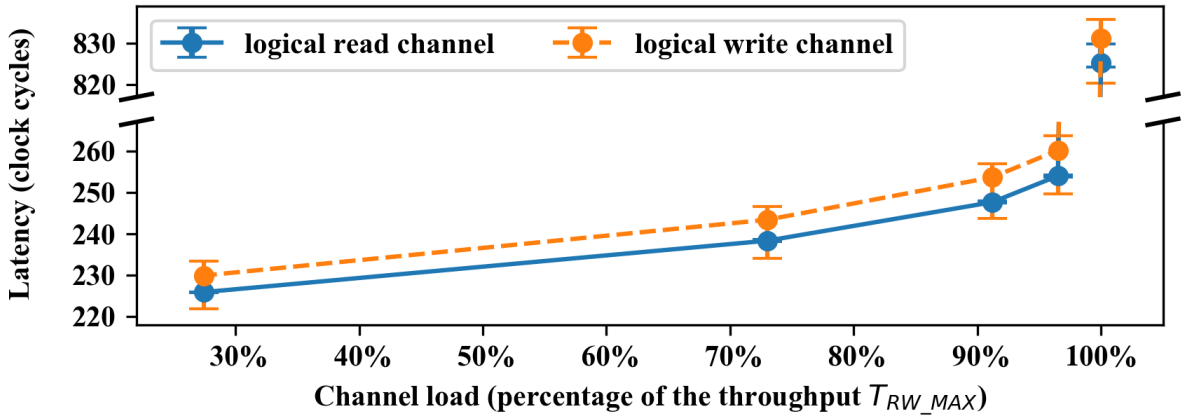

Fig.II.83. Load-latency curve for RingNet with 4 1st level rings and 15 PGs connected at each ring, and 4 parallel rings used at the network root (F=4 × G=15, R=4).

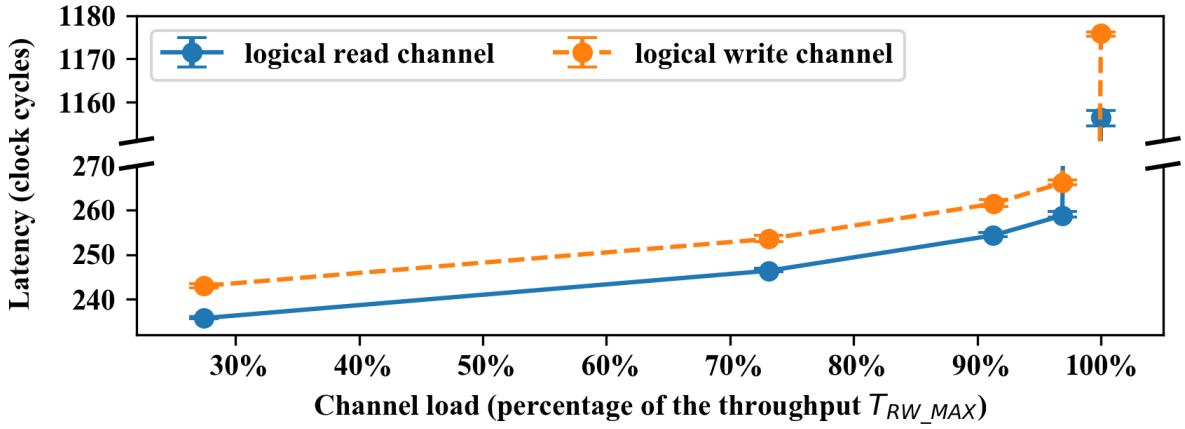

Fig.II.84. Load-latency curve for RingNet with 5 1st level rings and 15 PGs connected at each ring, and 4 parallel rings used at the network root (F=5 × G=15, R=4).

Appendix III. Simulation results of the test for the fairness of the RingNet access

In Appendix III, the results from the fairness experiment are presented.

Parameters of the test are:

- *R*: Multiplication degree of a root level, i.e., number of parallel rings used at the root level, set in the range of 1–4.
- *F*: The number of 1st level rings, set in the range of 1–5.
- *G*: The number of packet generators (PGs) connected to a single 1st level ring, set in the range of 1–15.
- Logical channel load. The aggregated load generated by all PGs is set in the range of 0%–100% of the theoretical throughput T_{RW_MAX} (1). The same logical channel load is set for the read and write channel.

PEs are simulated using Packet Generators (PGs). For each individual packet generator (PG), we collect the values of the average latency L_{PG} expressed in clock cycles, and the values of the throughput T_{PG} expressed in bits per clock cycle, for both logical channels. In tables, the values of the L_{PG}, and the values of T_{PG}, calculated over all PGs ($\overline{L_{PG}}$ and $\overline{T_{PG}}$, respectively), are presented together with a standard deviation for those variables ($\sigma_{\overline{L_{PG}}}$ and $\sigma_{\overline{T_{PG}}}$, respectively).

TABLE III.1
TRAFFIC STATISTICS FOR RINGNET WITH 2 1ST LEVEL RINGS AND 1 PG CONNECTED AT EACH RING, AND ONE RING USED AT THE NETWORK ROOT (F=2 × G=1, R=1).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	110	4.65	117	4.75	6.3	0.015	6.3	0.010
71%	111	3.60	118	3.55	16.5	0.005	16.5	0.015
88%	113	3.15	119	3.10	20.4	0.025	20.5	0.010
92%	113	2.95	120	2.90	21.3	0.010	21.3	0.005
100%	210	5.50	217	5.50	23.3	0.000	23.3	0.000

TABLE III.2
TRAFFIC STATISTICS FOR RINGNET WITH 3 1ST LEVEL RINGS AND 1 PG CONNECTED AT EACH RING, AND ONE RING USED AT THE NETWORK ROOT (F=3 × G=1, R=1).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	123	4.57	123	4.41	4.2	0.000	4.2	0.012
72%	126	3.50	131	3.15	11.1	0.052	11.1	0.029
89%	127	3.20	134	3.19	13.8	0.041	13.8	0.014
94%	129	3.06	135	3.11	14.6	0.012	14.6	0.005
100%	318	4.68	325	4.68	15.5	0.000	15.5	0.005

TABLE III.3
TRAFFIC STATISTICS FOR RINGNET WITH 4 1ST LEVEL RINGS AND 1 PG CONNECTED AT EACH RING, AND ONE RING USED AT THE NETWORK ROOT (F=4 × G=1, R=1).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	122	4.57	121	4.61	3.2	0.008	3.2	0.005
72%	125	3.50	130	3.48	8.4	0.011	8.4	0.015
90%	128	3.26	134	3.18	10.5	0.021	10.5	0.019
96%	130	3.10	136	3.11	11.1	0.024	11.1	0.008
100%	417	5.27	425	4.13	11.6	0.000	11.6	0.000

TABLE III.4
TRAFFIC STATISTICS FOR RINGNET WITH 5 1ST LEVEL RINGS AND 1 PG CONNECTED AT EACH RING, AND ONE RING USED AT THE NETWORK ROOT (F=5 × G=1, R=1).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	132	4.56	139	4.64	2.5	0.014	2.5	0.012
72%	136	3.55	142	3.71	6.7	0.022	6.7	0.010
90%	138	3.31	145	3.18	8.4	0.020	8.4	0.019
95%	140	3.05	147	3.17	8.8	0.016	8.8	0.020
100%	514	5.94	521	5.90	9.3	0.000	9.3	0.000

TABLE III.5
TRAFFIC STATISTICS FOR RINGNET WITH ONE 1ST LEVEL RING AND 2 PGS CONNECTED AT THE RING, AND ONE RING USED AT THE NETWORK ROOT (F=1 × G=2, R=1).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	115	0.35	114	0.45	6.3	0.015	6.3	0.010
71%	116	0.25	122	0.15	16.5	0.005	16.5	0.015
88%	118	0.20	124	0.10	20.4	0.025	20.5	0.010
92%	118	0.20	125	0.15	21.3	0.010	21.3	0.005
100%	166	0.50	172	0.50	23.3	0.000	23.3	0.000

TABLE III.6
TRAFFIC STATISTICS FOR RINGNET WITH 2 1ST LEVEL RINGS AND 2 PGS CONNECTED AT EACH RING, AND ONE RING USED AT THE NETWORK ROOT (F=2 × G=2, R=1).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	127	0.89	134	0.91	3.2	0.008	3.2	0.005
72%	130	1.98	137	1.95	8.4	0.011	8.4	0.015
90%	132	2.44	139	2.50	10.5	0.021	10.5	0.019
96%	134	2.64	141	2.63	11.1	0.024	11.1	0.008
100%	297	4.46	304	0.50	11.6	0.004	11.6	0.000

