



UNIVERSITÀ
degli STUDI
di CATANIA

DIPARTIMENTO di INGEGNERIA ELETTRICA,
ELETTRONICA e INFORMATICA

Low Power Techniques for Future Network-on-Chip Architectures

A thesis submitted for the degree of
Doctor of Philosophy in Ingegneria dei Sistemi, Energetica,
Informatica e delle Telecomunicazioni
XXIX Ciclo

Author

Andrea Mineo

Advisor:

Chiar.mo Prof. Vincenzo Catania

Co-Advisors:

Chiar.mo Prof. Maurizio Palesi

Chiar.mo Prof. Giuseppe Ascia

Coordinator:

Chiar.mo Prof. Paolo Arena

October 2016

“A mia moglie, Adriana”

“A mio nonno, Andrea”

“A mia zia, Pina”

“A mia madre, con affetto”

“A mio Padre”

Contents

1	Introduction	1
1.1	The Network-on-Chip Design Paradigm	2
1.1.1	Power Consumption	3
1.1.2	Scalability Issues	4
1.2	The Future of Wires	5
1.3	Scope and Contributions of the Thesis	8
1.4	Organization of the Thesis	9
2	Data Encoding Techniques in NoC Architectures	11
2.1	Link Energy Consumption	11
2.2	Low Power Coding	13
2.3	Duplicate Add Parity Code	14
2.4	DAP Codec	14
2.4.1	Voltage Swing Reduction	14
2.4.2	DAP on a Network-on-Chip	16
3	Reliability Aware Adaptive Voltage Swing Scaling	18
3.1	Probabilistic CMOS Technology	19
3.2	The Idea at a Glance	20
3.3	Contribution	21
3.4	Limitations and Applicability	22
3.5	Architectural and Microarchitectural Design	23
3.5.1	Reconfigurable Link with Duplication	23
3.5.2	Reconfigurable Link without Duplication	24
3.5.3	With Duplication <i>vs.</i> Without Duplication	25
3.5.4	Impact on the IC Design Flow	27
3.5.5	Control Circuitry	28

3.6	Design Configurations	29
3.7	Synthesis Results	30
3.8	Experiments	34
3.8.1	Energy Saving <i>vs.</i> QoS	36
3.8.2	Energy Saving <i>vs.</i> Link Length	37
3.8.3	Energy Saving <i>vs.</i> Packet Size	38
3.8.4	Energy Saving <i>vs.</i> Different Data Types	38
3.8.5	Energy Saving <i>vs.</i> Performance Degradation	39
3.8.6	Case Studies	40
3.9	Conclusions	45
4	Emerging Network-on-Chip Paradigms	48
4.1	3D NoC	48
4.1.1	3D Symmetric NoC	49
4.1.2	3D NoC-Bus Hybrid Architecture	50
4.1.3	Multi-layer 3D NoC Router Design.	51
4.2	Photonic NoCs	51
4.3	Wireless NoCs	52
4.3.1	Mesh-Topology Based WiNoCs	54
4.3.2	Small-World Network Based WiNoCs	55
4.3.3	Physical Layer Management	59
4.4	Comparative Analysis	62
5	Tunable Transmitting Power for WiNoC Architectures	63
5.1	Adaptive Transmitting Power Transceiver	64
5.1.1	Variable Gain Amplifier Controller	64
5.1.2	Determining the Minimal Transmitting Power	65
5.1.3	Overall Flow	68
5.1.4	The Mapping Problem	69
5.2	Experiments	71
5.2.1	Bandwidth and Radiation Pattern	72
5.2.2	Attenuation Maps	73
5.2.3	VGA Controller Analysis	75
5.2.4	Energy Saving in Mesh-Topology Based WiNoCs	78
5.2.5	Energy Saving in Small-World Network Based WiNoCs	81
5.2.6	Application Mapping	83

5.2.7	Case Study	84
5.3	Conclusions	84
6	Exploiting Antenna Directivity in WiNoCs	88
6.1	Antenna Directivity Optimization	89
6.1.1	Antenna Directivity	89
6.1.2	Formulation of the Problem	91
6.1.3	General Design Flow	92
6.2	Experimental Results	93
6.2.1	Simulation Methodology	93
6.2.2	Energy Saving Analysis	95
6.2.3	Case Study	98
6.3	Conclusions	99
7	Smart Transceiver for WiNoCs	101
7.1	SiESTA Power Reduction Strategy	102
7.1.1	Radio-Hub Architecture Overview	102
7.1.2	Radio-hub data flows	104
7.1.3	Detecting the <i>Radio Event Sleep</i> status	105
7.1.4	TX Power Management	107
7.1.5	RX Power Management	108
7.1.6	Hardware Implementation	109
7.2	Experiments	111
7.2.1	Simulation Setup	111
7.2.2	Effect of Packet Size and Packet Injection Rate	113
7.2.3	Effect of Flit Size	117
7.2.4	Effect of Buffers Size	121
7.2.5	Wireless Communication Energy Saving	122
7.2.6	Assessment under Real Traffic Scenarios	123
7.3	Conclusions	124
8	Conclusion	125
	Appendices	128
A	Scientific Production	129

Abstract

In a multi-many/core system, the Network-on-Chip (NoC) based communication backbone is responsible for a relevant fraction of the overall energy budget. In fact, the I/O buffers, the crossbars of the routers and the inter-router links are the main contributors of the NoC's energy dissipation. Specifically, electrical links will soon represent a bottleneck both in terms of energy dissipation and delay. For these reasons, several short and long terms solutions have been proposed from the NoCs research community. In particular, several techniques are based on reducing the voltage swing in links resulting in significant energy saving. Unfortunately, as voltage swing reduces, the bit-error-rate increases, that in turn compromises the communication reliability. Within this context, starting from the assumption that not all the communications need same level of reliability, in this dissertation we propose techniques and architectures for run-time tuning of the voltage swing of inter-router links. The proposed technique, is compared with the state of the art in link energy reduction through data encoding under both synthetic and real traffic scenarios. We found that the proposed techniques allow to significantly reduce the energy consumption of the NoC fabric without degrading the performance metrics. Energy savings ranging from 20% to 43% have been observed without any relevant impact on the performance metrics.

Since proposed short terms solutions will not longer satisfy the ever more aggressive energy consumption constraints, especially when the number of integrated processing elements (PEs) will be over the thousandths (as predicted by ITRS) new emerging paradigms have been proposed. In particular, wireless networks-on-chip (WiNoCs), have been recently proposed as candidate solutions for addressing the scalability limitations of conventional multi-hop NoC architectures. In a WiNoC, a subset of network nodes, namely, radio hubs, are equipped with a wireless interface that allows them to wire-

lessly communicate with other radio hubs. Thus, long-range communications, which would involve multiple hops in a conventional wireline NoC, can be realized by a single hop through the radio medium. Unfortunately, the energy consumed by the RF transceiver into the radio hub (i.e., the main building block in a WiNoC), and in particular by its transmitter, accounts for a significant fraction of the overall communication energy. In order to alleviate such contribution, two techniques have been proposed in this thesis. A first solution consists in a runtime tunable transmitting power technique for improving the energy efficiency of the transceiver. The basic idea is tuning the transmitting power based on the physical location of the recipient of the current communication. Specifically, based on the destination address of the incoming packet, the radio hub tunes its transmitting power to a minimum level, but high enough to reach the destination antenna without exceeding a certain bit error ratio. The proposed technique applied on different representative WiNoC architectures results in an average transmitter energy reduction up to 50% without any impact on performance and with a negligible overhead in terms of silicon area. A second solution focuses on the impact of antennas orientation on energy figures and performs a design space exploration for determining the optimal orientation of the antennas in such a way to minimize the communication energy consumption. When the antennas are optimally oriented, up to 80% transmitter energy saving has been observed.

Unfortunately, energy consumed by WiNoC transceiver does not depend by the transmitter but also by other modules including the receiver. In this sense, in order to obtain a further energy reduction in this thesis we propose a technique based on selectively turning off, for the appropriate number of cycles, all the radio-hubs that are not involved in the current wireless communication. The proposed energy managing technique is assessed on several network configurations under different traffic scenarios both synthetic and extracted from the execution of real applications. The obtained results show that, the application of the proposed technique allows up to 25% total communication energy saving without any impact on performance and with a negligible impact on the silicon area of the radio-hub.

Chapter 1

Introduction

Moore's Law ¹ has powered mainstream microelectronics for the past decades promising lower costs, lower power consumption and higher performance in terms of computational speed. Following the same trend, modern applications such as high performance computing, mobile computing and applications related to the Internet of Things (IoT), require higher and higher performance especially in terms of speed and energy efficiency. Unfortunately, in the ultra-deep sub-micron (UDSM) era, power and integration density in CMOS technologies do not scale at the same pace. In fact, as integration density increases by 2x every two years, power efficiency increases by only 1.4x. This trend is known as the end of Dennard's Scaling or as the begin of the *dark silicon* era [93]. The latter means that, mainly due to power dissipation issues, modern Systems-on-Chip (SoCs) cannot be anymore designed for operating at the maximum frequency while exploiting the entire chip resources at the same time. For facing such disruptive trend, in the last decade modern System-on-Chip are migrating from a single core (operating at high frequency) to a multicore architecture (with a lower operating frequency). In fact, nowadays, modern SoCs integrate ever more intellectual properties/processing (IPs/PEs) on the same chip. Following this trend, several chip makers, including Intel, AMD and TILERA, have already released commercial multi/many-core products. For instance, AMD has recently released the first native eight core processor for the desktop market [9], while TILERA and Intel have released a 72-core and a 60-core coprocessors, respec-

¹Moore's Law predicted that the number of transistors on a chip doubles every 18 to 24 months




Figure 1.1: Number of integrated processing elements inside a System-on-Chip [4].

tively [94, 23]. On the research side, Intel developed two prototypes in [98] and [21]. The former, developed in 2008, integrates 80 processing cores in a 65 nm CMOS technology, while the latter, after 6 years, integrates 256 cores in an 22 nm Tri-Gate CMOS technology. Fig. 1.1, reports a forecast on the number of integrated PEs over years showing quite clearly that such number is expected to reach the thousand within the 2020.

1.1 The Network-on-Chip Design Paradigm

Although the potential computational capabilities of many-core systems improve as the number of cores increases, on the other hand the on-chip communication subsystem becomes the actual bottleneck for what it concerns scalability, energy efficiency, performance and reliability of the entire system. The Network-on-Chip (NoC) design paradigm [12] is currently considered as the most viable solution for sustaining the communication demand of modern many-core architectures and for addressing the technological challenges in UDSM silicon nodes. As depicted in Fig. 1.2 a NoC is constituted by several processing elements (PEs) connected by a switched packet network made of switches (or routers) and electrical point-to-point links. Into the network, packets are fragmented in several small data units named *flits* which are stored and routed (according to a routing algorithm) by means of the main element of a NoC, namely, router. Due to severe timing constraints,

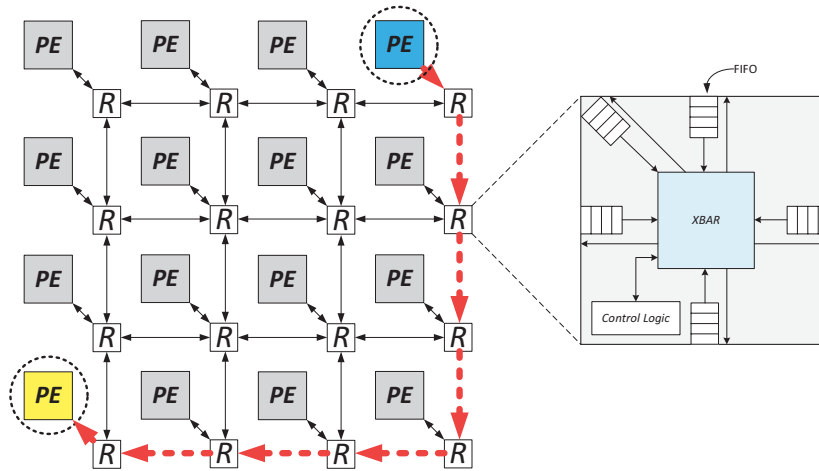


Figure 1.2: A 2D-Mesh NoC.

the latter must implement several pipeline stages. For this reason, each flit is routed through the network in a multi-hop fashion. In the picture mentioned above, for instance, a processing element (in blue) requires 7 hops to send a flit to the yellow PE.

1.1.1 Power Consumption

During the past years, several NoC architectures have been proposed by performing extensive design-space exploration involving topologies, routing algorithms, switching techniques, micro-architectural parameters, *etc.* From these studies it has been established that, although in the past the computation and the memory subsystems were considered as the main responsible of energy consumption, in NoC-based multi-/many-core architectures the communication subsystem accounts for a relevant fraction of the overall energy budget. For instance, in the Intel's 80-tiles TeraFLOPS processor [98], the communication power (due to routers and links) accounts for about 30% of the total power. Experimental results in [37] shown that the contribution of the interconnection network to the total power ranges from 15% (8-core tiled CMP) to 35% (16-core tiled CMP) on average, with some applications reaching up to 50%. In the AETHERIAL NoC the largest percentage of power dissipation (54%) is due to the NoC clock, followed by the NoC links (18%) [88]. In [33], it has been shown that on-chip interconnects account for a significant fraction (up to 50%) of the total on-chip energy consumption.




Figure 1.3: Router's power breakckdown [98].

Entering in more details, the power dissipated by a NoC depends on both physical wires and router internal components. As depicted in Fig. 1.2, the latter is essentially constituted by several FIFO (First In, First Out) memories to store incoming flits, a control logic and by a crossbar which is in charge of routing packets to a specific output port. In order to understand how the power is distributed into a router, Fig. 1.3 shows, for instance, the power breakdown of the Intel TeraFLOPS [98]. From this graph is clear that power consumption is dominated by FIFOs (22%), crossbar (15%) and electrical links (17%). For this reason, the power consumption of a router could be simplified as follows:

$$P = P_{fifo} + P_{xbar} + P_{link}. \quad (1.1)$$

From Eqn. (1.1) results clear now that, for the sake of power minimization a NoC designer should act on such three contributions. As it will be seen in the rest of the thesis, our research is mainly focused on such techniques or emerging NoC paradigms that try to optimize both delay and power metrics on physical links which are a bottleneck for today's and future NoC based SoCs.

1.1.2 Scalability Issues

In the NoC context is better thinking in terms of energy consumption rather than power dissipation. Starting from Eqn. (1.1) we can express the energy consumed to transfer a piece of information (a flit) in function of the number hops of the route between the source and the destination node:

$$E_{flit} = n \cdot (E_{fifo} + E_{xbar} + E_{link}) = n \cdot E_{router} + n \cdot E_{link} \quad (1.2)$$

where n is the number of routers and links crossed during a transfer between a source and destination node. The worst case happens when the number of hops equals to the network diameter (based on the position of source and destination). Considering the network depicted in Fig. 1.2 it results trivial computing the network diameter D as reported in the Eqn. 1.3.

$$D = 2 \cdot (\sqrt{N} - 1). \quad (1.3)$$

Since the latter depends by the number of integrated processing elements N , worst case energy can be finally computed as follows:

$$\max(E_{flit}) = D \cdot E_{hop} = 2 \cdot (\sqrt{N} - 1) \cdot E_{hop}. \quad (1.4)$$

From this last result emerges that when the number of cores increase energy consumed to transfer data from distant PEs increases as well.

The same behavior is valide for latency too. Due to pipeline structure of a router we can also compute latency in terms of the network diameter D and the number of pipeline stages $N_{pipeline}$ as follows:

$$\max(Latency) = N_{pipeline} \cdot D = 2N_{pipeline} \cdot (\sqrt{N} - 1) \quad (1.5)$$

If we consider, for instance, a 3 pipeline stages router and 100 cores the worst case latency results in 54 cycles. Thus, following the ITRS predictions as reported early, it results clear that around 2020 both latency and energy consumption will be a real issue for future NoC-based SoCs.

1.2 The Future of Wires

In this section we would to show how metallic/dielectrical wires will be a bottleneck, especially in terms of delay. In order to better understand this phenomena, a simple model to compute parasitics effects will be exposed. As it will be seen, different effects emerging in future technologies make metallic wires the weak point of today's NoC base systems.

In NoC-based SoCs, each point-to-point link consists of a bus made of n wires, where n is equal to the size of a flit. In fact, such flits are transmitted over the network in a parallel fashion through dedicated buses. Metallic wires are further arranged with minimal spacing in order to save area. Since in digital integrated circuits both delay and power consumption are proportional

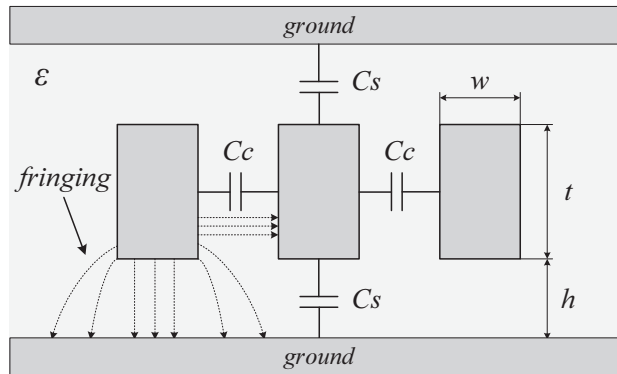


Figure 1.4: Stripline Model.

to parasitic capacitances introduced by devices and electrical wires [81, 44], it is useful to know how to compute distributed capacitance from technological parameters. From electromagnetic theory, to compute such parasitics in CMOS metal/dielectric wires, a strip-line model should be used [69]. Fig. 1.4 shows a model of a generic bit-line inside a bus in which capacitance is thus modeled by four parallel-plate capacitors for the top, bottom, right, and left sides, plus a constant term for fringing capacitance. In modern technologies, vertical and horizontal capacitors may have different relative dielectrics using low- k materials.

The “far” plates for the top and bottom capacitors are typically modeled as being grounded, since they represent a collection of orthogonally routed conductors that, averaged over the length of the wire, maintain a constant voltage. Capacitors to the left and right, on the other hand, have data-dependent effective capacitances that can vary: if the left and right neighbors switch in the opposite direction as the wire, the effective sidewall capacitances double, and if they switch with the wire, the effective sidewall capacitances approach zero. This effect is known as “Miller multiplication”. These left and right neighbors are also the worst offenders for noise injection. The fringe term depends weakly on geometry and for today’s technologies is about $40fF/mm$. For the very top layers of metal with no upper layers, we can use three parallel plates with extra fringing terms on the two horizontal capacitors.

In order to estimate the capacitance per unit length we can thus write:

$$c_{dist} = \epsilon_0 \epsilon \left(2K \frac{t}{spacing} + 2\epsilon \frac{w}{h} \right) + fringe(\epsilon) \quad (1.6)$$

Where, ϵ_d is the dielectric constant assumed to be homogeneously distributed both between layers and between metal lines within a layer, ϵ_0 is the permit of free space, K takes in accounts of the Miller multiplication (varying from 0 to 2), while t , w and h are geometrical parameters as shown in Fig. 1.4.

A convenient simple model, enough for first-order hand calculation can be derived from the above equation by means of the following assumptions:

- The interlevel dielectric thickness h is assumed to be the same as the metal thickness t .
- Intralevel dielectric thickness and wire width w is assumed to be half of the pitch ($spacing + w$).
- An average dielectric constant value is used for the case where a range of values is suggested in the roadmap for a given technology node.
- The capacitance values represent the worst switching scenario when two adjacent wires, on the same level, are simultaneously switching in the apposite direction as the signal line, hence doubling the intralevel capacitance contribution.

The former computed capacitance per unit length is thus given by Eqn. 1.7.

$$C_{dist} = 2\epsilon_d\epsilon_0 \left(\frac{1 + 2AR^2}{AR} \right) + fringe(\epsilon) \quad (1.7)$$

where AR is the aspect ratio of the wire defined as the thickness to width ratio of the metal. The introduction of such constant is particularly useful to understand the capacitances trend for future technology steps.

In fact, ITRS [3] predicts that aspect ratio is destined to increase, and thus, from Eqn. 1.7 is clear that in the UDSM regime parasitics introduced by wires id destined to increase too. Consequently, as technology shrinks, the ratio between the wire capacitance and input gate capacitance increases [47] making the interconnection network the principal actor with respect to the energy/power metric. This particular aspect can be also viewed in terms of delay. In fact, as shown in Fig. 1.5, for each technology step delay due to logical gate (in terms of FO4) decreases while wire delay increases even with repeaters insertion. It means that, for future technologies the wire

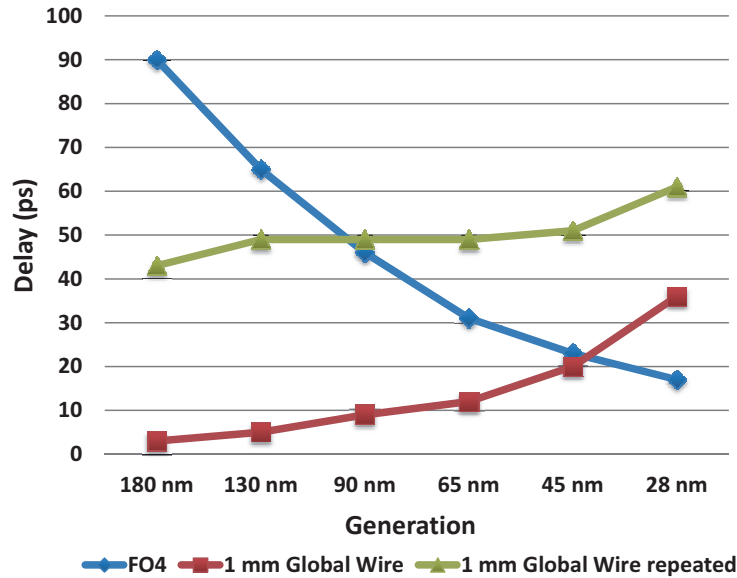


Figure 1.5: Gate and wire scaling [47].

contribution will be the real bottleneck both in terms of speed and energy efficiency.

1.3 Scope and Contributions of the Thesis

For facing scalability problems discussed before, several solutions have been proposed in literature. In particular, in this dissertation, we aim to minimize the negative impact in terms of energy consumption of a NoC by proposing both short and long term solutions. The research presented here falls in two topics on minimizing energy in modern NoC based systems: data encoding techniques and an emerging interconnect paradigm named Wireless Networks-on-Chip (WiNoCs). While the former represents a short term solution the latter are based on new technologies beyond traditional CMOS processes. Both topics and related proposed solutions will be briefly introduced in the rest of this section.

1. Data Encoding Techniques in Network-on-Chip architectures.

According with these techniques, energy consumption could be minimized by reducing possible responsible of power consumption such as the switching activity or voltage swing in physical lines. In this thesis it is proposed a technique for minimizing the energy consumption by

reducing voltage swing. Unfortunately, reducing the voltage swing has a negative impact on the signal integrity (i.e., on the content of the packets traveling in the NoC) which results in a reduction of the communication reliability due to the increase of the bit-error-rate (BER). As for the computation counterpart in which the concept of probabilistic CMOS has been recently investigated, [71], we believe that, in the context of communication, it is not always necessary to ensure that the BER must be the same for all the transmitted packets. Based on this, we propose to extend the general concept of probabilistic CMOS to the communication side of a NoC based system. In particular, the concept is applied to two of the main components of a NoC architecture which mostly affect the NoC energy, namely, links and crossbar. We present link architectures able to be configured on-line for working at two different voltage swings, namely, full and low-voltage swing. The basic assumption is that not all the communications involved in an application have the same reliability requirements. Thus, the links and the crossbars traversed by the packets of a communication which does not have stringent reliability requirements are configured on the fly to work at low-swing voltage.

2. **Emerging interconnect technologies for on-chip networks.** To face with the scalability problems especially in terms of latency, new technologies such as 3D NoC, optical and wireless NoC (WiNoC) are emerging as technological alternatives to the wire based NoCs [15]. In particular a WiNoC constituted a wireless backbone upon the traditional wire-based NoC [30]. WiNoCs introduce new hardware structures such as antenna and transceivers that represent an overhead in terms of area and power. In this dissertation we present three techniques for reducing power consumption on the transceiver which is the most power hungry device of this kind of systems.

1.4 Organization of the Thesis

The remainder of the thesis is as follows. Before detailing the different contributions of our work, Chapter 2 and 4 introduce some concepts of Data Encoding Techniques and of emerging technology respectively.

Chapter 3 exposes a data encoding technique which reduces energy by varying, at runtime, the voltage swing on physical lines.

Chapter 5, 6 and 7 describe three proposed techniques, developed for Wireless Network-on-Chip paradigm. In particular as it will be seen, transceiver power consumption is predominant if compared to other WiNoC devices. Proposed techniques thus work on the reduction of such power. To do this, the technique exposed in the Chapter 5 reduces energy by varying the transmitting power of transmitter based on the position of source and destination. Technique proposed in the Chapter 6 provides a design flow to determine the optimum orientation of antennas in a WiNoC. Further, in the Chapter 7 we propose a mechanism to reduce energy consumption when a receiver is not a recipient of a message. Finally, Chapter 8 concludes the dissertation.

Chapter 2

Data Encoding Techniques in Network-on-Chip Architectures

As stated before, in this dissertation we will expose a method for reducing energy consumption due to physical lines of a NoC. To better understand the proposed technique, we start to show how to compute power dissipation on electrical wires. After that a background has been given, the proposed scheme will be finally discussed in the next chapter.

2.1 Link Energy Consumption

The average switching energy dissipated by a digital circuit can be estimated using the following formula [81]:

$$E_{dyn} = \alpha \cdot C_L \cdot V_{dd}^2, \quad (2.1)$$

where α is the switching activity of the nodes of the circuit, C_L is the total load capacitance, and V_{dd} is the supply voltage. From Eqn. 2.1 it can be observed that V_{dd} has a squared effect on the dynamic energy consumption.

A physical line can be seen as a distributed load for a driving gate. However, for the sake of power estimation, it can be assumed as a lumped load [81]:

$$C_L = c_l \cdot L,$$

where c_l is the the capacitance per unit length that can be computed as in the Eqn. 1.7 (pp. 7), and L is the line length.

Table 2.1: Effective capacitance, C_{eff} , for each type of transition on a victim and aggressors lines [80]. The \uparrow and \downarrow represents the versus of transition and the symbol ‘-’ is used for representing no transition in the specified line.