TABLE III.7
TRAFFIC STATISTICS FOR RINGNET WITH 3 1ST LEVEL RINGS AND 2 PGS CONNECTED AT EACH RING, AND ONE RING USED AT THE NETWORK ROOT (F=3 × G=2, R=1).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	138	0.97	138	1.19	2.1	0.007	2.1	0.005
73%	142	2.09	148	2.53	5.6	0.022	5.6	0.011
90%	145	2.60	152	2.70	7.0	0.017	7.0	0.022
96%	148	2.84	155	2.96	7.4	0.027	7.4	0.021
100%	459	6.89	469	3.68	7.8	0.000	7.8	0.000

TABLE III.8

TRAFFIC STATISTICS FOR RINGNET WITH 4 1ST LEVEL RINGS AND 2 PGS CONNECTED AT EACH RING, AND ONE RING USED AT THE NETWORK ROOT (F=4 × G=2, R=1).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	138	0.88	138	0.91	1.6	0.007	1.6	0.007
73%	143	2.28	149	2.49	4.2	0.012	4.2	0.015
91%	147	2.74	154	2.75	5.3	0.014	5.3	0.011
97%	151	2.84	158	3.03	5.6	0.014	5.6	0.015
100%	602	21.06	611	20.46	5.8	0.003	5.8	0.000

TABLE III.9

TRAFFIC STATISTICS FOR RINGNET WITH 5 1ST LEVEL RINGS AND 2 PGS CONNECTED AT EACH RING, AND ONE RING USED AT THE NETWORK ROOT (F=5 × G=2, R=1).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	147	3.57	154	3.77	1.3	0.008	1.3	0.005
73%	153	2.85	160	3.15	3.4	0.012	3.4	0.008
91%	158	3.08	165	3.15	4.2	0.021	4.2	0.011
97%	163	3.03	169	3.19	4.5	0.012	4.5	0.014
100%	750	19.04	762	15.14	4.7	0.000	4.7	0.004

TABLE III.10

TRAFFIC STATISTICS FOR RINGNET WITH ONE 1ST LEVEL RING AND 3 PGS CONNECTED AT THE RING, AND ONE RING USED AT THE NETWORK ROOT (F=1 × G=3, R=1).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	115	0.78	114	0.67	4.2	0.012	4.2	0.009
72%	117	0.66	123	0.45	11.1	0.012	11.1	0.012
89%	119	0.53	126	0.42	13.8	0.009	13.8	0.028
94%	120	0.45	127	0.37	14.6	0.005	14.6	0.016
100%	209	0.82	216	0.82	15.5	0.005	15.5	0.005

TABLE III.11
TRAFFIC STATISTICS FOR RINGNET WITH 2 1ST LEVEL RINGS AND 3 PGS CONNECTED AT EACH RING, AND ONE RING USED AT THE NETWORK ROOT (F=2 × G=3, R=1).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	132	4.72	139	5.00	2.1	0.005	2.1	0.013
73%	136	3.44	143	4.84	5.6	0.012	5.6	0.015
90%	139	3.19	146	5.22	7.0	0.007	7.0	0.014
96%	142	2.70	149	5.41	7.4	0.020	7.4	0.021
100%	390	5.55	397	3.40	7.8	0.000	7.8	0.000

TABLE III.12
TRAFFIC STATISTICS FOR RINGNET WITH 3 1ST LEVEL RINGS AND 3 PGS CONNECTED AT EACH RING, AND ONE RING USED AT THE NETWORK ROOT (F=3 × G=3, R=1).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	141	4.45	141	4.44	1.4	0.008	1.4	0.008
73%	147	3.52	153	3.84	3.8	0.010	3.8	0.011
91%	151	3.20	158	3.91	4.7	0.016	4.7	0.015
96%	155	3.05	162	3.93	5.0	0.009	5.0	0.016
100%	594	13.56	601	9.40	5.2	0.000	5.2	0.000

TABLE III.13
TRAFFIC STATISTICS FOR RINGNET WITH 4 1ST LEVEL RINGS AND 3 PGS CONNECTED AT EACH RING, AND ONE RING USED AT THE NETWORK ROOT (F=4 × G=3, R=1).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	141	4.18	140	4.16	1.1	0.008	1.1	0.009
73%	147	3.53	153	3.78	2.8	0.012	2.8	0.013
91%	153	3.37	159	3.45	3.5	0.015	3.5	0.009
96%	157	3.34	163	3.62	3.7	0.014	3.7	0.018
100%	780	14.11	790	9.92	3.9	0.000	3.9	0.000

TABLE III.14
TRAFFIC STATISTICS FOR RINGNET WITH 5 1ST LEVEL RINGS AND 3 PGS CONNECTED AT EACH RING, AND ONE RING USED AT THE NETWORK ROOT (F=5 × G=3, R=1).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	151	3.89	158	3.92	0.9	0.008	0.9	0.008
73%	158	3.24	165	3.60	2.3	0.011	2.3	0.011
91%	165	3.42	172	3.44	2.8	0.009	2.8	0.013
97%	169	3.00	177	3.52	3.0	0.016	3.0	0.010
100%	980	28.97	979	6.81	3.1	0.004	3.1	0.002

TABLE III.15
TRAFFIC STATISTICS FOR RINGNET WITH ONE 1ST LEVEL RING AND 4 PGS CONNECTED AT THE RING, AND ONE RING USED AT THE NETWORK ROOT (F=1 × G=4, R=1).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	115	0.94	114	0.95	3.2	0.011	3.2	0.013
72%	118	0.82	124	0.65	8.4	0.022	8.4	0.018
90%	120	0.67	127	0.78	10.5	0.013	10.5	0.011
96%	122	0.66	129	0.67	11.1	0.018	11.1	0.028
100%	252	1.12	260	1.12	11.6	0.000	11.6	0.000

TABLE III.16
TRAFFIC STATISTICS FOR RINGNET WITH 2 1ST LEVEL RINGS AND 4 PGS CONNECTED AT EACH RING, AND ONE RING USED AT THE NETWORK ROOT (F=2 × G=4, R=1).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	132	5.00	138	4.95	1.6	0.008	1.6	0.009
73%	137	3.64	144	3.51	4.2	0.015	4.2	0.012
91%	141	3.04	147	3.08	5.3	0.017	5.3	0.020
97%	145	2.82	152	2.45	5.6	0.019	5.6	0.011
100%	478	5.61	485	5.63	5.8	0.000	5.8	0.000

TABLE III.17
TRAFFIC STATISTICS FOR RINGNET WITH 3 1ST LEVEL RINGS AND 4 PGS CONNECTED AT EACH RING, AND ONE RING USED AT THE NETWORK ROOT (F=3 × G=4, R=1).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	144	4.68	144	4.29	1.1	0.008	1.1	0.009
73%	151	3.54	157	3.65	2.8	0.012	2.8	0.013
91%	157	3.25	164	3.07	3.5	0.015	3.5	0.009
96%	161	3.34	167	2.89	3.7	0.014	3.7	0.018
100%	729	21.70	736	19.26	3.9	0.000	3.9	0.000

TABLE III.18
TRAFFIC STATISTICS FOR RINGNET WITH 4 1ST LEVEL RINGS AND 4 PGS CONNECTED AT EACH RING, AND ONE RING USED AT THE NETWORK ROOT (F=4 × G=4, R=1).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	143	4.89	143	4.81	0.8	0.007	0.8	0.007
73%	150	3.77	156	3.71	2.1	0.009	2.1	0.012
91%	157	3.35	164	3.33	2.7	0.012	2.7	0.008
97%	163	3.36	170	3.06	2.8	0.009	2.8	0.010
100%	958	19.52	969	22.79	2.9	0.000	2.9	0.000

TABLE III.19
TRAFFIC STATISTICS FOR RINGNET WITH 5 1ST LEVEL RINGS AND 4 PGS CONNECTED AT EACH RING, AND ONE RING USED AT THE NETWORK ROOT (F=5 × G=4, R=1).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
28%	153	4.60	160	4.78	0.6	0.004	0.6	0.005
73%	160	3.75	168	3.69	1.7	0.010	1.7	0.007
91%	168	3.56	176	3.49	2.1	0.007	2.1	0.007
97%	175	3.08	183	3.21	2.3	0.008	2.3	0.009
100%	1193	22.65	1201	40.66	2.3	0.002	2.3	0.003

TABLE III.20
TRAFFIC STATISTICS FOR RINGNET WITH ONE 1ST LEVEL RING AND 7 PGS CONNECTED AT THE RING, AND ONE RING USED AT THE NETWORK ROOT (F=1 × G=7, R=1).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	136	1.86	135	1.92	1.8	0.011	1.8	0.007
73%	140	1.72	146	1.64	4.8	0.016	4.8	0.006
91%	144	1.55	151	1.53	6.0	0.021	6.0	0.018
96%	147	1.56	155	1.49	6.4	0.021	6.4	0.012
100%	394	2.00	401	2.00	6.6	0.000	6.6	0.000