C_{eff}	Transition Pattern	p
C_s	$(\uparrow, \uparrow, \uparrow) (\downarrow, \downarrow, \downarrow)$	0
$C_s + C_c$	$(-, \uparrow, \uparrow) (-, \downarrow, \downarrow) (\uparrow, \uparrow, -) (\downarrow, \downarrow, -)$	1
$C_s + 2C_c$	$(-, \uparrow, -) (-, \downarrow, -) (\downarrow, \uparrow, \uparrow) (\downarrow, \downarrow, \uparrow)$ $(\uparrow, \uparrow, \downarrow) (\uparrow, \downarrow, \downarrow)$	2
$C_s + 3C_c$	$(-, \uparrow, \downarrow) (-, \downarrow, \uparrow) (\uparrow, \downarrow, -) (\downarrow, \uparrow, -)$	3
$C_s + 4C_c$	$(\uparrow, \downarrow, \uparrow) (\downarrow, \uparrow, \downarrow)$	4

For the physical lines, the average power dissipation also depends on the type of transition that occur on the current line and in the neighboring lines (contribution due to the crosstalk). Specifically, in terms of energy consumption, the Eqn. (2.1) can be written as follow:

$$E_{dyn} = C_{eff} \cdot V_{dd}^2, \quad (2.2)$$

where C_{eff} is the effective capacitance that not only depends on the transition in the current line, but also by the specific transition of the aggressor line. C_{eff} is often expressed as:

$$C_{eff} = C_s + p \cdot C_c, \quad (2.3)$$

where C_s is the line’s self capacitance, and C_c is the coupling capacitance. The factor p can be 0, 1, 2, 3, or 4 according to Tab. 2.1. The elements in the table are 3-tuples in which the central term represents the victim line (*i.e.*, the line under examination) while the left and the right terms represent the neighboring aggressor lines. If present, the arrow indicates the transition type (\uparrow rising transition, \downarrow falling transition). The symbol “-” represents the case in which no transition occurs in the current line. It should be pointed out that, if the victim line makes a transition while the neighboring lines remain steady (*i.e.*, $(-, \uparrow, -)$ or $(-, \downarrow, -)$), the effective capacitance is $C_s + 2C_c$, but from the noise point of view, there is not injection of charge in the victim line.

2.2 Low Power Coding

From previous section results clear that for reducing energy due to electrical wires (disposed in a bus fashion), a designer can act on three main constituting factors: by reducing the switching activity (α term), reducing effective capacitance (C_{eff}) or decreasing the voltage swing. For alleviating these terms, data that should be transmitted on a bus can be coded properly. Initially proposed for shared bus based System-on-Chips (SoCs), several coding techniques have been developed by research community [86]. In particular, we can distinguish among three main coding families:

- Low Power Coding (LPC): the power dissipation in the bus depends on data transition activity (α). This kind of techniques try to reduce such activity. Their effectiveness are relevant only when coupling capacitance between lines are negligible. Interesting research papers that fall in this topic can be found in [87, 72].
- Crosstalk Avoidance Coding (CAC): both power dissipation and delay of a wire in a bus depend to its transitions and to adjacent wires activity. From Eqn. 2.3, the worst-case delay of a wire is when p term is equal to 4. The purpose of the crosstalk avoidance coding is to limit such term to 2. This goal can be achieved if some transitions are forbidden by mean of a dedicated codec [77].
- Joint Crosstalk Avoidance and Forward Error Correction (CAC+FEC): a joint crosstalk avoidance and error-correction code can be obtained by combining a crosstalk avoidance code with an error-correcting code. p term is reduced from 4 to 2 and voltage swing can be scaled down by means of an error correction code.

Since in the UDSM regime coupling capacitance has a higher impact, both CAC and CAC+FEC techniques are more effective in terms of obtained energy saving if compared to low power codes. Further, from recent studies [42] results that CAC+FEC techniques are the best candidate for saving energy especially in the NoC context. Since the data encoding technique proposed in this thesis will be compared with several CAC+FEC techniques, a representative scheme named Duplicate Add Parity, will be described in the next section.

2.3 Duplicate Add Parity Code

As mentioned in the in the previous section, CAC techniques obtain an energy reduction by avoiding critical transitions on two adjacent physical lines. The technique *Duplicate Add Parity* obtain such requisite by doubling each bit-line of a bus. In this way, for each line, we have the original data and a copy of it. It should be noted that added lines act as shielding lines from the original data that should be transmitted on the bus. Furthermore, an error detection mechanism is introduced to detect transmission errors by mean of a parity bit. With this technique, an n bits of data become coded in a $2n + 1$ bits. Since line have been duplicated, with the introduction of such mechanism, if an error occurs on a specific bit-line, a copy of this data can be selected. The just gained robustness can be therefore exploited to scale down the voltage swing on the lines. Thus, according with the Eqn. 2.1 a squared effects on energy reduction will be achieved.

2.4 DAP Codec

Logical scheme of the codec which implements the DAP technique on a n bit bus is shown in Fig. 2.1. The encoder is constituted by a combinational block name *Parity*. The former computes the logical XOR between each bit-line as follows:

$$Parity = b_0 \oplus b_1 \oplus \dots b_{k-1}$$

With the introduction of additional line coming from this block, the number of bits 1 on the bus should be even for any words combinations. The decoder is thus implemented with the *Parity* block followed by a logical XOR which driver the selector of a series of multiplexers. Therefore, if an error occurs into the bus, the decoder will selects the right lane selected as data valid.

2.4.1 Voltage Swing Reduction

As mentioned before, the introduction of a error detection mechanism is translated in a higher robustness. The probability that an error occurs on a single bite, named BER (Bit Error Rate), is strictly tied to the signal to

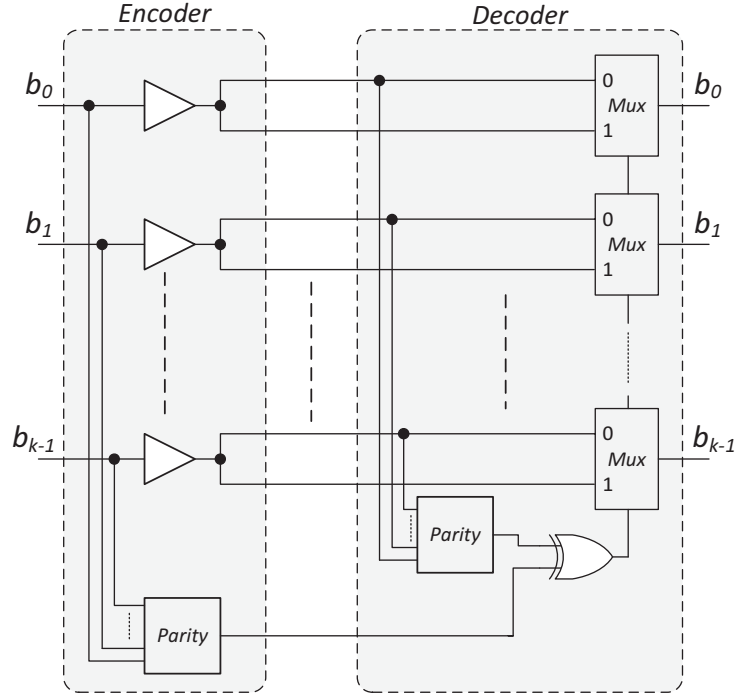


Figure 2.1: CAC+ECC (DAP) Codec.

noise ration presents on the line. Thus, to an higher voltage swing follows an higher errors immunity. The probability ϵ that a single error happens on a line (BER) can modeled as follows:

$$\epsilon = Q\left(\frac{V_{dd}}{2\sigma_N^2}\right) \quad (2.4)$$

where Q is a well known equation (Eqn. 5.4).

$$Q(x) = \frac{1}{\sqrt{2\pi}} \int_x^\infty e^{-\frac{y^2}{2}} dy \quad (2.5)$$

The probability of having an error on a word is a function of the BER for each single bit-line (ϵ). If we refer with $P_{\text{unc}}(\epsilon)$ the BER without any implemented error correction mechanism and with $P_{\text{ecc}}(\epsilon)$ the BER obtained after that coding is applied, we will have:

$$P_{\text{ecc}}(\epsilon) \leq P_{\text{unc}}(\epsilon) \quad (2.6)$$

or using the Eqn. 2.4 and by fixing $P_{\text{ecc}}(\epsilon) = P_{\text{unc}}(\epsilon)$, we can reduce the voltage swing as reported below.

$$\widehat{V}_{dd} = V_{dd} \frac{Q^{-1}(\hat{\epsilon})}{Q^{-1}(\epsilon)} \quad (2.7)$$

In the Eqn. 2.7, V_{dd} represents the nominal voltage without any correction while \widehat{V}_{dd} is a reduce swing that can achieve the same BER by means of the error corrections mechanism.

Without any correction the probability that an error occurs in to a word is tied to the ϵ as reported below:

$$P_{\text{unc}}(\epsilon) = k\epsilon \quad (2.8)$$

where k represent the number of bits presents on the word. If DAP strategy is applied, residual error probability can be computed like in [77] (reported in the Eqn. 2.9).

$$P_{\text{DAP}}(\epsilon) = 1 - \sum_{i=0}^k \binom{k}{i} \epsilon^i (1 - \epsilon)^{2k-1+i} - \sum_{i=0}^{k/2} \binom{2k+1}{2i+1} \epsilon^{2i+1} (1 - \epsilon)^{2k-2i} \quad (2.9)$$

The latter could be simplify only wor small ϵ values as reported in the Eqn. 2.10.

In literature are present other techniques to reduce voltage swing further by detecting more than just one error. In Fig. 2.2 is reported the voltage swing as function of the number of detected error.

$$P_{\text{DAP}}(\epsilon) = \frac{3k(k+1)}{2} \epsilon^2 \quad (2.10)$$

2.4.2 DAP on a Network-on-Chip

As discussed on previous sections, applying a DAP codec on a bus the voltage swing could be reduced. Since a NoC is constituted by point-to-point links, a encoder and decoder pair should be instantiated for each oh them. From this, it results clear that the price to pays for implementing such technique will be both in terms of silicon area and latency. In Fig.2.3 is depicted DAP codec on a traditional NoC switch. By adding the DAP codec, tho more pipeline stages ha been added. For facing the latency issue in the next section we propose a technique that reduces voltage on the lines without affects the pipeline structure of a router. As it will be seen, a trade-off between latency and bit error rate (BER) will be introduced.




Figure 2.2: DAP: Voltage swing with a BER of 10^{-20} following Eqn. 2.7 with $\sigma = 65$ mV.

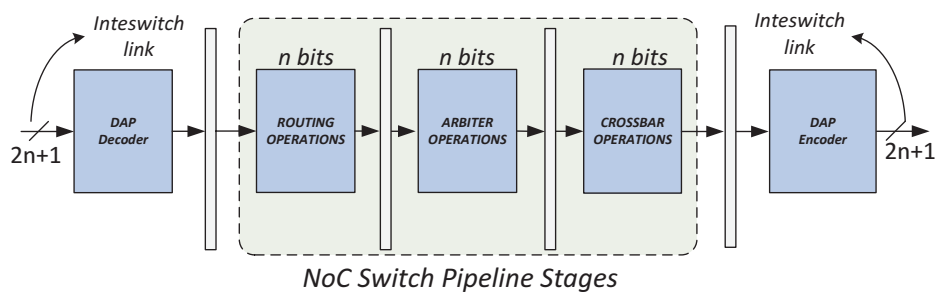


Figure 2.3: DAP codec implemented in a NoC router. Two pipeline stages have been added.

Chapter 3

Reliability Aware Adaptive Voltage Swing Scaling

As mentioned in the previous section, state-of-art data encoding techniques (such as DAP technique) are based on voltage reduction that can be obtained by mean of an error detection mechanism. Such techniques are, in fact, based on a sequence of actions that can be summarized as follows:

1. Decrease the link voltage swing.
2. The BER increases.
3. Encode data and/or use an error correction scheme.
4. The BER decreases to the nominal value (*i.e.*, that observed at nominal voltage swing).

Since performing operation described at the point 3 leads to a penalty both in terms of additional silicon area and delay (as shown in Sec. 2.4.2), in the research presented here we propose a techniques that stops at the point 2. We, in fact, start from the assumption that for some specific communications, a higher BER can be tolerated without the need of using any error detection/correction mechanism because the possible data corruption does not affect the functionality of the application but only the quality of its results. For instance, suppose a packet whose payload is a macroblock of a frame in a video application. Even if the macroblock reaches its destination affected by errors, it results in a degradation of the image quality (*e.g.*,

presence of spurious pixels in the image) but it does not affect the functionality of the algorithm that will process (use) the macroblock. Basically, the proposed technique exploits the trade-off between energy saving and communication reliability. Such communication reliability, in practical cases, results in a quality/accuracy/precision metric based on the underlying application. Conversely, data encoding techniques do not consider such trade-off. In fact, for them, all the communications must provide the same nominal BER and the maximum quality/accuracy/precision metric is assumed.

It should be pointed that, proposed idea is not new at all. In fact, on the the low reliability level in USDM CMOS devices (very susceptible to the process variation and to the perturbations due to the noise) has been recently exploited in the area of *probabilistic CMOS* for defining interesting trade-off between reliability and energy consumption [71].

Proposed technique tries to extend the general concept of probabilistic CMOS, which has been traditionally investigated in the context of computation, to the communication side of a NoC based SoC. Specifically, we present link architectures able to be configured on-line for working at two different voltage swings, namely, full and low-voltage swing. Thus, the links traversed by the packets of a communication which does not have stringent reliability requirements are configured on the fly to work at low-swing voltage. Conversely, if a communication has an higher constraint in terms of required BER, electrical links will be configured to work at the nominal voltage.

Experiments carried out on both synthetic and real traffic scenarios show the effectiveness of the proposed technique in terms of energy saving. Differently from the state-of-the-art in link energy reduction through data encoding schemes, which have a negative impact on the communication latency, the proposed technique provides higher energy saving without impacting the performance metrics.

3.1 Probabilistic CMOS Technology

In the USDM era the device scaling faces several hurdles. Such devices are more and more susceptible to the process variation and to the perturbations due to the noise. For this reason, in [71] the authors introduced the concept of the *probabilistic CMOS* devices or PCMOS for short. Introducing the

probabilistic devices, the ordinary boolean functions has been substituting with the probabilistic switching function in which the output results are not deterministic but have a given probability of correctness p to obtain the ordinary operations.

Principles of statistical thermodynamics may be applied to such devices to quantify their energy consumption. While a deterministic switch consumes at least $KT \ln(2)$ (with T the temperature and K the Boltzmann constant) Joule of energy, a probabilistic switch can realize a switching functions with $KT \ln(2p)$ Joule of energy [8]. As in the standard static-CMOS circuitry the energy consumption depends by the supply voltage, several research works [71, 8] use the voltage scaling to obtain a squared reduction of energy consumption introducing a tradeoff with the operation's correctness [70]. In fact, the introduction of the concept of probabilistic CMOS allows to split the applications domain in applications which can tolerate the probabilistic behavior (such as multimedia applications) and applications that can even benefit (or harness) from probabilistic behavior at the devices level naturally such as Bayesian inference [55], Probabilistic Cellular Automata [39], and Hyper Encryption [31]. Interested reader in [46] can find more recent developments in this research field. For the best of our knowledge, the concept of probabilistic CMOS has been exploited only for the computation. In the research presented here, we want to exploit the modularity and the flexibility of the NoC design paradigm for extending the concept of probabilistic CMOS to the communication side of the system.

3.2 The Idea at a Glance

The basic idea of the proposed technique can be summarized as follows. Let us suppose that the communication system makes it available two different kinds of communication channels, namely, the *default* (δ) channels and the *low power* (λ) channels. A δ channel uses the nominal voltage swing whereas a λ channel uses a lower voltage swing. The flits of the packets of communications travel on δ channels as usual. However, the flits of the packets of communications which do not have stringent reliability requirements, will use the λ channels with a consequent reduction of energy consumption.

The implementation of this mechanism requires just a single bit of infor-




Figure 3.1: General scheme implementing the proposed idea.

mation to be stored in the head flit of the packet, namely, the *robustness* flag (R). When R is set, all the flits of the packet will use only δ channels, whereas when R is not set, the body flits of the packet will use the λ channel (Fig. 3.1). It should be pointed out that, head flits are always transmitted on δ channels. This is because head flits carry critical information (*e.g.*, destination address) that cannot tolerate errors.

A general way for modeling at high level of abstraction an application to be mapped on a multicore SoC is by means of its communication graph (CG). A CG is a graph whose vertices represent the tasks and edges represent the communications. The edges of the CG are usually annotated with attributes which characterize that communications. Typical attributes are the average communication bandwidth and the traffic volume. In this research, we assume that the edges of the CG are annotated with an additional information, namely, the *robustness*. With the term robustness of a communication, c , we indicate the chance for c to be not affected by communication errors. Precisely, we consider two robustness levels denoted as *low* and *high*. A communication with a robustness *low* admits that the BER is higher than the nominal case. Conversely, a communication with a robustness *high* means that such communication expects a nominal BER. Thus, the robustness flag of the packets belonging to a communication is set in accordance with the robustness attribute of that communication.

3.3 Contribution

Reducing the voltage swing of NoC links is not a new technique in the context of energy efficient NoC design. However, the contribution of our proposed

technique is not simply the selective reduction of voltage swing of NoC links. Differently from the previous approaches (chapter 2), the NoC link voltage swing is not only selective but also dynamic. That is, the voltage swing of a link is dynamically tuned based on the nature of the packets currently transmitted.

For the best of our knowledge, all the reviewed proposals in the context of NoC link energy reduction which exploit the possibility of reducing the voltage swing with a consequent trade-off in terms of reliability differ from the proposed scheme for an important and essential point. All the proposal discussed in the previous section assume that data transmitted through the NoC must be delivered to the destination cores unaffected by any error. Conversely, our contribution exploits the concepts of probabilistic CMOS in which there could be circumstances for which receiving data affected by errors does not have any impact on the functionalities of the system but they only affect some quality indexes that sometimes can be tolerated (*e.g.*, image or audio quality in a multimedia application). Based on this, the state-of-the-art in data encoding makes use of error detection/correction mechanisms to guarantee a certain reliability level (usually expressed in terms of a maximum allowed BER) whereas our approach does not make use of any error detection/correction mechanism with a consequent reduction of complexity, cost, power and performance as it will be shown in the experimental section (Sec. 3.8).

3.4 Limitations and Applicability

The proposed scheme, to be used, needs the availability of a communication graph annotated with reliability information. Although for some applications it might be simple determining the reliability level of a certain communication, in general cases it is not a trivial task. In this research we do not address the important issue on how to map the communications involved in an application as reliable or unreliable. Such important issue is left for future work. The goal of our work is providing a mechanism which enables the user (*i.e.*, the application developer) to improve the energy efficiency of the application whenever he/she is able to classify the communications determined by the application as reliable or unreliable.

Further, it is assumed that the software communication library provides a new send primitive which exposes an additional parameter that allows to specify the reliability level of the current transfer. For instance, we assume the availability of a send primitive like `send(dst, data, robustness)`. The *robustness* parameter is a boolean value that allows the programmer to specify whether the *data* transmitted to *dst* will use a reliable but power hungry links or an unreliable but low power links. In fact, the robustness flag is simply mapped into the robustness bit of the packet.

3.5 Architectural and Microarchitectural Design

The proposed technique discussed in the previous section is now elaborated and several implementations are analyzed. First we describe two runtime reconfigurable link architectures. In the first implementation, the δ and λ channels are mapped into two different physical links, namely, the δ and the λ links which work at nominal and low voltage swing, respectively. The overhead due to the duplication of the physical links of the first implementation is overcome by the second implementation in which a single physical link is used for both δ and λ channels.

3.5.1 Runtime Reconfigurable Link Architecture with Duplication

The proposed link architecture is shown in Fig. 3.2. The transceiver for the single bitline [Fig. 3.2(a)] is formed by two lines working at high voltage swing (δ) and low voltage swing (λ), respectively. A demultiplexer (demux) is used to transmit the current bit over the δ line or λ line based on a selection command stored in the head flit. Several techniques can be used to implement the low swing driver [107]. The low swing signal has to be converted back to a high swing signal before entering into the multiplexer (mux). Such conversion is performed by the module denoted as *level restorer* [107] in the figure.

Fig. 3.2(b) shows the proposed link architecture. As it can be observed, the δ and λ lines are interleaved. Such organization provides an additional

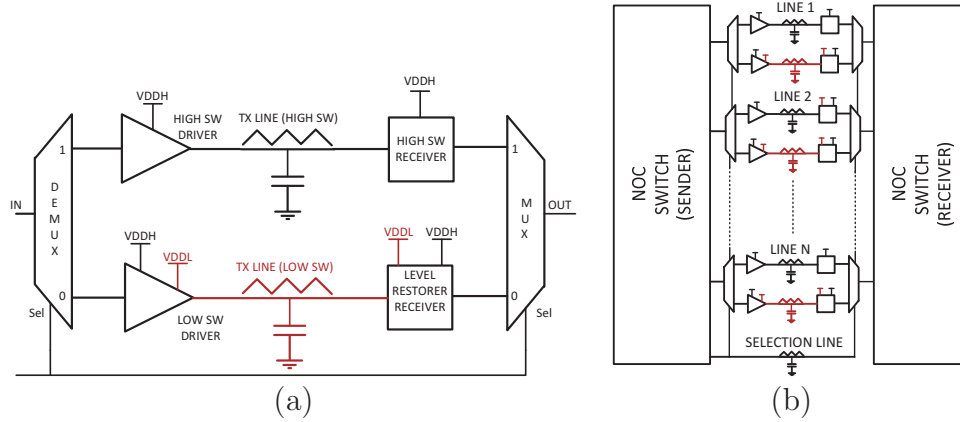


Figure 3.2: Link architecture with line duplication. Bitline (a), link (b).

positive effect due to the fact that, in a given clock cycle, if δ lines are active, λ lines are inactive and viceversa (with the term *active lines* we mean the lines currently involved in a transfer). It means that, an active line will have two neighboring inactive lines. Such inactive lines can be seen as a shield for the current line. That is, there is no switching in the neighboring lines causing $p = 2$ in Eqn. (2.3). In addition, such organization, can increase the throughput of the line due to the fact that the Miller multiplication [47] is avoided. Also, it results in less noisy links since crosstalk effects are reduced.

It should be pointed out that, the proposed circuitry remains valid even when the insertion of repeaters becomes necessary. In such case, the repeaters of the δ and λ lines will be supplied with the nominal and low voltage, respectively.

3.5.2 Runtime Reconfigurable Link Architecture without Duplication

Although the link architecture discussed in the previous subsection has the advantage of i) reducing the crosstalk effects and ii) reducing the factor p in Eqn. (2.3) to 2, it has the drawback of duplicating the bit lines. Let us now present another link organization which does not require bit lines duplication.

The scheme of the single link architecture is shown in Fig. 3.3(a). As it can be observed, the bitline is preceded by a chain formed by a demultiplexer, two tapered buffers as line drivers and two tristate buffers. The low swing conversion can be situated within the demultiplexer or in the tapered buffer

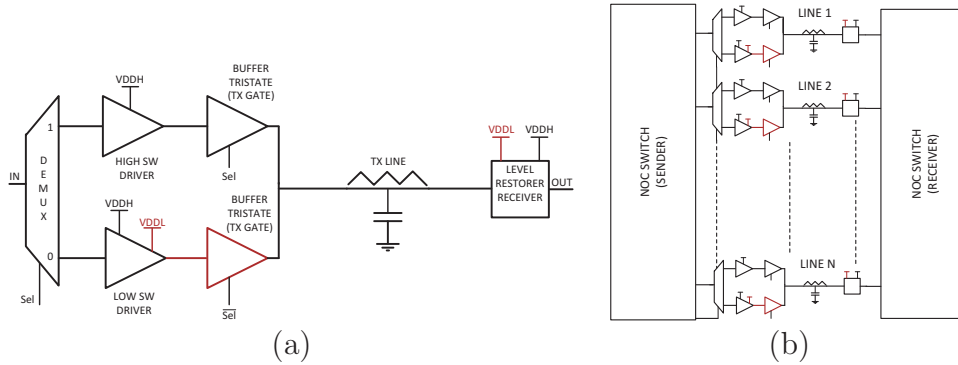


Figure 3.3: Link architecture without line duplication. Bitline (a), link (b).

as in one of many solutions proposed in [107]. Note that, if the demultiplexer is well sized, the introduction of a new stage not only does not affect the delay, but instead has a positive effect on performance if the proper number of inverter stages that compose the tapered buffer is chosen [89]. The tristate buffers are based on transmission gate logic. With this solution, if the select input is high (low), the full (low) swing path is active and the low (high) swing path is disconnected by the high impedance state introduced by tristate buffer. The level restorer circuit is similar to the sense amplifier used in RAM memories. It restores the signal at full swing if the signal on the line is set to low swing, or maintain the original swing if the signal is in full swing mode. In particular, when designing the level restorer, the logic threshold has to be set at half of low swing. In many cases, if the level restorer is well optimized, the speed of the circuits will be improved. Several circuitual solutions for the design of the level restorer can be found in [107].

3.5.3 Runtime Reconfigurable Link Architecture: With Duplication *vs.* Without Duplication

In the previous two subsections, two runtime reconfigurable link architectures, namely, with duplication and without duplication, have been presented. In this subsection, we compare them in qualitative terms for what it concerns their impact on delay, area, and energy. A quantitative analysis is provided in the experimental section (Sec. 3.8).

- *Delay.* The duplication of the bitlines in the link avoids the occurrence of the the Miller multiplication which impacts the effective capacitance.

Thus, as it will be shown in the experiments, bitlines duplication allows to reduce the transmission delay. In addition, the absence of the Miller effect alleviates the crosstalk noise contribution with a consequent improvement in terms of communication robustness.

- *Link Area.* The link architecture without duplication reduces the wires congestion. With regard to the silicon area usage, bitlines use the higher metalization layers so as it is possible to route them over functional modules without using dedicated area for them [75]. However for long links which need the insertion of repeaters along the bitlines, since they use the same type of resources as other functional blocks, the area overhead of the repeaters should be taken into account and the use of the link architecture with duplication is preferred.
- *Energy* As it will be shown in the experimental results, the two link architectures exhibit similar average energy consumption for the same low voltage value used by the δ line. This can be simply explained by comparing their effective capacitance as it determine the average energy consumption. In the link architecture with duplication, the effective capacitance is $C_{eff} = C_s + 2C_c$ (Sec. 2.1, pp. 11). On the other side, in the link architecture without duplication, all the cases listed in Tab. 2.1 are possible. Now, considering a uniform random switching probability, the average number of switching type will be for $p = 2$. For this reason, even for the link architecture without duplication, on average, the effective capacitance is $C_{eff} = C_s + 2C_c$.

Overall, the two link architectures mainly differ in terms of delay and area. With regard to the delay, the link architecture with duplication exposes the minimal delay. In terms of area, the link architecture with duplication is significantly better than that with duplication especially when there is the need of repeaters along the bitlines. In terms of energy, the two link architectures perform almost the same for the same low voltage link level. For this reason, in the rest of this chapter, we do not distinguish between the two link architectures when energy saving results are presented.

3.5.4 Impact on the IC Design Flow

In the previous sections, proposed architectures have been discussed from an architectural point of view. Since switch line driver could be implemented by using both semi-custom or fully automated design flow, results now mandatory to clarify how designers can manage the physical implementation of proposed schemes.