TABLE III.21
TRAFFIC STATISTICS FOR RINGNET WITH 2 1ST LEVEL RINGS AND 7 PGS CONNECTED AT EACH RING, AND ONE RING USED AT THE NETWORK ROOT (F=2 × G=7, R=1).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	152	5.04	159	5.26	0.9	0.007	0.9	0.006
73%	160	3.65	167	3.59	2.4	0.012	2.4	0.006
91%	169	2.14	175	2.36	3.0	0.010	3.0	0.010
97%	181	3.18	187	3.19	3.2	0.011	3.2	0.012
100%	749	16.24	756	17.24	3.3	0.005	3.3	0.005

TABLE III.22
TRAFFIC STATISTICS FOR RINGNET WITH 3 1ST LEVEL RINGS AND 7 PGS CONNECTED AT EACH RING, AND ONE RING USED AT THE NETWORK ROOT (F=3 × G=7, R=1).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	162	4.91	162	4.92	0.6	0.004	0.6	0.006
73%	170	3.94	175	4.17	1.6	0.008	1.6	0.011
91%	178	3.26	185	3.49	2.0	0.009	2.0	0.010
97%	185	3.41	193	3.37	2.2	0.011	2.2	0.009
100%	1127	3.42	1134	3.38	2.2	0.000	2.2	0.002

TABLE III.23
TRAFFIC STATISTICS FOR RINGNET WITH 4 1ST LEVEL RINGS AND 7 PGS CONNECTED AT EACH RING, AND ONE RING USED AT THE NETWORK ROOT (F=4 × G=7, R=1).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
28%	161	4.84	161	4.51	0.5	0.005	0.5	0.006
73%	170	3.74	175	3.98	1.2	0.008	1.2	0.009
91%	179	3.67	186	3.72	1.5	0.010	1.5	0.010
97%	189	3.75	196	3.67	1.6	0.009	1.6	0.008
100%	1516	15.14	1523	18.88	1.7	0.000	1.7	0.000

TABLE III.24
TRAFFIC STATISTICS FOR RINGNET WITH 5 1ST LEVEL RINGS AND 7 PGS CONNECTED AT EACH RING, AND ONE RING USED AT THE NETWORK ROOT (F=5 × G=7, R=1).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	171	4.24	178	4.43	0.4	0.006	0.4	0.006
73%	181	3.79	188	3.94	1.0	0.007	1.0	0.008
91%	192	4.00	200	3.82	1.2	0.008	1.2	0.009
97%	201	3.68	210	3.75	1.3	0.007	1.3	0.008
100%	1877	33.38	1879	25.62	1.3	0.000	1.3	0.000

TABLE III.25
TRAFFIC STATISTICS FOR RINGNET WITH ONE 1ST LEVEL RING AND 15 PGS CONNECTED AT THE RING, AND ONE RING USED AT THE NETWORK ROOT (F=1 × G=15, R=1).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	176	4.13	176	4.06	0.9	0.008	0.9	0.008
73%	183	3.84	189	3.87	2.3	0.011	2.3	0.011
91%	189	3.90	196	3.95	2.8	0.009	2.8	0.013
97%	194	3.70	202	3.76	3.0	0.016	3.0	0.010
100%	764	4.32	771	4.32	3.1	0.000	3.1	0.000

TABLE III.26
TRAFFIC STATISTICS FOR RINGNET WITH 2 1ST LEVEL RINGS AND 15 PGS CONNECTED AT EACH RING, AND ONE RING USED AT THE NETWORK ROOT (F=2 × G=15, R=1).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
28%	192	6.34	199	6.37	0.4	0.005	0.4	0.005
73%	203	5.27	210	5.08	1.1	0.009	1.1	0.006
91%	214	4.64	222	4.91	1.4	0.008	1.4	0.008
97%	225	4.12	232	4.43	1.5	0.010	1.5	0.008
100%	1528	6.99	1536	6.99	1.5	0.000	1.5	0.000

TABLE III.27
TRAFFIC STATISTICS FOR RINGNET WITH 3 1ST LEVEL RINGS AND 15 PGS CONNECTED AT EACH RING, AND ONE RING USED AT THE NETWORK ROOT (F=3 × G=15, R=1).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	205	6.14	205	6.14	0.3	0.005	0.3	0.005
73%	216	5.57	221	5.42	0.8	0.007	0.8	0.008
92%	230	5.00	237	4.53	0.9	0.007	0.9	0.006
97%	241	4.96	249	5.07	1.0	0.007	1.0	0.009
100%	2300	6.75	2307	6.75	1.0	0.004	1.0	0.004

TABLE III.28
TRAFFIC STATISTICS FOR RINGNET WITH 4 1ST LEVEL RINGS AND 15 PGS CONNECTED AT EACH RING, AND ONE RING USED AT THE NETWORK ROOT (F=4 × G=15, R=1).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	204	6.40	204	6.27	0.2	0.004	0.2	0.004
73%	214	5.65	220	5.36	0.6	0.006	0.6	0.006
92%	231	5.62	237	5.47	0.7	0.006	0.7	0.007
97%	245	5.33	252	5.59	0.8	0.007	0.8	0.006
100%	3058	6.99	3064	6.99	0.8	0.004	0.8	0.004

TABLE III.29
TRAFFIC STATISTICS FOR RINGNET WITH 5 1ST LEVEL RINGS AND 15 PGS CONNECTED AT EACH RING, AND ONE RING USED AT THE NETWORK ROOT (F=5 × G=15, R=1).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	213	6.44	220	6.42	0.2	0.003	0.2	0.003
73%	224	5.78	231	5.39	0.5	0.005	0.5	0.005
92%	242	5.37	249	5.74	0.6	0.006	0.6	0.006
97%	258	5.01	266	5.57	0.6	0.006	0.6	0.007
100%	3672	40.46	3712	33.86	0.6	0.000	0.6	0.000

TABLE III.30
TRAFFIC STATISTICS FOR RINGNET WITH 2 1ST LEVEL RINGS AND 1 PG CONNECTED AT EACH RING, AND 2 PARALLEL RINGS USED AT THE NETWORK ROOT (F=2 × G=1, R=2).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	127	0.70	134	0.60	12.5	0.020	12.5	0.000
69%	129	2.55	136	1.95	32.0	0.070	32.0	0.005
85%	130	3.55	138	2.60	39.4	0.010	39.4	0.010
92%	132	3.75	140	2.60	42.7	0.085	42.7	0.015
100%	124	3.50	132	3.50	46.5	0.005	46.5	0.000

TABLE III.31
TRAFFIC STATISTICS FOR RINGNET WITH 3 1ST LEVEL RINGS AND 1 PG CONNECTED AT EACH RING, AND 2 PARALLEL RINGS USED AT THE NETWORK ROOT (F=3 × G=1, R=2).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	134	4.54	134	4.44	8.4	0.012	8.4	0.019
69%	137	3.53	142	3.23	21.4	0.033	21.3	0.045
87%	138	3.18	145	3.11	26.9	0.014	27.0	0.014
92%	139	3.02	146	2.91	28.5	0.012	28.4	0.000
100%	272	5.21	296	5.19	31.0	0.000	31.0	0.000

TABLE III.32
TRAFFIC STATISTICS FOR RINGNET WITH 4 1ST LEVEL RINGS AND 1 PG CONNECTED AT EACH RING, AND 2 PARALLEL RINGS USED AT THE NETWORK ROOT (F=4 × G=1, R=2).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	136	4.12	135	3.93	6.3	0.011	6.3	0.008
71%	138	3.44	144	3.01	16.5	0.015	16.5	0.012
88%	140	3.22	147	3.03	20.5	0.015	20.5	0.008
92%	141	3.02	148	2.94	21.3	0.031	21.3	0.015
100%	353	5.00	379	4.76	23.3	0.000	23.3	0.000

TABLE III.33
TRAFFIC STATISTICS FOR RINGNET WITH 5 1ST LEVEL RINGS AND 1 PG CONNECTED AT EACH RING, AND 2 PARALLEL RINGS USED AT THE NETWORK ROOT (F=5 × G=1, R=2).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	147	3.82	154	3.84	5.1	0.014	5.1	0.011
71%	150	3.25	158	3.22	13.1	0.034	13.1	0.009
89%	153	2.87	160	3.13	16.5	0.020	16.5	0.011
95%	155	2.91	162	3.05	17.7	0.021	17.7	0.030
100%	436	5.42	469	4.58	18.6	0.000	18.6	0.000

TABLE III.34
TRAFFIC STATISTICS FOR RINGNET WITH 2 1ST LEVEL RINGS AND 2 PGS CONNECTED AT EACH RING, AND 2 PARALLEL RINGS USED AT THE NETWORK ROOT (F=2 × G=2, R=2).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	138	1.05	145	0.95	6.3	0.011	6.3	0.008
71%	141	2.78	148	2.97	16.5	0.015	16.5	0.012
88%	143	3.73	151	4.55	20.5	0.015	20.5	0.008
92%	144	3.68	152	4.85	21.3	0.031	21.3	0.015
100%	193	5.52	200	5.52	23.3	0.000	23.3	0.000