- *Semi-custom Design Flow.* In order to optimize power, area and delay metrics is common in the digital flow to implement some components such as high performance adders, RAMs, PLLs, I/O pads, and so on by mean of both hard macros and standard cells. In fact, the insertion of this custom building blocks inside the automated design flow in not trivial in conjunction with existing physical implementation tools [14, 90]. In particular, a NoC platform can be optimized by using customized crossbars, FIFO RAMs and line drivers. Since regulars NoCs present well controlled parasitic contribution, proposed line drivers could be designed once (one time effort) for a given technology and a given number of integrated IPs. As said still valid for crossbars.
 - *Fully Automated Flow.* In several cases, in order to speed up design phase, is recommended a fully automated implementation by mean of standard cells based flow. Several Process Design Kit provides standard cell that can operate whit different operating voltages. During the physical implementation, several power domain can be specified by specifying such voltage. Different power domain can communicate together by mean of level shifter. For these reasons, proposed architectures can be implemented by using level shifters (used to implement level restorer) before/after line drivers, tristate buffer. It should be noted that, in the case of the architecture whit duplicated lines, repeater insertion does not implement level shifter. In this case low voltage bit-lines should be inserted inside actual power domain. Furthermore, since level shifter are usually implemented by transmission gate logic, several input/output operating voltage can be implemented. Further, in practical case, since standard cell are well characterized for different operating corners (PVT), when a low level path not is involved in an aggressive voltage scaling, the same standard-cell library can be
-

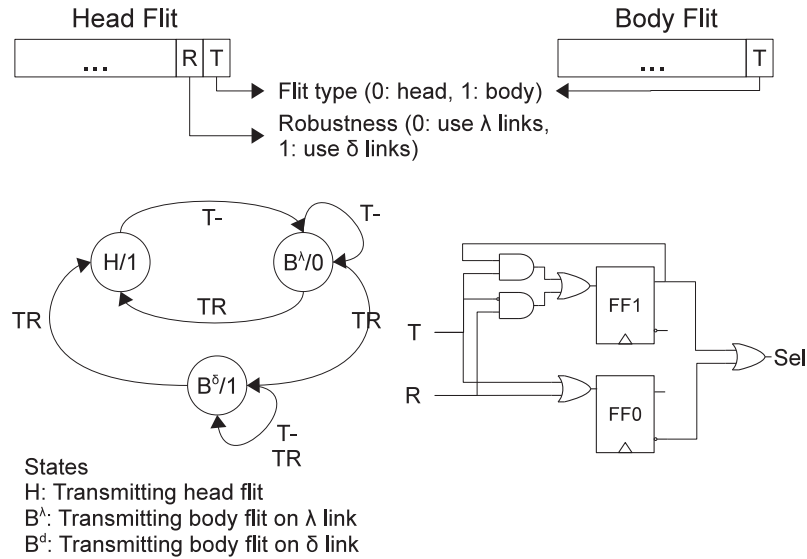


Figure 3.4: Finite State Machine implementing the selection logic in Fig. 3.2.

reused for both high and low voltage. Due to delay degradation such scenario is suitable only when crossbar and lines are not part of the router critical path.

To avoid any ambiguity, if not otherwise specified, a full custom implementation of line drivers and crossbar have been considered in the rest of this chapter.

3.5.5 Control Circuitry

The selection logic for driving the selection signal in Figs. 3.2 and 3.3 can be implemented by means of a simple finite state machine (FSM) as shown in Fig. 3.4. We assume that the head flit provides a T bit which defines the flit type (either head or body), and a R bit which defines the robustness of the current communication. A body flit, does not have the R bit but only the T bit. As it can be observed, head flits ($T = 0$) determine $Sel = 1$ irrespectively of the value of bit R . Body flits determine $Sel = 1$ ($Sel = 0$) if the R bit in the previously transmitted head flit was $R = 1$ ($R = 0$).

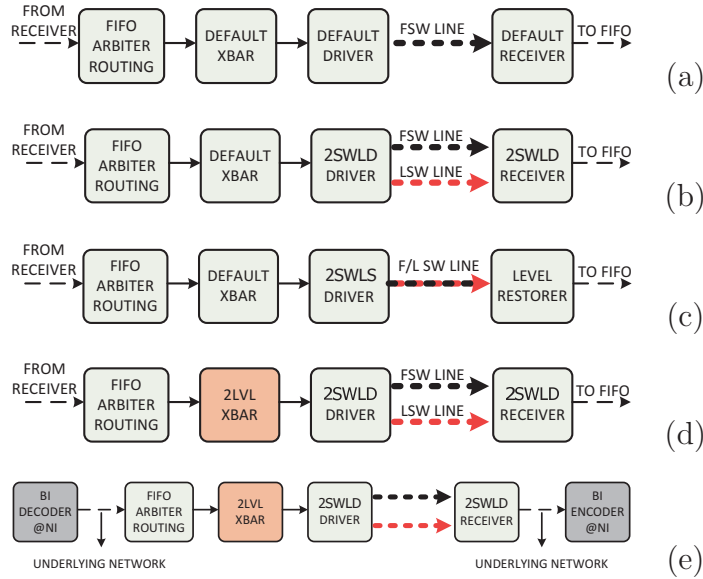


Figure 3.5: Design configurations.

3.6 Design Configurations

Concepts introduced for links could be used also for the routers crossbar. In fact, the latter is essentially composed by several wires driven by large multiplexers. A squared effects on power reduction could be also obtained, not only for wires but also in to the router in order to obtain an higher energy saving. Since applying proposed technique to the crossbar is a naive extension of structure already presented for links, we will not enter in detail to the circuital implementation of a modified crossbar, but we prefer the same, to provide several possible configurations obtained by mixing links organizations discussed so far with other that can be obtained applying the scheme to the crossbar. Such configurations are characterized by a different trade-off between performance, power, reliability, and area. Specifically, the configurations that will be analyzed in the experimental section are shown in Fig. 3.5 and are introduced in the following.

1. Baseline [Fig. 3.5(a)]: A traditional transceiver modeled with a chain of inverters as driver and a single inverter as receiver. This is the baseline scenario and represents the standard implementation in a traditional digital design flow when there is no need for the insertion of repeaters. Within the router a conventional crossbar is used.

2. 2SWLD [Fig. 3.5(b)]: The proposed link transceiver with lines duplication presented in Sec. 3.5.1. Within the router a conventional crossbar is used.
3. 2SWLS [Fig. 3.5(c)]: The proposed link transceiver without duplication presented in Sec. 3.5.2. Here the shielding property obtained with the duplication is absent, but the link area is reduced.
4. 2SWL+2SWXBAR [Fig. 3.5(d)]: Like the 2SWLD configuration in which the two swing signaling is enabled in the crossbar.
5. 2SWL+2SWXBAR+BI [Fig. 3.5(e)]: Like the 2SWL+2SWXBAR in which the data encoding Bus Invert [87] is used. The Encoder and decoder are situated within the network interfaces as proposed in [72].

We omit the remaining two configurations obtained by coupling 2SWLS with the two voltage swing crossbar (2SWXBAR) and with the bus invert data encoding scheme (BI). In fact, in terms of energy saving, they perform almost the same as the 2SWLD counterpart as it will be shown in the experimental section.

3.7 Synthesis Results

The link and crossbar architectures discussed so far have been designed and analyzed. The designs have been targeted for working at a clock frequency of 2 GHz (which is the target clock speed of our baseline router). The analysis has been carried out with HSPICE using a 45 nm CMOS LVT library from Nangate [2] which provides 10 metal layers. The parasitics extraction from layout has been made using Cadence Virtuoso. With the same tool we estimate the silicon area occupied by the links and by crossbar switch. The results are reported in Tab. 3.1.

We considered a 8×8 2D-mesh based NoC architecture in a $20 \text{ mm} \times 20 \text{ mm}$ silicon die. The link length can be computed as [76, 42]:

$$l = \frac{\sqrt{Area}}{\sqrt{M} - 1} \quad (3.1)$$

where M is the numbers of tiles (*i.e.*, 64 in our case). Based on our parameters, the line length is 2.8 mm. We used the seventh metal layer for the

Table 3.1: HSPICE simulation results for a bitline of the link.

	Conventional full-swing	Proposed low-swing			
		Double link		Single link	
		@ 0.8 V	@ 0.7 V	@ 0.9 V	@ 0.8 V
Technology		1.1 V, 10 metal, 45 nm CMOS LVT			
Interconnect		Metal 7: Width 0.4 μm , Space 0.32 μm , Length 2.8 mm, Rwire 225 Ω			
	Cwire 946 fF [†]	Cwire 512 fF [†]		Cwire 946 fF [*]	
Supply (VDDH/VDDL)	1.1 V/NA	1.1/0.8 V	1.1/0.7 V	1.1/0.9 V	1.1/0.8 V
Worst case total delay ($T_{d50\%}$)	214 ps	280 ps	380 ps	375 ps	410 ps
Avg. Energy/Transition	512 fJ	520 fJ/274 fJ	520 fJ/230 fJ	527 fJ/304 fJ	527 fJ/258 fJ
BER	1.3E-17	1.3E-17/3.8E-10	1.3E-17/3.6E-8	1.3E-17/2.2E-12	1.3E-17/3.8E-10

* Due to Miller multiplication.

† Miller multiplication cannot occur.

inter-router link (upper levels are used for supply and clock distribution). The physical line and the transceiver are simulated using the same setup in [107]. The multiplexer shown in Fig. 3.2, is based on well sized static-CMOS logic for improving the driving capability of the driver [89]. With regards to the drivers of the δ lines, they are implemented by a chain of cascaded inverters with last transistor size of $W_p = 30 \mu\text{m}$ and $W_n = 10 \mu\text{m}$. For the low-swing voltage path (λ) we used the driver of the CLC transceiver proposed in [86] where the driver size is the same of that used in the δ lines. For what it concerns the receiving section at the end of the line, we use a minimal size inverter for the full-swing voltage path (δ) and by a level restorer [86] for the low-swing voltage path included in the CLC scheme. Similarly, for the single link architecture (Sec. 3.5.2) the shared δ and λ lines are implemented by a chain of cascaded inverters with last transistor size of $W_p = 35 \mu\text{m}$ and $W_n = 12 \mu\text{m}$.

In our experimental set-up there has not been the need of inserting repeaters along the bitlines of the link. In fact, for these relatively short lines, repeaters reduce the delay only marginally [47] as the dominant time-constant of the interconnect itself is only $1/2R_{wire}C_{wire} = 57 \text{ ps}$ (Tab. 3.1). However for such line lengths, in chip manufacturing, gate oxide can be easily damaged by electrostatic discharge. The static charge that is collected on wires during the multilevel metalization process can damage the device or lead to a total chip failure. The phenomena of an electrostatic charge being discharged into the device is referred to as either *antenna* or *charge-collecting antenna problems*. This issue can be addressed with the insertion of protection diodes connected to the transistor gate, or cutting the wires in several small segment, changing properly the level between each segments pairs [90]. In this work, we consider the former approach for facing with such antenna effect.

For the double link configuration we considered two cases with VDDL 0.8 V and 0.7 V. For the single link configuration we considered two cases with VDDL 0.9 V and 0.8 V. In fact, in the single link configuration the performance requirements of 2 GHz could be met for a low swing voltage higher than 0.75 V.


For what it concerns the performance, the delay introduced by the line increases as the voltage level decreases. This is a well known issue mainly introduced by the level restorer of the above mentioned CLC scheme. It should be pointed out that, however, there are other more sophisticated solutions for implementing the level restorer which have a limited impact on the delay [107] (or, in some cases, improve the performance). Further, the double link transceiver is faster than the single link transceiver. This is due to the reduced crosstalk capacitance due to the shielding structure discussed in Sec. 3.5.1. For this reason, looking again at Tab. 2.1 (Sec. 2.1, page 12), for the double line transceiver the effective capacitance is $C_{eff} = C_s + 2C_c$ as the other transition types cannot occur. For the single line transceiver, worst case delay happens when the aggressor lines make an opposed transition with respect to the victim line ($C_{eff} = C_s + 4C_c$).

In terms of energy consumption, the single link configuration is slightly better than the double link configuration. For instance, considering the case at 0.8 V present in both the configurations, the average energy saving per transition is 46% and 55% for the double link and single link, respectively. In fact, the average energy consumption is computed considering all the possible transition patterns for a given configuration. In particular, for the cases $p = 0$ and $p = 1$ in Tab. 2.1, which are possible for the single link configuration, the effective capacitance is lower than that exposed by double link implementation in which only the case $p = 2$ is possible.

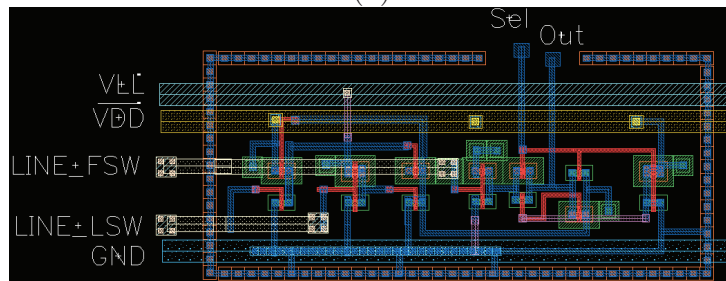
Similarly, in the Tab. 3.2 we report the results obtained for a 5×5 crossbar. Most of the considerations made for the link are still valid for the crossbar. Differently from the link, the impact of crosstalk in the crossbar is less evident than in the link as it is partially masked by the large junction capacitance introduced by the tristate buffers multiplexer. Fig. 3.6 shows the layout of the proposed link driver and receiver for a single bitline. In the transmitter, we can distinguish the cascade of inverters (tapered buffer) supplied at two different voltage levels namely, VDDH and VDDL. We can also

Table 3.2: HSPICE simulation results for a bitline of the crossbar.

	Conventional full-swing	Proposed low-swing	
		@ 0.8 V	@ 0.7 V
Technology	1.1 V, 10 metal, 45 nm CMOS LVT		
Interconnect	Metal 6-7: Pitch 0.28-0.8 μm		
Ports	5 Ports; No U-turns are permitted		
Tristate logic style	Transmission Gate		
Supply VDDH/VDDL	1.1 V/NA	1.1 V/0.8 V	1.1 V/0.7 V
Worst case total delay ($T_{d50\%}$)	200 ps	335 ps	350 ps
Avg. Energy/Transition	42 fJ	44 fJ/20 fJ	44 fJ/17 fJ
BER	1.3E-17	1.3E-17/3.8E-10	1.3E-17/3.8E-10



(a)



(b)

Figure 3.6: Layout of the proposed link driver (a) and receiver (b) for a single bitline.

Table 3.3: Silicon area occupation in μm^2 and percent overhead.

Configuration	Router (A)	Line Driver (B)	Line Receiver (C)	Codec/Logic (D)	Link Area (E)	Total (A+B+C+D)	Overhead (%)
Baseline	47687	3712	544	0	71680	51943	-
2SWLS	47757	10368	2880	0	145600	58514	12.3%
2SWLD	47757	7872	2880	0	73920	59250	11.2%
2SWL+2SWXBAR	50887	7872	2880	0	145600	61714	15.8%
2SWL+2SWXBAR+BI	50887	7872	2880	810	145600	62449	16.8%
DAP	47687	7680	2304	600	145600	58346	10.9%
JTEC	47687	9240	2772	2208	172480	62183	16.4%
JTEC-SQED	47687	9480	2808	2532	176960	62823	17.3%

notice the selection terminal SEL used for selecting the full- or the low-swing signaling path. At the bottom of the same pictures it is shown the receiver. In both driver and receiver we can notice the two lines, namely, LINEFSW and LINELSW for the full-swing and low-swing, respectively. The silicon area occupation is $246 \mu\text{m}^2$ and $90 \mu\text{m}^2$ for the transmitter and receiver, respectively. For the sake of comparison, we analyze a set of representative data encoding techniques that will be considered in the experiments section below. Precisely, we designed and synthesized the encoding and decoding logic of DAP presented in the previous section and for other similar techniques such as JTEC and JTEC-SQED [42]. A complete report is shown in Tab. 3.3 which reports not only the overhead as respect to the entire NoC but also the absolute value of the overall occupied area. Please note that the links area (contribute E in the table) is not taken into account in the total area computation. In fact, the higher metalization layers have been used for the bitlines of the links, so as it is possible to route them over functional modules, that is, route bitlines of the links without using dedicated area for them [75].

3.8 Experiments

In this section we present the results of experiments carried out on both synthetic and real traffic scenarios. Noxim [35, 17] NoC simulator has been extended to support the proposed link and crossbar architectures. The power model implemented in Noxim has been updated with the power figures extracted from DSENT [34] and HSPICE simulations. In particular, we used DSENT for estimating the the energy contribution of the various components of the router. The energy consumed by the links and the crossbar (for all the configurations presented in Sec. 3.6) have been estimated by using




Figure 3.7: NoC energy breakdown for different link lengths and configurations. (a) Baseline , (b) 2SWLD @ 0.7V (b) 2SWLD @ 0.7V + 2SWXBAR @ 0.7V.

HSPICE. We consider a baseline router implementation clocked at 2 GHz, 4-flit input buffers, and 32-bit flit. A 8×8 mesh-based NoC is considered in the experiments.

For the sake of clarity Fig. 3.7 shows energy contributions for each switch building block. In particular, three possible scenario are reported: baseline, proposed scheme with duplication and aggressive scaling (2SWLD @ 0.7V), proposed scheme with aggressive scaling in both crossbar and link (2SWLD @ 0.7V + 2SWXBAR @ 0.7V). It should be pointed out that Fig. 3.7-(b) and (c) consider only energy contribution due to the lower swing path. Since energy contribution due to the electrical wires is reduced decreasing link length, proposed scheme applied only in the link (2SWLD @ 0.7V) is less effective in the case of 1.8 and 0.8 mm. At the same time both FIFO and crossbar become dominant for such short lines. From these considerations an important result emerges by applying proposed scheme also in the crossbar (2SWLD @ 0.7V + 2SWXBAR @ 0.7V). In fact as link length decrease and crossbar energy is reduced, the overall router energy results reduced. By considering the latter case a theoretical maximum energy reduction of about 44%, 36% and 32% can be obtained for a link length of 2.8, 1.8 and 0.8 mm respectively.

For the sake of clarity, most of the experiments are carried out on configuration 2SWL+2SWXBAR (*cf.* Sec. 3.6). In fact, it is a representative configuration as it couples the proposed link architecture and the proposed crossbar architecture. In addition, without loss of generality, in most of the experiments we consider the runtime reconfigurable link architecture with duplication as, in terms of energy saving, it performs almost the same as the single link implementation. In the following, if not different specified, energy

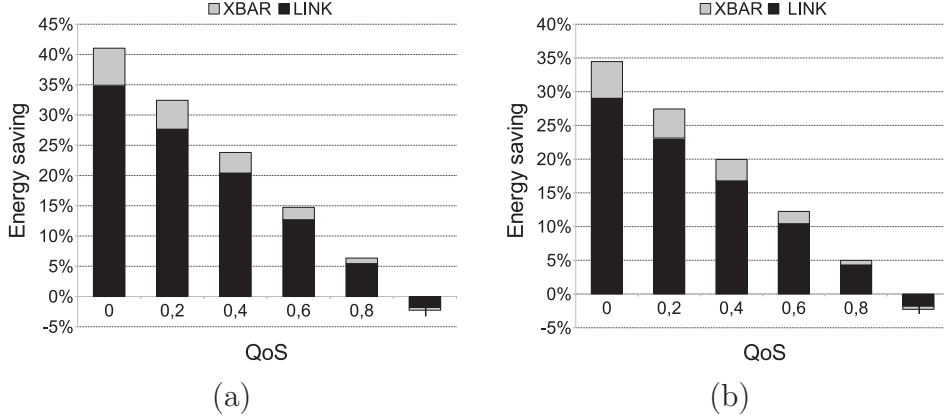


Figure 3.8: Energy saving for different QoS. (a) VDDL 0.7 V, (b) VDDL 0.8 V.

saving results are with respect to the baseline configuration. In particular, such savings are obtained from Noxim simulations after considering energy results and by applying the following:

$$E_{saving} = 100 \times \frac{E_{baseline} - E_{proposed}}{E_{baseline}},$$

where $E_{baseline}$ and $E_{proposed}$ are energy values obtained after running Noxim with baseline and proposed configuration respectively. It should be pointed out that since Noxim has been back-annotated with energy contribution of every NoC building block and additional circuitry (if required), reported saving refer to the entire communication infrastructure. Energy dissipated by processing elements has not been considered because is out of the scope of our study.

3.8.1 Energy Saving vs. QoS

Fig. 3.8 shows the percentage energy saving obtained by using configuration 2SWL+2SWXBAR. The percentage energy saving is reported for different quality-of-service (QoS). Here, with the term QoS we mean the fraction of communications marked robust over the total communications. A QoS of one (zero) means that all the communications in the communication graph of the application are marked with the robustness flag high (low). Uniform random traffic and random data patterns are considered. Packets size are randomly generated between 2 and 10 flits. The graph highlights the energy saving contribution due to the crossbar and to the link individually. As

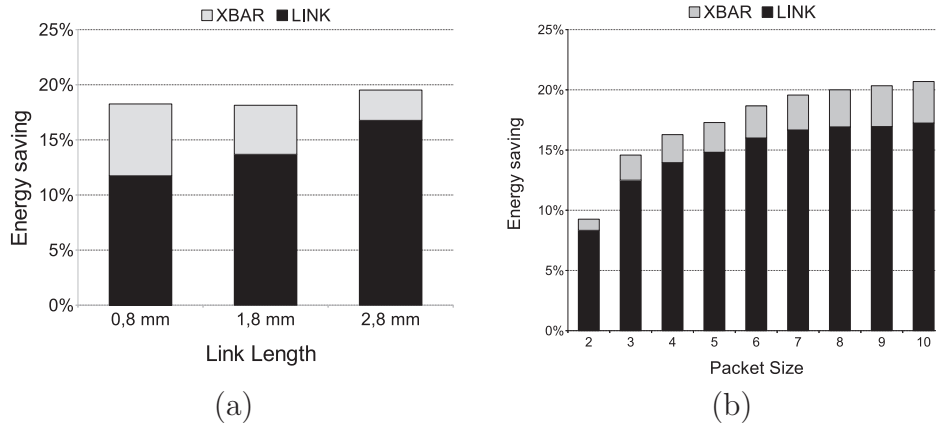


Figure 3.9: Energy saving: (a) for different link lengths; (b) for different packet sizes.

expected, as the percentage of communications marked with the robustness flag low increases (QoS decreases), the percentage of energy saving increases. It should be pointed out that, when QoS is 1 (*i.e.*, all the communications have the robustness flag set), the energy consumption when the proposed scheme is used is higher than that in the baseline NoC. This is due to the fact that, the power overhead in the proposed transceiver is higher than the power reduction due to the elimination of the crosstalk effects in the link.

3.8.2 Energy Saving vs. Link Length

Let us now analyze the energy saving when the length of the link is made to vary. Fig. 3.9-a shows the percentage energy saving when configuration 2SWL+2SWXBAR at 0.7 V is used for different link lengths. We considered three link lengths of 2.8 mm, 1.8 mm, and 0.8 mm. We used the same traffic scenario considered in the previous subsection for a QoS of 0.5. We can observe that, if link length decreases, the energy saving decreases. This is because the link energy contribution becomes less dominant as respect to the overall energy consumption. On the other side, the contribute in energy saving due to the crossbar increases. Overall, the total energy saving remain almost the same as link length decreases.

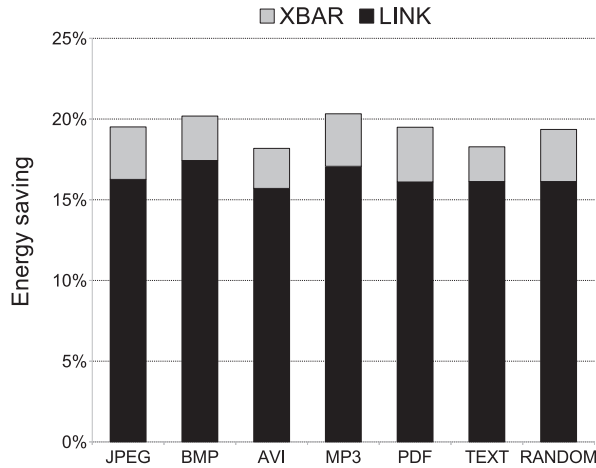


Figure 3.10: Energy saving for different data types.

3.8.3 Energy Saving *vs.* Packet Size

The packet size impacts the energy saving since the head flit is transmitted through the high-swing path even if its robustness flag is set or not. Fig. 3.9-b shows the percentage energy saving under uniform random traffic and random data patterns for a QoS of 0.5 and different packet sizes when the organization 2SWL+2SWXBAR at 0.7 V is used. As expected, as packet size increases, energy saving increases as well. As packet size goes beyond 8 flits, there is not any relevant improvement in energy saving. That is, the energy penalty due to the transmission of the head flit through the high-swing path is rapidly absorbed by the positive effect of transmitting the body flits through the low-swing path.

3.8.4 Energy Saving *vs.* Different Data Types

Link energy consumption depends on the kind of data traveling on the links which determine different switching activities. In the previous experiments, we considered randomly generated data patterns. Let us now consider the case in which the data patterns belong to a set of data streams from eight different media formats namely ASCII text, PDF, gray scale image and true color image (both in BMP and JPEG formats), MP3 audio and MPEG video. For each class, ten data streams are considered and average values are reported. Fig. 3.10 shows the energy saving for QoS of 0.5 at 0.7 V. As it

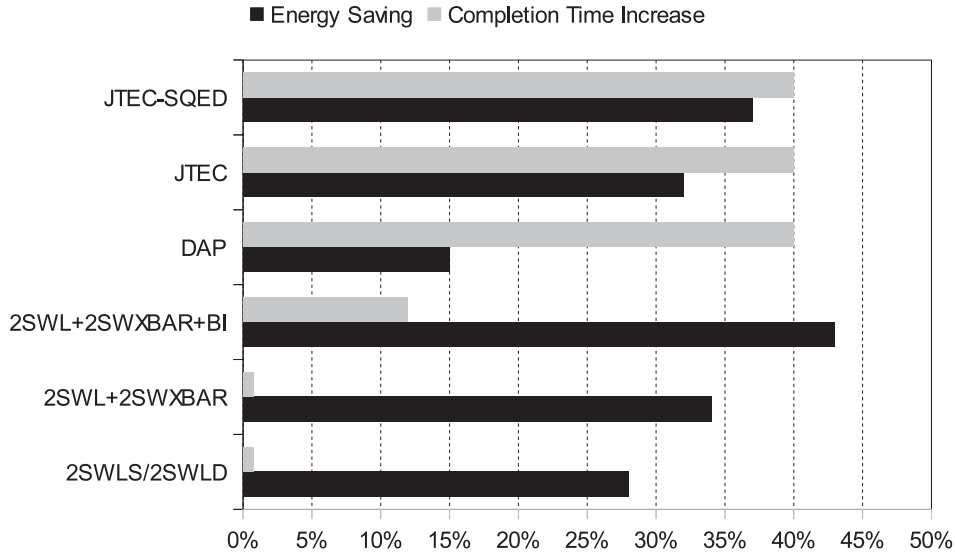


Figure 3.11: Energy saving vs. completion time increase.

can be observed, on average, the energy saving is about 17% if the proposed technique is applied only to the links and increase to about 21% when the proposed technique is applied to the crossbar too.