TABLE III.35
TRAFFIC STATISTICS FOR RINGNET WITH 3 1ST LEVEL RINGS AND 2 PGS CONNECTED AT EACH RING, AND 2 PARALLEL RINGS USED AT THE NETWORK ROOT (F=3 × G=2, R=2).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	149	0.96	149	1.13	4.2	0.017	4.2	0.007
72%	153	2.52	159	2.67	11.1	0.022	11.1	0.011
89%	155	2.98	162	2.82	13.8	0.024	13.8	0.018
94%	156	3.11	163	2.80	14.6	0.020	14.6	0.017
100%	341	2.64	365	1.77	15.5	0.005	15.5	0.004

TABLE III.36
TRAFFIC STATISTICS FOR RINGNET WITH 4 1ST LEVEL RINGS AND 2 PGS CONNECTED AT EACH RING, AND 2 PARALLEL RINGS USED AT THE NETWORK ROOT (F=4 × G=2, R=2).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	149	1.02	149	1.00	3.2	0.012	3.2	0.008
72%	153	2.48	159	2.61	8.4	0.017	8.4	0.019
90%	156	2.95	163	2.98	10.5	0.019	10.5	0.015
96%	158	3.09	165	3.00	11.1	0.016	11.1	0.013
100%	449	8.59	478	9.68	11.6	0.000	11.6	0.000

TABLE III.37
TRAFFIC STATISTICS FOR RINGNET WITH 5 1ST LEVEL RINGS AND 2 PGS CONNECTED AT EACH RING, AND 2 PARALLEL RINGS USED AT THE NETWORK ROOT (F=5 × G=2, R=2).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	160	1.06	167	0.87	2.5	0.010	2.6	0.007
72%	165	2.41	172	2.47	6.7	0.020	6.7	0.019
90%	168	2.83	175	2.80	8.4	0.012	8.4	0.020
95%	169	3.16	176	3.02	8.8	0.020	8.8	0.012
100%	550	8.48	589	14.32	9.3	0.000	9.3	0.000

TABLE III.38
TRAFFIC STATISTICS FOR RINGNET WITH 2 1ST LEVEL RINGS AND 3 PGS CONNECTED AT EACH RING, AND 2 PARALLEL RINGS USED AT THE NETWORK ROOT (F=2 × G=3, R=2).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	138	1.16	145	1.12	4.2	0.010	4.2	0.015
72%	142	2.76	149	3.01	11.1	0.018	11.1	0.020
89%	145	3.80	152	4.64	13.8	0.013	13.8	0.015
94%	146	4.00	154	5.28	14.7	0.017	14.6	0.024
100%	236	5.56	238	0.82	15.5	0.005	15.5	0.004

TABLE III.39
TRAFFIC STATISTICS FOR RINGNET WITH 3 1ST LEVEL RINGS AND 3 PGS CONNECTED AT EACH RING, AND 2 PARALLEL RINGS USED AT THE NETWORK ROOT (F=3 × G=3, R=2).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	149	1.10	149	1.22	2.8	0.011	2.8	0.008
72%	153	2.61	160	2.82	7.4	0.021	7.4	0.013
90%	156	2.93	163	3.08	9.3	0.031	9.3	0.020
95%	159	3.01	165	3.01	9.9	0.021	9.8	0.017
100%	407	4.15	430	6.16	10.3	0.003	10.3	0.003

TABLE III.40
TRAFFIC STATISTICS FOR RINGNET WITH 4 1ST LEVEL RINGS AND 3 PGS CONNECTED AT EACH RING, AND 2 PARALLEL RINGS USED AT THE NETWORK ROOT (F=4 × G=3, R=2).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	152	4.19	152	4.07	2.1	0.009	2.1	0.013
73%	157	3.38	163	3.36	5.6	0.019	5.6	0.017
90%	160	3.18	167	3.31	7.0	0.016	7.0	0.013
96%	163	3.07	170	3.29	7.4	0.014	7.4	0.016
100%	541	5.94	571	6.71	7.8	0.000	7.8	0.000

TABLE III.41
TRAFFIC STATISTICS FOR RINGNET WITH 5 1ST LEVEL RINGS AND 3 PGS CONNECTED AT EACH RING, AND 2 PARALLEL RINGS USED AT THE NETWORK ROOT (F=5 × G=3, R=2).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	162	3.84	169	3.89	1.7	0.007	1.7	0.011
72%	168	3.31	175	3.31	4.5	0.012	4.5	0.014
91%	172	3.09	179	3.38	5.6	0.012	5.6	0.012
96%	175	3.08	182	3.29	6.0	0.013	6.0	0.022
100%	665	7.40	699	8.64	6.2	0.004	6.2	0.004

TABLE III.42
TRAFFIC STATISTICS FOR RINGNET WITH 2 1ST LEVEL RINGS AND 4 PGS CONNECTED AT EACH RING, AND 2 PARALLEL RINGS USED AT THE NETWORK ROOT (F=2 × G=4, R=2).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	143	4.75	150	4.64	3.2	0.010	3.2	0.017
72%	148	2.91	155	2.56	8.4	0.016	8.4	0.014
90%	151	1.94	159	1.13	10.5	0.018	10.5	0.019
96%	154	1.41	162	0.67	11.1	0.015	11.1	0.035
100%	280	5.61	287	5.61	11.6	0.003	11.6	0.000

TABLE III.43
TRAFFIC STATISTICS FOR RINGNET WITH 3 1ST LEVEL RINGS AND 4 PGS CONNECTED AT EACH RING, AND 2 PARALLEL RINGS USED AT THE NETWORK ROOT (F=3 × G=4, R=2).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	152	4.53	152	4.35	2.1	0.009	2.1	0.013
73%	158	3.27	164	3.26	5.6	0.019	5.6	0.017
90%	161	2.96	168	3.04	7.0	0.016	7.0	0.013
96%	163	3.05	170	3.11	7.4	0.014	7.4	0.016
100%	476	6.04	500	5.41	7.8	0.000	7.8	0.000

TABLE III.44
TRAFFIC STATISTICS FOR RINGNET WITH 4 1ST LEVEL RINGS AND 4 PGS CONNECTED AT EACH RING, AND 2 PARALLEL RINGS USED AT THE NETWORK ROOT (F=4 × G=4, R=2).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	154	4.81	154	4.70	1.6	0.007	1.6	0.010
73%	160	3.58	166	3.55	4.2	0.015	4.2	0.009
91%	164	3.29	171	3.30	5.3	0.014	5.3	0.016
97%	167	3.14	174	3.16	5.6	0.012	5.6	0.015
100%	632	12.20	661	12.42	5.8	0.000	5.8	0.000

TABLE III.45
TRAFFIC STATISTICS FOR RINGNET WITH 5 1ST LEVEL RINGS AND 4 PGS CONNECTED AT EACH RING, AND 2 PARALLEL RINGS USED AT THE NETWORK ROOT (F=5 × G=4, R=2).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	166	4.68	173	4.75	1.3	0.006	1.3	0.010
73%	172	3.68	179	3.60	3.4	0.011	3.4	0.011
91%	177	3.34	184	3.24	4.2	0.016	4.2	0.015
97%	180	3.22	187	3.24	4.5	0.012	4.5	0.010
100%	780	11.03	812	17.41	4.7	0.004	4.7	0.004

TABLE III.46
TRAFFIC STATISTICS FOR RINGNET WITH 2 1ST LEVEL RINGS AND 7 PGS CONNECTED AT EACH RING, AND 2 PARALLEL RINGS USED AT THE NETWORK ROOT (F=2 × G=7, R=2).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	162	3.93	169	3.93	1.8	0.010	1.8	0.009
73%	170	3.09	177	2.93	4.8	0.012	4.8	0.016
91%	175	2.40	183	1.85	6.0	0.017	6.0	0.015
96%	179	2.22	187	1.42	6.4	0.015	6.4	0.012
100%	422	5.85	428	5.85	6.7	0.000	6.7	0.000

TABLE III.47
TRAFFIC STATISTICS FOR RINGNET WITH 3 1ST LEVEL RINGS AND 7 PGS CONNECTED AT EACH RING, AND 2 PARALLEL RINGS USED AT THE NETWORK ROOT (F=3 × G=7, R=2).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	171	3.53	171	3.45	1.2	0.007	1.2	0.008
73%	178	3.44	184	3.40	3.2	0.012	3.2	0.009
91%	183	3.18	190	3.37	4.0	0.016	4.0	0.011
96%	186	3.46	193	3.31	4.3	0.012	4.3	0.015
100%	667	11.84	691	11.18	4.4	0.000	4.4	0.000

TABLE III.48
TRAFFIC STATISTICS FOR RINGNET WITH 4 1ST LEVEL RINGS AND 7 PGS CONNECTED AT EACH RING, AND 2 PARALLEL RINGS USED AT THE NETWORK ROOT (F=4 × G=7, R=2).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	173	4.52	174	4.70	0.9	0.006	0.9	0.007
73%	180	3.95	186	3.72	2.4	0.010	2.4	0.011
91%	186	3.79	193	3.62	3.0	0.010	3.0	0.013
97%	190	3.74	197	3.62	3.2	0.015	3.2	0.015
100%	892	7.94	916	8.07	3.3	0.005	3.3	0.005

TABLE III.49
TRAFFIC STATISTICS FOR RINGNET WITH 5 1ST LEVEL RINGS AND 7 PGS CONNECTED AT EACH RING, AND 2 PARALLEL RINGS USED AT THE NETWORK ROOT (F=5 × G=7, R=2).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	184	4.45	191	4.44	0.7	0.006	0.7	0.006
73%	190	3.83	197	3.80	1.9	0.011	1.9	0.010
91%	197	3.69	204	3.55	2.4	0.010	2.4	0.011
97%	202	3.74	209	3.59	2.6	0.013	2.6	0.012
100%	1108	12.34	1153	5.75	2.7	0.000	2.7	0.000

TABLE III.50
TRAFFIC STATISTICS FOR RINGNET WITH 2 1ST LEVEL RINGS AND 15 PGS CONNECTED AT EACH RING, AND 2 PARALLEL RINGS USED AT THE NETWORK ROOT (F=2 × G=15, R=2).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	204	6.22	211	6.06	0.9	0.008	0.9	0.007
73%	213	4.63	221	4.45	2.3	0.010	2.3	0.007
91%	221	4.12	228	3.94	2.8	0.012	2.8	0.010
97%	226	3.90	234	3.82	3.0	0.011	3.0	0.014
100%	792	6.99	798	6.99	3.1	0.000	3.1	0.000

TABLE III.51
TRAFFIC STATISTICS FOR RINGNET WITH 3 1ST LEVEL RINGS AND 15 PGS CONNECTED AT EACH RING, AND 2 PARALLEL RINGS USED AT THE NETWORK ROOT (F=3 × G=15, R=2).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	213	6.04	213	6.14	0.6	0.005	0.6	0.006
73%	221	5.22	227	5.01	1.5	0.008	1.5	0.008
91%	229	5.09	236	5.21	1.9	0.010	1.9	0.009
97%	235	5.03	242	4.78	2.0	0.008	2.0	0.012
100%	1257	15.60	1281	13.84	2.1	0.000	2.1	0.000

TABLE III.52
TRAFFIC STATISTICS FOR RINGNET WITH 4 1ST LEVEL RINGS AND 15 PGS CONNECTED AT EACH RING, AND 2 PARALLEL RINGS USED AT THE NETWORK ROOT (F=4 × G=15, R=2).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
28%	215	6.29	215	6.22	0.4	0.006	0.4	0.005
73%	223	5.51	229	5.41	1.1	0.006	1.1	0.007
91%	232	5.15	238	5.23	1.4	0.009	1.4	0.009
97%	238	4.91	246	5.00	1.5	0.009	1.5	0.009
100%	1671	28.19	1700	25.02	1.6	0.000	1.6	0.000

TABLE III.53
TRAFFIC STATISTICS FOR RINGNET WITH 5 1ST LEVEL RINGS AND 15 PGS CONNECTED AT EACH RING, AND 2 PARALLEL RINGS USED AT THE NETWORK ROOT (F=5 × G=15, R=2).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	227	6.31	234	6.37	0.3	0.005	0.3	0.005
73%	235	5.47	242	5.42	0.9	0.006	0.9	0.007
91%	244	5.19	252	5.24	1.1	0.009	1.1	0.007
97%	253	5.06	261	5.00	1.2	0.009	1.2	0.008
100%	2003	6.91	2037	9.90	1.2	0.000	1.2	0.000

TABLE III.54
TRAFFIC STATISTICS FOR RINGNET WITH 3 1ST LEVEL RINGS AND 1 PG CONNECTED AT EACH RING, AND 3 PARALLEL RINGS USED AT THE NETWORK ROOT (F=3 × G=1, R=3).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	142	4.57	142	4.54	12.5	0.033	12.5	0.019
69%	147	3.29	153	3.07	32.0	0.019	32.0	0.019
85%	150	2.94	157	2.57	39.4	0.034	39.4	0.037
92%	152	2.71	159	2.03	42.7	0.037	42.7	0.050
100%	152	7.79	163	9.72	46.5	0.005	46.5	0.000

TABLE III.55
TRAFFIC STATISTICS FOR RINGNET WITH 4 1ST LEVEL RINGS AND 1 PG CONNECTED AT EACH RING, AND 3 PARALLEL RINGS USED AT THE NETWORK ROOT (F=4 × G=1, R=3).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	141	4.02	141	3.80	9.5	0.022	9.5	0.004
70%	145	2.81	151	2.68	24.4	0.039	24.4	0.025
86%	147	2.37	155	2.19	30.1	0.040	30.1	0.032
92%	149	2.41	157	2.20	32.0	0.034	32.0	0.029
100%	311	4.15	334	5.50	34.9	0.000	34.9	0.000

TABLE III.56
TRAFFIC STATISTICS FOR RINGNET WITH 5 1ST LEVEL RINGS AND 1 PG CONNECTED AT EACH RING, AND 3 PARALLEL RINGS USED AT THE NETWORK ROOT (F=5 × G=1, R=3).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	152	3.67	159	3.67	7.5	0.026	7.5	0.019
71%	156	2.56	163	2.85	19.7	0.030	19.7	0.031
87%	158	2.30	166	2.70	24.4	0.029	24.4	0.023
92%	159	2.21	167	2.64	25.6	0.014	25.6	0.029
100%	374	4.40	402	7.57	27.9	0.000	27.9	0.000

TABLE III.57
TRAFFIC STATISTICS FOR RINGNET WITH 3 1ST LEVEL RINGS AND 2 PGS CONNECTED AT EACH RING, AND 3 PARALLEL RINGS USED AT THE NETWORK ROOT (F=3 × G=2, R=3).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	157	0.92	158	1.05	6.3	0.016	6.3	0.009
71%	163	2.40	170	2.55	16.5	0.026	16.5	0.018
88%	167	3.40	175	2.86	20.5	0.012	20.5	0.026
92%	169	3.33	177	3.01	21.3	0.016	21.3	0.024
100%	220	9.00	224	5.20	23.3	0.000	23.3	0.000

TABLE III.58
TRAFFIC STATISTICS FOR RINGNET WITH 4 1ST LEVEL RINGS AND 2 PGS CONNECTED AT EACH RING, AND 3 PARALLEL RINGS USED AT THE NETWORK ROOT (F=4 × G=2, R=3).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	155	2.64	155	2.55	4.7	0.015	4.7	0.012
72%	160	2.07	167	2.01	12.5	0.026	12.5	0.014
89%	163	2.25	171	1.98	15.5	0.031	15.5	0.027
95%	165	2.38	173	1.97	16.5	0.020	16.5	0.025
100%	365	3.22	393	4.79	17.4	0.000	17.5	0.005

TABLE III.59
TRAFFIC STATISTICS FOR RINGNET WITH 5 1ST LEVEL RINGS AND 2 PGS CONNECTED AT EACH RING, AND 3 PARALLEL RINGS USED AT THE NETWORK ROOT (F=5 × G=2, R=3).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	165	3.34	172	3.21	3.8	0.007	3.8	0.018
72%	170	2.85	178	2.87	10.0	0.018	10.0	0.018
89%	173	2.94	181	2.97	12.5	0.019	12.5	0.024
94%	174	2.95	183	2.89	13.1	0.014	13.1	0.021
100%	451	1.73	475	3.60	14.0	0.005	14.0	0.003

TABLE III.60
TRAFFIC STATISTICS FOR RINGNET WITH 3 1ST LEVEL RINGS AND 3 PGS CONNECTED AT EACH RING, AND 3 PARALLEL RINGS USED AT THE NETWORK ROOT (F=3 × G=3, R=3).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	158	1.86	159	1.97	4.2	0.009	4.2	0.016
72%	164	2.83	172	2.93	11.1	0.014	11.1	0.017
89%	169	3.55	178	3.10	13.8	0.019	13.8	0.024
94%	171	3.77	181	3.06	14.6	0.016	14.6	0.020
100%	264	9.02	275	7.40	15.5	0.005	15.5	0.005