3.8.5 Energy Saving vs. Performance Degradation

The tradeoff between the energy saving with the completion time (*i.e.*, the amount of the time needed to drain a given amount of traffic volume) is an important characteristic of the system. The percentage increase of completion time is defined as the percentage increase of the time needed to drain a given amount of traffic.

We assume 32-bit links, and packets of 4 flits (flit size is 32 bits). The configurations 2SWLS, 2SWLD, and 2SWL+2SWXBAR require an additional 1 bit in the head flit for carrying the robustness flag. Thus, the overhead in the packet is $1/(32 \times 4) = 0.8\%$. In configuration 2SWL+2SWXBAR+BI, we apply an end-to-end data encoding scheme based on bus invert as presented in [72] where we found that the best configuration in terms of energy/performance tradeoff is partitioning the link in four sub-links of 8 bits. In this case, other than the robustness flag to be stored in the head flit, the four

invert bits have to be stored in each body flit with a consequent overhead in the packet of $(1 + 4 \times 4)/(32 \times 4) = 13\%$. The data encoding schemes DAP, JTEC, JTEC-SQED do not result in any overhead in the packet as dedicated physical control lines are used instead. Precisely, from the baseline configuration with 32-bit links, DAP, JTEC, and JTEC-SQED require links of 65-, 77-, and 78-bit, respectively. In addition, while 2SWLS, 2SWLD, 2SWL+2SWXBAR, and 2SWL+2SWXBAR-BI do not introduce any latency in the router, DAP, JTEC, JTEC-SQED increase the pipeline depth of the router with two additional stages for data decoding and encoding tasks. We considered router with five pipeline stages (buffer, routing, switch allocation, switch traversal, link traversal) for 2SWL, 2SWL+2SWXBAR, and 2SWL+2SWXBAR+BI, whereas we considered a router with seven pipeline stages (buffer, data decoding, routing, switch allocation, switch traversal, data encoding, link traversal) for DAP, JTEC, and JTEC-SQED. Fig. 3.11 shows the tradeoff between energy saving and completion time increase. As it can be observed, the proposed configurations are characterized by the best tradeoff energy saving/completion time increases. In fact, although all the techniques provide comparable energy savings, the proposed configurations have a limited impact on performance and Pareto dominate the other considered techniques.

3.8.6 Case Studies

Let us now assess the proposed techniques on two real case studies, namely, a multimedia system and a JPEG codec.

Multimedia System

We consider an heterogeneous multimedia system which includes a H.263 video encoder, a H.263 video decoder, a MP3 audio encoder and a MP3 audio decoder [48]. The communication graph of the multimedia application is shown in Fig. 3.12. Communications represented with a dashed arrows are those communications with the robustness flag set as low. The application is mapped on a 4x4 mesh-based NoC using GAMAP [74].

In this case study, both packet size and packet injection rate vary with communication flow. For instance, the communication flows involved in

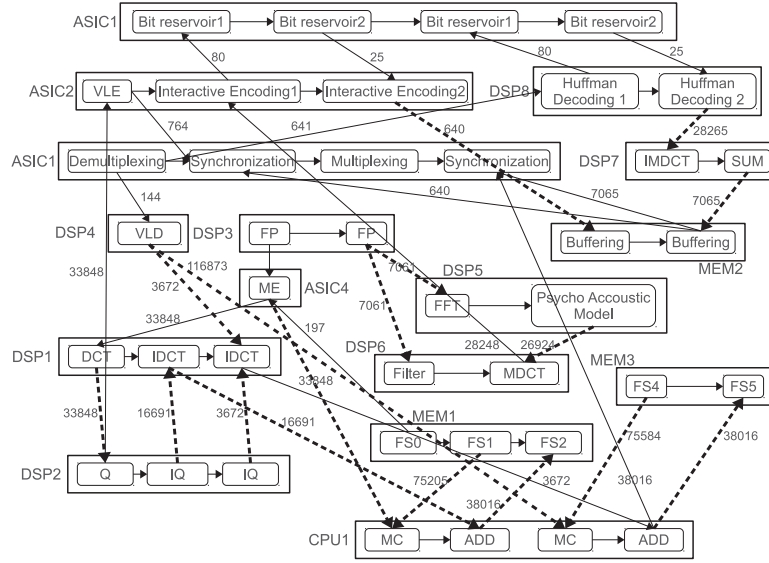


Figure 3.12: Communication graph of the multimedia application.

MMS-Enc and MMS-Dec use a packet size tuned on the basis of a macroblock. Packet injection rate has been computed for each communication flow on the basis of the bandwidth requirements for each application as reported in [48].

Fig. 3.13 shows the energy saving when the proposed technique is used for different inter-router line lengths. It also shows the energy saving obtained when data encoding techniques are used. Specifically, we considered the state-of-the-art data encoding techniques [42] such as DAP, JETC, and JTEC-SQED.

As it can be observed, the proposed 2SWL+2SWXBAR configuration allows to save up to 28% and 34% of energy when the low path swing works at 0.8 V and 0.7 V, respectively. The energy saving increases up to 39% and 43% when the flits are encoded using bus invert (2SWL+2SWXBAR+BI). It should be pointed out that, although the BI technique has been designed to be applied in the context of off-chip bus (*i.e.*, long and high capacitive bus where the coupled capacitance is negligible as compared to the self capacitance), it becomes a good on-chip solution when coupled with the proposed technique. In fact, using the proposed technique, the crosstalk effects are drastically reduced, making the energy contribution due to the self capacitance again relevant. Thus, using BI in conjunction with the proposed technique, allows to reduce the self switching activity that otherwise would




Figure 3.13: Energy saving for the multimedia application.

be ignored. Although the energy saving obtained with JTEC/JTECSQED is comparable for a link length of 2.8 mm, this is not true for shorter line lengths for which energy saving exhibited by 2SWL+2SWXBAR is higher.

Furthermore, it should be pointed out that, all the data encoding techniques considered in this analysis have a negative impact on the timing of the router. In fact, their application require the introduction of two pipeline stages in the router [42] with a consequent increase in communication latency. This might not be tolerated in applications with tight time constraints. Conversely, in our proposed techniques, the delay introduced by the multiplexer/demultiplexer does not affect the pipeline depth of the router.

In the last experiment, we assess the trade-off between energy saving and perceived video quality. The latter is measured by computing the probability that a video frame is affected by an error (frame error rate). Here, with the term error, we refer to the case in which at least a bit of the RGB components of a pixel is affected by an error. Precisely, starting from the communication graph shown in Fig. 3.12, we progressively transform the no robust communications (dashed arrows in the figure) to robust communications, in descending order respect to the bandwidth requirement. Fig. 3.14 shows the tradeoff between energy and frame error rate (FER). Energy values are normalized by the energy measured when all the communications are marked as robust. As expected, the energy required to guarantee a certain FER increases as FER decreases. Please note that, even in the case of min-




Figure 3.14: Tradeoff between energy and the probability that a video frame is affected by an error. Energy values normalized by the energy measured when all the communications are marked as robust. (a) 2SWLD @ 0.8 V, (b) 2SWLD @ 0.7 V.

imum energy consumption, FER values are in the order of 10^{-3} and 10^{-5} which does not result in any appreciable degradation of video quality.

JPEG Codec

The second case study is a JPEG codec. Firstly, let us analyse the impact on the image quality when the 2SWL+2SWXBAR configuration is used. Fig. 3.15(a) shows the simulation setup. The left side of the figure shows the steps performed by the encoder, whereas the right part shows the steps performed by the decoder. Each computational step of the flow is mapped on a node of a 4×4 mesh-based NoC as shown in Fig. 3.15(b). When the nominal voltage swing is used, we assume that there are no communication errors. When the low voltage swing is used, the bit error rate is computed by means of Eqn. (2.4). Thus, we encode and decode an image in two cases: the case in which communications are mapped on links working at nominal voltage swing, and the case in which they are mapped on links working at low voltage swing. Finally, the two images are compared qualitatively and quantitatively.

From a qualitative viewpoint, Fig. 3.16 shows the images obtained from the decoding stage when two low voltage swing at 0.5 V and 0.6 V are used. For higher voltage swing (>0.6 V), there is not any appreciable difference

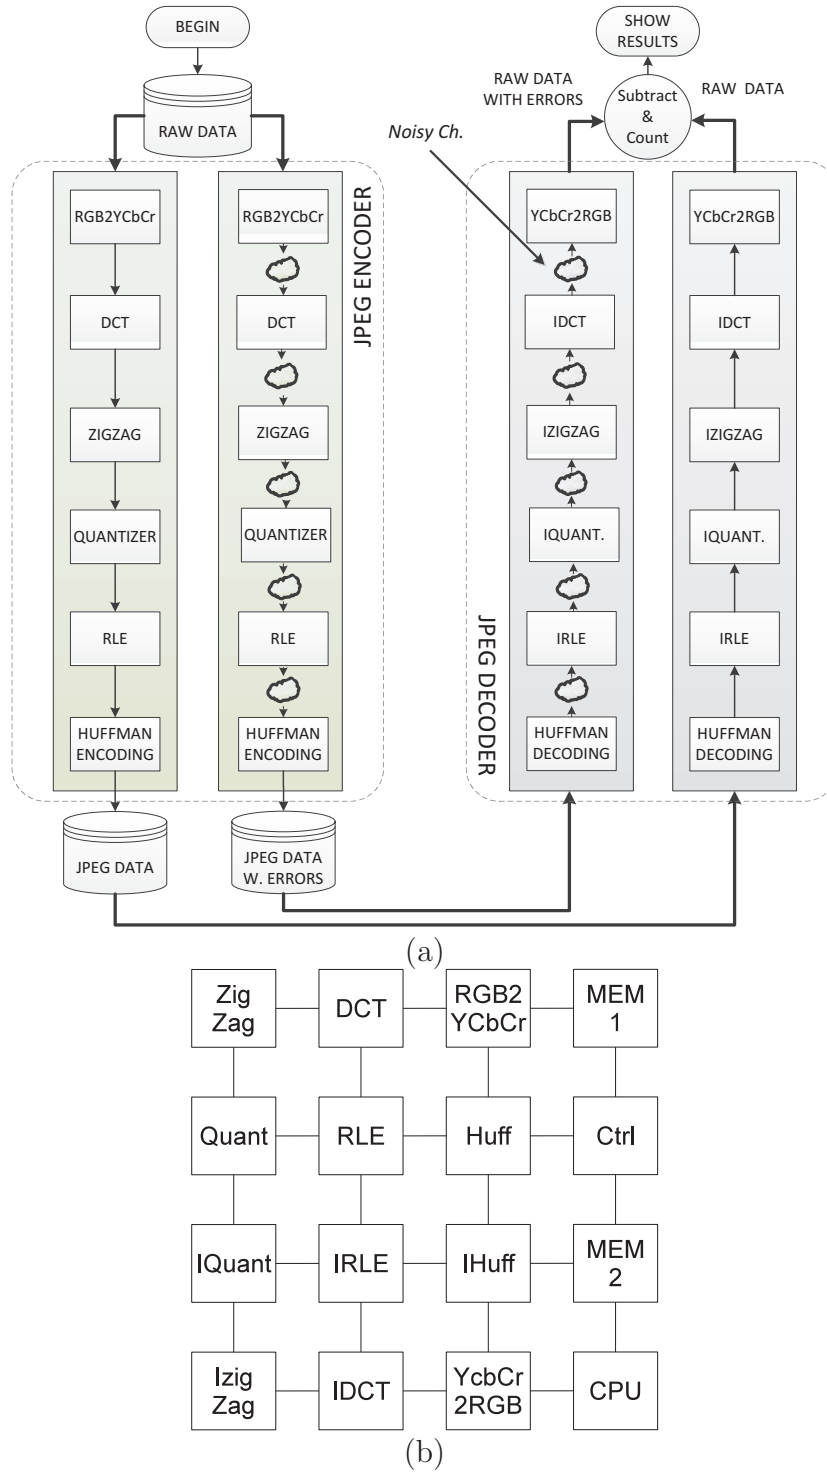


Figure 3.15: Simulation setup used for the JPEG codec (a), and mapping of the JPEG codec on the NoC (b).



Figure 3.16: Image encoded and decoded based on the flow shown in Fig. 3.15 using two low voltage swing at 0.5 V (a) and 0.6 V (b).

between the original image and the encoded and decoded one. Quantitatively, Fig. 3.17 shows the percentage of the image affected by errors (spurious pixels) for different low voltage swing values.

Fig. 3.18 shows the percentage energy saving when the proposed technique is used. As it can be observed, more than 49% of energy can be saved using the proposed techniques. In fact, using the proposed 2SWL+2SWXBAR+BI with low voltage swing of 0.7 V, the total communication energy saving is 49%. For such low voltage swing values, as shown in Fig. 3.17, the perceived image quality is the same as the original image (the percentage of spurious pixels is zero).

3.9 Conclusions

The on-chip communication network accounts for a significant fraction of the overall energy budget of a multi/many-core system. The crossbar into the routers and the links which connects the routers are the main responsible for the energy consumption of the NoC. While reducing the voltage swing of these energy hungry elements has a positive effect in terms of energy saving, on the other side, the communication reliability decreases due to the increase of the bit-error-rate (BER). Starting from the assumption that, in general, not all the communications in an application have the same reliability re-

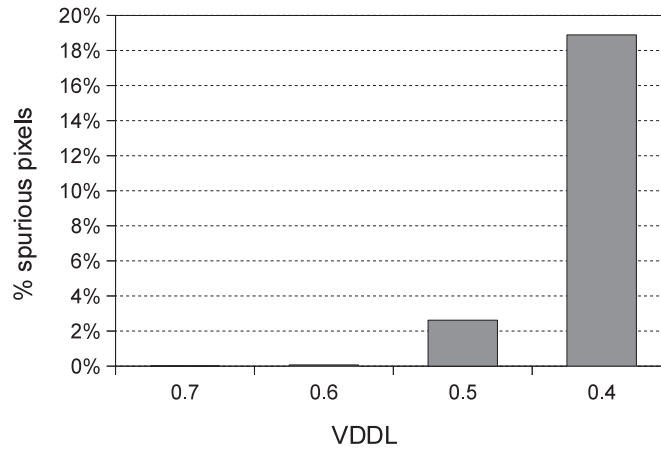


Figure 3.17: Percentage of spurious pixels in the image for different low voltage swing values.