TABLE III.61
TRAFFIC STATISTICS FOR RINGNET WITH 4 1ST LEVEL RINGS AND 3 PGS CONNECTED AT EACH RING, AND 3 PARALLEL RINGS USED AT THE NETWORK ROOT (F=4 × G=3, R=3).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	159	1.79	159	1.76	3.2	0.010	3.2	0.012
72%	164	1.91	172	1.90	8.4	0.021	8.4	0.014
90%	167	2.15	176	2.09	10.5	0.020	10.5	0.020
96%	170	2.30	178	2.16	11.1	0.018	11.1	0.024
100%	432	3.89	455	6.16	11.6	0.003	11.6	0.000

TABLE III.62
TRAFFIC STATISTICS FOR RINGNET WITH 5 1ST LEVEL RINGS AND 3 PGS CONNECTED AT EACH RING, AND 3 PARALLEL RINGS USED AT THE NETWORK ROOT (F=5 × G=3, R=3).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	168	3.68	175	3.69	2.5	0.009	2.6	0.010
72%	174	3.09	182	3.22	6.7	0.016	6.7	0.014
90%	177	3.09	185	3.00	8.4	0.022	8.4	0.016
95%	179	2.98	187	2.98	8.8	0.018	8.8	0.015
100%	533	8.59	558	10.69	9.3	0.000	9.3	0.000

TABLE III.63
TRAFFIC STATISTICS FOR RINGNET WITH 3 1ST LEVEL RINGS AND 4 PGS CONNECTED AT EACH RING, AND 3 PARALLEL RINGS USED AT THE NETWORK ROOT (F=3 × G=4, R=3).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	162	4.40	162	4.13	3.2	0.010	3.2	0.012
72%	169	3.25	177	2.91	8.4	0.021	8.4	0.014
90%	174	2.77	183	2.51	10.5	0.020	10.5	0.020
96%	177	3.01	187	2.33	11.1	0.018	11.1	0.024
100%	308	9.05	318	10.42	11.6	0.003	11.6	0.000

TABLE III.64
TRAFFIC STATISTICS FOR RINGNET WITH 4 1ST LEVEL RINGS AND 4 PGS CONNECTED AT EACH RING, AND 3 PARALLEL RINGS USED AT THE NETWORK ROOT (F=4 × G=4, R=3).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	161	3.61	162	3.51	2.4	0.010	2.4	0.011
72%	168	2.63	175	2.45	6.3	0.018	6.3	0.010
90%	171	2.29	180	2.07	7.9	0.020	7.9	0.015
96%	174	2.37	183	2.11	8.4	0.021	8.4	0.023
100%	494	4.89	515	4.89	8.7	0.002	8.7	0.000

TABLE III.65
TRAFFIC STATISTICS FOR RINGNET WITH 5 1ST LEVEL RINGS AND 4 PGS CONNECTED AT EACH RING, AND 3 PARALLEL RINGS USED AT THE NETWORK ROOT (F=5 × G=4, R=3).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	172	3.41	179	3.29	1.9	0.011	1.9	0.011
73%	178	2.86	186	2.41	5.1	0.016	5.1	0.014
91%	182	2.69	191	2.14	6.3	0.019	6.3	0.018
97%	186	2.59	194	2.01	6.7	0.016	6.7	0.017
100%	610	4.85	633	7.90	7.0	0.000	7.0	0.000

TABLE III.66
TRAFFIC STATISTICS FOR RINGNET WITH 3 1ST LEVEL RINGS AND 7 PGS CONNECTED AT EACH RING, AND 3 PARALLEL RINGS USED AT THE NETWORK ROOT (F=3 × G=7, R=3).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	181	3.97	181	3.91	1.8	0.008	1.8	0.009
73%	190	3.66	198	3.38	4.8	0.017	4.8	0.014
91%	197	3.38	206	3.12	6.0	0.017	6.0	0.016
96%	202	3.42	211	2.45	6.4	0.019	6.4	0.019
100%	449	9.20	456	9.20	6.7	0.000	6.7	0.000

TABLE III.67
TRAFFIC STATISTICS FOR RINGNET WITH 4 1ST LEVEL RINGS AND 7 PGS CONNECTED AT EACH RING, AND 3 PARALLEL RINGS USED AT THE NETWORK ROOT (F=4 × G=7, R=3).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	180	3.41	181	3.23	1.4	0.010	1.4	0.007
73%	188	3.03	196	2.91	3.6	0.009	3.6	0.011
91%	193	2.84	202	2.67	4.5	0.014	4.5	0.012
96%	196	2.82	205	2.82	4.8	0.017	4.8	0.014
100%	674	3.69	696	2.51	5.0	0.003	5.0	0.003

TABLE III.68
TRAFFIC STATISTICS FOR RINGNET WITH 5 1ST LEVEL RINGS AND 7 PGS CONNECTED AT EACH RING, AND 3 PARALLEL RINGS USED AT THE NETWORK ROOT (F=5 × G=7, R=3).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	190	4.02	197	3.85	1.1	0.007	1.1	0.007
73%	197	3.44	205	3.23	2.9	0.011	2.9	0.013
91%	203	3.37	211	2.99	3.6	0.012	3.6	0.013
97%	207	3.25	215	2.90	3.9	0.014	3.8	0.013
100%	830	7.14	856	6.31	4.0	0.000	4.0	0.000

TABLE III.69
TRAFFIC STATISTICS FOR RINGNET WITH 3 1ST LEVEL RINGS AND 15 PGS CONNECTED AT EACH RING, AND 3 PARALLEL RINGS USED AT THE NETWORK ROOT (F=3 × G=15, R=3).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	223	5.94	224	5.74	0.9	0.006	0.9	0.007
73%	235	4.80	243	4.68	2.3	0.011	2.3	0.007
91%	243	4.62	254	4.52	2.8	0.010	2.8	0.012
97%	249	5.01	260	4.29	3.0	0.012	3.0	0.010
100%	819	9.97	830	11.64	3.1	0.002	3.1	0.000

TABLE III.70
TRAFFIC STATISTICS FOR RINGNET WITH 4 1ST LEVEL RINGS AND 15 PGS CONNECTED AT EACH RING, AND 3 PARALLEL RINGS USED AT THE NETWORK ROOT (F=4 × G=15, R=3).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	222	5.50	223	5.60	0.6	0.005	0.6	0.006
73%	232	4.68	240	4.67	1.7	0.009	1.7	0.010
91%	239	4.50	248	4.45	2.1	0.009	2.1	0.010
97%	244	4.70	254	4.28	2.3	0.009	2.3	0.010
100%	1188	11.59	1210	15.96	2.3	0.001	2.3	0.001

TABLE III.71
TRAFFIC STATISTICS FOR RINGNET WITH 5 1ST LEVEL RINGS AND 15 PGS CONNECTED AT EACH RING, AND 3 PARALLEL RINGS USED AT THE NETWORK ROOT (F=5 × G=15, R=3).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	233	5.42	240	5.21	0.5	0.005	0.5	0.006
73%	241	4.89	249	4.75	1.4	0.009	1.4	0.008
91%	249	4.86	258	4.59	1.7	0.009	1.7	0.009
97%	254	4.72	263	4.48	1.8	0.009	1.8	0.010
100%	1441	7.66	1466	6.17	1.9	0.000	1.9	0.000

TABLE III.72
TRAFFIC STATISTICS FOR RINGNET WITH 4 1ST LEVEL RINGS AND 1 PG CONNECTED AT EACH RING, AND 4 PARALLEL RINGS USED AT THE NETWORK ROOT (F=4 × G=1, R=4).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	145	4.36	144	4.43	12.5	0.029	12.5	0.025
69%	149	3.43	154	3.08	32.0	0.034	32.0	0.029
85%	152	3.13	159	2.68	39.4	0.020	39.4	0.011
92%	155	2.88	162	2.27	42.7	0.043	42.7	0.038
100%	156	10.59	168	12.39	46.5	0.005	46.5	0.004

TABLE III.73
TRAFFIC STATISTICS FOR RINGNET WITH 5 1ST LEVEL RINGS AND 1 PG CONNECTED AT EACH RING, AND 4 PARALLEL RINGS USED AT THE NETWORK ROOT (F=5 × G=1, R=4).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{L_{PG}}$	$\sigma_{L_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$	$\overline{T_{PG}}$	$\sigma_{T_{PG}}$
27%	154	4.55	161	4.65	10.0	0.012	10.1	0.012
69%	158	3.80	166	4.34	25.6	0.014	25.6	0.029
86%	162	3.80	168	4.00	32.0	0.046	32.0	0.057
92%	163	3.45	170	3.93	34.2	0.038	34.1	0.022
100%	349	4.80	368	4.46	37.2	0.005	37.2	0.005

TABLE III.74
TRAFFIC STATISTICS FOR RINGNET WITH 4 1ST LEVEL RINGS AND 2 PGS CONNECTED AT EACH RING, AND 4 PARALLEL RINGS USED AT THE NETWORK ROOT (F=4 × G=2, R=4).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	159	3.05	158	3.35	6.3	0.026	6.3	0.017
71%	164	3.19	171	2.82	16.5	0.020	16.5	0.025
88%	169	3.01	177	2.78	20.5	0.031	20.5	0.031
92%	170	3.28	179	2.64	21.3	0.034	21.3	0.040
100%	223	9.13	230	12.72	23.3	0.000	23.3	0.000

TABLE III.75
TRAFFIC STATISTICS FOR RINGNET WITH 5 1ST LEVEL RINGS AND 2 PGS CONNECTED AT EACH RING, AND 4 PARALLEL RINGS USED AT THE NETWORK ROOT (F=5 × G=2, R=4).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	168	4.40	175	4.47	5.1	0.015	5.1	0.016
71%	173	4.34	180	4.02	13.1	0.014	13.1	0.021
89%	177	4.02	184	4.26	16.5	0.021	16.5	0.025
95%	179	4.02	185	4.01	17.7	0.025	17.7	0.024
100%	401	4.33	423	5.01	18.6	0.000	18.6	0.000

TABLE III.76
TRAFFIC STATISTICS FOR RINGNET WITH 4 1ST LEVEL RINGS AND 3 PGS CONNECTED AT EACH RING, AND 4 PARALLEL RINGS USED AT THE NETWORK ROOT (F=4 × G=3, R=4).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	162	2.82	162	2.57	4.2	0.009	4.2	0.014
72%	169	3.12	175	2.90	11.1	0.018	11.1	0.024
89%	174	2.96	182	2.49	13.8	0.025	13.8	0.023
94%	176	2.91	185	2.03	14.6	0.023	14.6	0.024
100%	272	9.49	279	9.49	15.5	0.005	15.5	0.004

TABLE III.77
TRAFFIC STATISTICS FOR RINGNET WITH 5 1ST LEVEL RINGS AND 3 PGS CONNECTED AT EACH RING, AND 4 PARALLEL RINGS USED AT THE NETWORK ROOT (F=5 × G=3, R=4).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	171	4.36	178	4.22	3.4	0.010	3.4	0.007
72%	178	4.01	185	4.03	9.0	0.013	9.0	0.015
90%	181	3.84	188	3.80	11.1	0.019	11.1	0.014
96%	184	3.88	191	3.63	11.9	0.025	11.9	0.017
100%	459	6.88	481	6.18	12.4	0.000	12.4	0.000

TABLE III.78
TRAFFIC STATISTICS FOR RINGNET WITH 4 1ST LEVEL RINGS AND 4 PGS CONNECTED AT EACH RING, AND 4 PARALLEL RINGS USED AT THE NETWORK ROOT (F=4 × G=4, R=4).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	165	4.15	165	4.23	3.2	0.011	3.2	0.012
72%	172	3.36	179	2.90	8.4	0.021	8.4	0.023
90%	177	3.25	186	2.58	10.5	0.026	10.5	0.032
96%	181	2.97	191	2.13	11.1	0.019	11.1	0.020
100%	316	10.11	323	15.35	11.6	0.002	11.6	0.002

TABLE III.79
TRAFFIC STATISTICS FOR RINGNET WITH 5 1ST LEVEL RINGS AND 4 PGS CONNECTED AT EACH RING, AND 4 PARALLEL RINGS USED AT THE NETWORK ROOT (F=5 × G=4, R=4).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	175	4.34	182	4.24	2.5	0.013	2.6	0.010
72%	182	3.91	189	3.64	6.7	0.016	6.7	0.017
90%	186	3.61	194	3.42	8.4	0.017	8.4	0.016
95%	188	3.73	195	3.46	8.8	0.020	8.8	0.021
100%	515	8.75	538	7.64	9.3	0.000	9.3	0.000

TABLE III.80
TRAFFIC STATISTICS FOR RINGNET WITH 4 1ST LEVEL RINGS AND 7 PGS CONNECTED AT EACH RING, AND 4 PARALLEL RINGS USED AT THE NETWORK ROOT (F=4 × G=7, R=4).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	184	3.90	184	3.60	1.8	0.009	1.8	0.008
73%	193	3.51	200	2.99	4.8	0.013	4.8	0.014
91%	199	3.32	207	2.88	6.0	0.017	6.0	0.018
96%	204	3.14	213	2.00	6.4	0.015	6.4	0.019
100%	454	12.46	462	12.46	6.7	0.000	6.7	0.000

TABLE III.81
TRAFFIC STATISTICS FOR RINGNET WITH 5 1ST LEVEL RINGS AND 7 PGS CONNECTED AT EACH RING, AND 4 PARALLEL RINGS USED AT THE NETWORK ROOT (F=5 × G=7, R=4).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	193	4.67	200	4.61	1.5	0.008	1.5	0.008
73%	202	4.31	209	4.10	3.9	0.014	3.9	0.016
91%	208	4.04	215	3.80	4.8	0.013	4.8	0.015
96%	211	4.15	218	3.75	5.1	0.011	5.1	0.017
100%	689	6.11	708	6.42	5.3	0.000	5.3	0.000

TABLE III.82
TRAFFIC STATISTICS FOR RINGNET WITH 4 1ST LEVEL RINGS AND 15 PGS CONNECTED AT EACH RING, AND 4 PARALLEL RINGS USED AT THE NETWORK ROOT (F=4 × G=15, R=4).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	226	5.86	226	5.85	0.9	0.008	0.9	0.007
73%	238	5.19	246	4.64	2.3	0.009	2.3	0.011
91%	247	4.93	257	4.39	2.8	0.013	2.8	0.014
97%	254	4.57	263	3.97	3.0	0.011	3.0	0.014
100%	824	13.04	834	14.96	3.1	0.002	3.1	0.002

TABLE III.83
TRAFFIC STATISTICS FOR RINGNET WITH 5 1ST LEVEL RINGS AND 15 PGS CONNECTED AT EACH RING, AND 4 PARALLEL RINGS USED AT THE NETWORK ROOT (F=5 × G=15, R=4).

Load (percentage of the throughput T_{RW_MAX} (1))	Average latency (clock cycles)				Average throughput (bits per cycle)			
	Read		Write		Read		Write	
	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{L_{PG}}$	$\sigma_{\overline{L_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$	$\overline{T_{PG}}$	$\sigma_{\overline{T_{PG}}}$
27%	236	5.97	243	5.92	0.7	0.007	0.7	0.006
73%	246	5.67	254	5.29	1.8	0.010	1.8	0.009
91%	254	5.25	262	5.03	2.3	0.012	2.3	0.009
97%	259	5.37	267	5.00	2.4	0.011	2.4	0.010
100%	1157	6.93	1175	8.54	2.5	0.000	2.5	0.000

Appendix IV. Detailed description of the key network components

In Appendix IV, three key components of the RingNet ring are presented: Leaf Interface, Root Interface, and Root-to-Leaf Manager (cf. Fig. 2 in the paper).

A. Leaf-to-Root Manager

The access to the Leaf-to-Root (L2R) channel slots is controlled by the L2R Manager. In order to access this channel, the Leaf Interface (LI) sends a request to the L2R Manager and obtains permission. The requests and permissions are sent using the L2R control channel. Permission is sent synchronously to the header flit of a granted packet slot. Requests are sent in the remaining time slots (cf. Fig. 4 in the paper). In Fig. IV.1 fields of the permission and request flits are presented.

	Field name	Number of bits	
PERM	Permission Valid (PV)	1	12
	Packet Length (L)	1	
	Packet Priority (PP)	2	
	Leaf interface ID (LID)	4	
	Request ID (RID)	4	
REQ	Request Valid (PV)	1	12
	Packet Length (L)	1	
	Packet Priority (PP)	2	
	Leaf interface ID (LID)	4	
	Request ID (RID)	4	

Fig. IV.1. Permission and request format.

The permission and the request flits have similar format. One bit is used to indicate the permission and request validity (Permission Valid and Request Valid fields). One bit indicates if the granted or requested slot is a long or a short one (Packet Length field). RingNet supports 4 packet priorities, ordered from the highest priority 3 to the lowest priority 0. Each packet and the associated request and permission has an assigned priority (Packet Priority field). Leaf Interface ID field identifies the source of the request or the destination of a permission. Value of the Request ID field from the request flit is copied by the L2R Manager to the corresponding permission. It helps the Leaf Interface to identify which buffered packet the permission relates to, especially when it sends many overlapping in time request for many slots.

Block diagram of the L2R Manager is depicted in Fig. IV.2. The L2R Manager keeps the requests in LUTRAM-based buffers. Individual buffers are used for requests of different priorities. If the L2R Manager detects a free packet slot, it generates a permission according to the buffered requests. The permissions are granted based on the priority of buffered requests and their order of arrival.

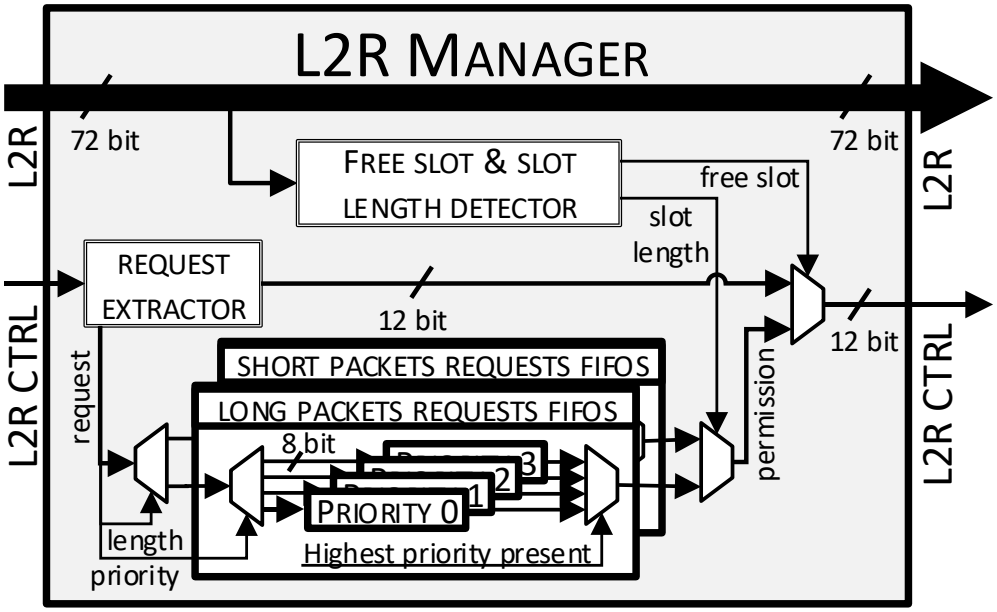

Fig. IV.2. L2R Manager.

B. Leaf Interface

Both Root Interface (RI) and Leaf Interface (LI) are 3-port switches and they insert packets at a ring using a 2×1 multiplexer and take packets from a ring using 1×2 demultiplexers (see Fig. IV.3 and Fig. IV.4).

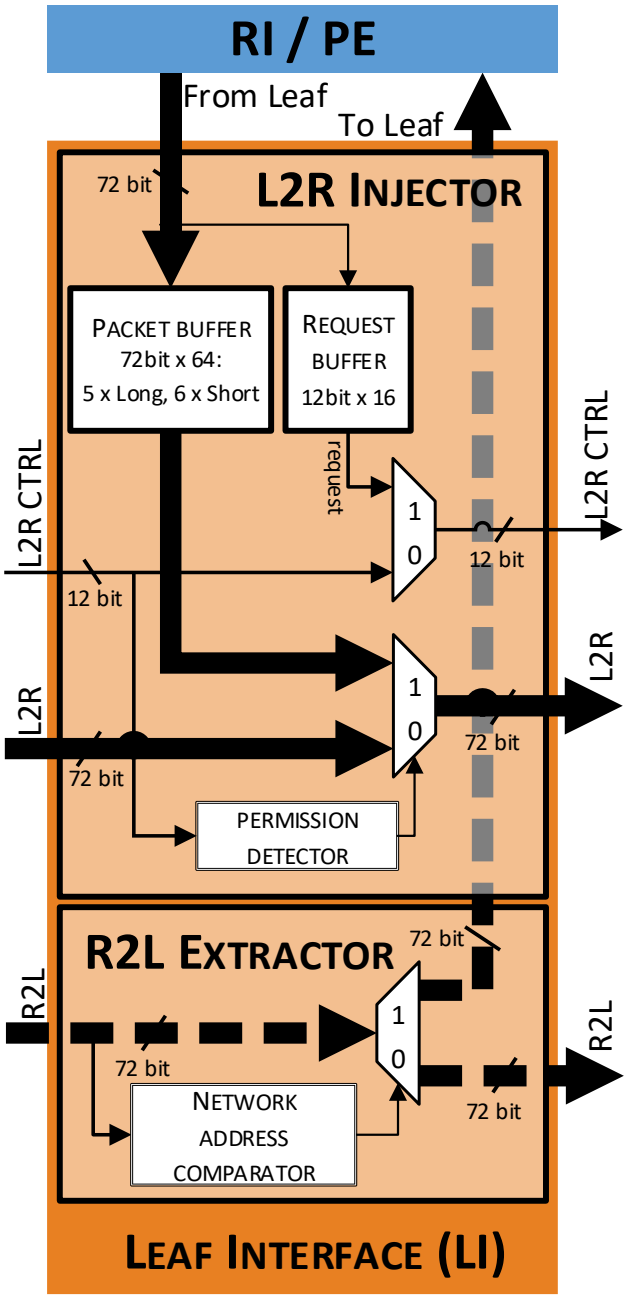

Fig. IV.3. Leaf Interface.

The LI is divided into L2R Injector part and R2L Extractor part. The L2R Injector buffers packets, send requests, and insert buffered packets onto the L2R ring after acquiring permission. The LI can buffer multiple packets while waiting for the permissions. The packet buffer used in LIs and depicted in Fig. VI.9, utilizes 64-words deep LUTRAM. It provides enough capacity for 6 short packets and 5 long packets. The request buffer utilizes 16-words deep LUTRAM. The requests are stored in the request buffer. Each of the requests corresponds to one packet stored in the packet buffer. When a permission reaches the Leaf Interface, it carries a Request

ID number that identifies the corresponding request from the request buffer and the buffered packet.

The R2L Extractor part of the LI extracts a packet from the R2L channel or bypasses the packet to next LI on a ring. The respective decision is taken according to the addressing information encapsulated in the packet header. In contradiction to the L2R Injector part of the LI, the L2R Extractor part has no buffer for packets. The extracted packet will be accepted by the attached processing element (PE) or the attached Root Interface, depending on which one is connected to the LI. If a PE is connected to the LI, it should accept the packet as a result of the primary idea of the RingNet, expressed in Section V.A.1 (a processing element controls the traffic load at a R2L channel, therefore, the packet has been requested by the PE). If, instead of a PE, a RI of higher network level is connected to the LI, this RI provides a small buffer used for synchronizing the extracted R2L packet to slots on the higher ring level.

C. Root Interface

Block diagram of the RI is depicted in Fig. IV.4.

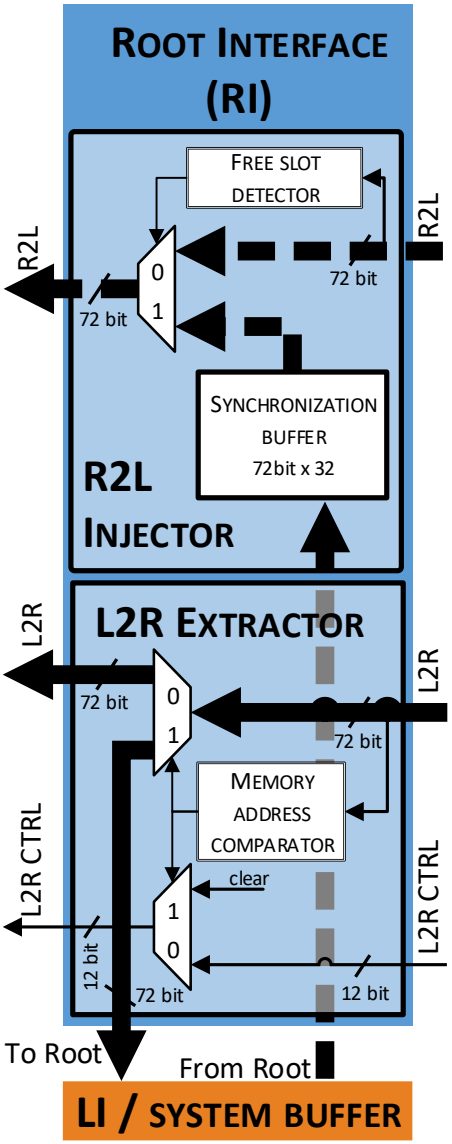

Fig. IV.4. Root Interface (RI).

The R2L Injector part of the RI accepts packets from the attached device (LI or system buffer) and stores them in the internal synchronization buffer. The synchronization buffer is used for synchronizing the packet with slots at the R2L channel.

The L2R extractor part of the RI identifies packet that should be extracted from the L2R channel, and checks if the attached LI or system buffer has buffer space available for the packet. If the buffer space is available then the packet is extracted. Otherwise, the RI marks the packet as rejected and bypass it to the L2R channel. The packet marked as rejected is recognized by

the L2R Manager and appropriate mechanism, described in Section VI.A is initialized with the aim of extracting the rejected packet.