Figure 3.18: Energy saving for JPEG codec application.

quirements, in this chapter we have presented methods and architectures for run-time tuning the voltage swing for signaling in links traversed by the flits of a packet based on the communication reliability requirement of that particular communication. Results obtained for links have also extended for crossbar. Experiments carried out both synthetic and real traffic scenarios have shown the effectiveness of the proposed technique in terms of energy saving. As compared to the state-of-the-art in link energy reduction through data encoding schemes, the proposed technique provides higher energy saving without impacting the performance metrics of the system.

Chapter 4

Emerging Network-on-Chip Paradigms

As stated in the introduction, in nowadays SoCs, as the number of integrated cores increases (according to ITRS projections), traditional interconnect paradigms no longer satisfy the actual performance requirements. The continued progress of interconnect performance will require approaches that introduce materials and structures beyond the conventional metal/dielectric system, and may require information carriers other than charge. Multiple options have been envisioned to provide alternatives to the metal/dielectric system. In particular, three emerging interconnect technologies are three-dimensional (3-D) integration, nanophotonic communication, and RF/wireless on chip interconnects [15, 50]. Starting from these technologies, several emerging NoC paradigms have been developed, such as *3D NoCs*, *Photonic NoCs* and *Wireless Networks-on-Chip*. The purpose of this section is to provide a brief taxonomy for each of the above-mentioned emerging paradigms. Since techniques developed in the research presented here are only suitable for wireless Networks-on-Chip (WiNoCs), the latter will be discussed with more emphasis in the final part of this chapter.

4.1 3D NoC

Three-dimensional integrated circuits (3D ICs) offer an attractive solution for overcoming the barriers to interconnect scaling. Despite several technological options have been introduced in order to implement 3D on-chip structures,




Figure 4.1: 3D Symmetric NoC.

Trough-Silicon-Via (TSV) approach is emerging as the most promising [43]. In particular, the latter consists in a stack of silicon dies interconnected through several metallic vias.

Today's silicon design CAD tools already implement the support for this technology [14]. From a designer point of view, in fact, synthesized gate-level netlists can be spanned among several layers in a quite trivial way. Each logical module of a design can be thus physically positioned upon a particular layer and interconnected to each others by means of defined inter-layers I/O ports. Traditional NoC concepts could be therefore used in conjunction with the above-mentioned integration solutions in order to build a new interconnect paradigm named *3D NoC*. In particular, the NoC research community proposes several alternatives in this sense. Some of them will be reported in the following part of this section.

4.1.1 3D Symmetric NoC

The most easy and natural way to extend a baseline NoC router to the 3D technology it consists on the introduction of two additional ports. As depicted in the picture 4.1, a router port can be used to connect the upper layer while another is used for the lower level. The basic pipeline structure of a NoC will be thus unaffected and communications still remain in a multi-hop fashion. This architecture is simple to implement but has two major

Table 4.1: Area and power comparison of the crossbar switches in a 90nm technology [15].

Xbar Type	Area	Power (500 MHz)
5×5	$8523 \mu m^2$	$4.21 mW$
6×6	$11579 \mu m^2$	$5.06 mW$
7×7	$17289 \mu m^2$	$9.41 mW$

inherent drawbacks:

1. It wastes the beneficial attribute of a negligible inter-wafer distance in 3D chips. In fact, the thickness of a die could be as small as 10s of μm while interconnect for going in horizontal is long on the order of few mm.
2. The addition of two extra ports necessitates a larger 7×7 crossbar instead of a 5×5 one. In the NoC context it is known that crossbars scale upward very inefficiently, as illustrated in Table 4.1 [15]. The latter includes the area and power budgets of crossbars synthesized in a 90 nm technology.

For the reasons stated before, it is clear that a 3D Symmetric implementation is not an optimal solution. For this reason NoC research community developed other optimized architectures, some of them are shown in the following part of this section.

4.1.2 3D NoC-Bus Hybrid Architecture

Since inter-wafer distance is in order of 10s of micrometers, communication involved in a vertical direction could happen in a single hop fashion. To enable this feature, vertical links could be shared among layers [52]. This realization opens the door to a very popular shared-medium interconnect, the bus. In this manner a 6×6 crossbar could be used (rather than a 7×7). The price to pay is in terms of bandwidth. In fact, since each vertical link is shared, an arbitration is required when concurrent routers would to use the link at the same time. Further, additional hardware is required. A dedicated interface with a dedicated queue is indeed implemented for each router. This architecture is not indicated when the application require an high bandwidth.

4.1.3 Multi-layer 3D NoC Router Design.

All the architectures discussed above, start from the assumption that the processing element is designed in a 2D fashion. To augment the granularity of a 3D design, a processing element could be split among different layers [95, 102]. With such multi-layer stacking of processing elements, also the router could be divided among different layers [78]. In this way, from a topological point of view, a 3D NoC is equal to a 2D NoC. Since both router and *PEs* are smaller, the saving in chip area can be used for enhancing the router capability, for example, adding express link between non adjacent *PEs* to reduce the average hop count.

4.2 Photonic NoCs

Optical communication is widely accepted as an interconnection medium for long and medium-reach distances typically above 10 m. Thanks to recent technological improvements, today's chips can integrate devices that permit to implement nanophotonic communication to establish both inter and intra-chip links [11]. In fact, the continued advances in photonic technology have resulted in the decrease of CMOS-compatible photonic device sizes that have become comparable to electrical components.

The main components of on-chip nanophotonic communication include a light source, the waveguide where the light is routed, a modulator for electrical to optical signal conversion, and a detector for the optical to electrical conversion. For the modulators and the detectors, micro-ring resonator-based technology is commonly used. The main advantages of this paradigm consists in the fact that very high speed link, over 1 Tbps, can be practically achieved independently to the distance among sources and destinations. Silicon optical communication can leverage on two important properties:

- **bit-rate transparency:** differently from electronic electrical communications, in which power consumption depends on the switching activity, a nanophotonics link power consumption is independent from the bitrate. Thus, a photonic modulator consumes energy only if it is activated.
- **low loss in optical waveguides:** at the chip and board scale, the

Table 4.2: ITRS projections for the transition frequency f_t and maximum oscillating frequency f_{max} [5].

Year	2012	2013	2014	2015	2016	2017
f_t (GHz)	315	315	345	360	375	390
f_{max} (GHz)	420	455	490	525	560	595

power that is dissipated on a photonic link is independent of the transmission distance. As reported from literature, on chip optical links could now achieve propagation losses low as about 1.7 dB/cm.

Thanks to the earlier positive properties, a huge effort has been made by the NoC community to implement the *photonic NoC*, where light is coupled on chip, modulated at the transmitter, and guided to the receiver through integrated waveguides. The main challenge in this sense regards the integration of laser sources inside a chip. Since a huge silicon area is required, laser source is usually positioned out of the chip. Authors in [85] propose in fact an hybrid 3D nanophotonic NoC consisting in a multilayer chip, in which the first layers is used among processing elements and memory resources, while the top of the chip is used by photonic NoC communications.

Another issue in such NoCs is the limitation in terms of computing and storage capabilities. For this reason, a hybrid structure has been also proposed. In particular, in [85] a high-bandwidth circuit-switched photonic network is combined with a low-bandwidth packet-switched electronic network. While the electronic network carries small-size control (and data) packets, the photonic network transfers large-size data messages between pairs of cores.

4.3 Wireless NoCs

On-chip radio communication is a novel technique born initially for distributing clock signals into the chip for reducing clock skew related problems [38]. The main drawback until then, was the capability of integrating an antenna in a standard silicon substrate compatibly with CMOS technology. This is linked by the capability for transistors of operating at high frequencies. Tab. 4.2 shows the trend for the cut-off and oscillating frequency for MOS transistors as foreseen by the International Technology Roadmap for Semiconductors (ITRS) [5]. The meaning of such projection is that, over the

time, active devices can operate at higher and higher frequencies. Since, the dimension of an antenna has to be comparable with the wavelength, the first consequence of an higher operating frequency is that the dimension of an antenna will decrease. For instance, the dimension of a dipole antenna (simply formed by two conductors) operating at 60 GHz would have a length of $632 \times 2 \mu\text{m}$ when integrated in a silicon substrate [45]; while, if operating in 5.8 GHz, the dimension increases to $6.5 \times 2 \text{ mm}$, which is comparable with the entire die size. Furthermore, the scaling is not only limited to the antenna but it also affects the passive elements inside the main building blocks of the RF front-end which are responsible for a relevant fraction of its silicon area. For instance, as reported in [19], at 20 GHz the size of the inductor is approximately $50 \mu\text{m} \times 50 \mu\text{m}$ while at 400 GHz it can be reduced to $12 \mu\text{m} \times 12 \mu\text{m}$. Based on the above considerations, several research groups have proven the possibility of integrating every building block of the RF front-end (including the antenna) into the same chip [38, 67, 53]. Another important aspect that has been investigated regards the interferences introduced by metallic structures present in CMOS chips such as metal/dielectric wires and dummy fills. In fact, in [108] the authors reported a complete characterization of several integrated antenna by mean of a test chip in which several metal structures were inserted. Further, the effects of antenna rotation has been also discussed. Another interesting research study can be found in [45] that reports several antenna analysis obtained with the aid of a field solver.

In the context of on-chip communication, the capability of integrating an antenna with its transceiver into a silicon die [53] has lead several research groups on assessing the advantages of having long range wireless links upon the traditional wire-based NoC introducing the Wireless Network-on-Chip paradigm (WiNoC for short). The main advantage of using radio is the capability of a single hop communication. In fact, as discussed in the introduction of the thesis, traditional NoCs suffer of scalability problems (in terms of latency) due to the pipeline structure of electrical wires and routers. Since the research proposed in this thesis here falls in this context, it results mandatory to provide some notions about Wireless Network-on-Chip paradigm. In the following subsections several architecture will be shown.

As it will be discussed in the Chapters 5, 6 and 7, our research con-

tribution have been developed for different WiNoC architectures. Such architectures could be mainly divided, based on the adopted topology, in two main groups: *Mesh-Topology based* and *Small-World Network based* WiNoCs. Both kind of architectures are discussed in the following subsections.

4.3.1 Mesh-Topology Based WiNoCs

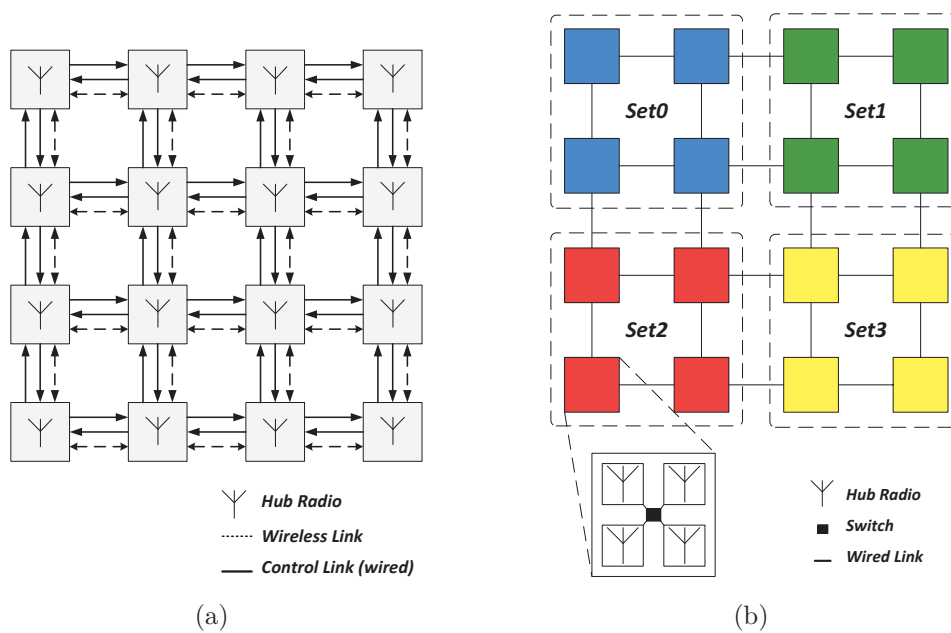


Figure 4.2: WiNoC Technologies: (a) McWiNoC; (b) iWise64 architecture.

These architectures are variants of the traditional mesh topology. In [109] authors proposed *McWiNoC*, an Ultra Wide Band (UWB) transceiver based WiNoC. Traditional router have been replaced by a new device named radio-hub which consists in a traditional router augmented with a transceiver and a respective antenna. As depicted in Fig. 4.2(a) each link between processing elements is wireless and can happen via a single or multiple-hops while control links are based on electrical wires. This NoC is demonstrated to achieve 65.3% average end-to-end latency reduction over a baseline mesh-based wire line NoC consisting of 64 cores.

Another interesting architecture is named *iWISE*. Such architecture is able to reduces power consumption and area overhead while improving performance in terms of network latency [32]. The latter, shown in Fig. 4.2(b),

consists in a 2D-Mesh divided in several clusters (16 in picture). Further, groups of clusters are partitioned in sets. Further, each set is tuned to a particular channel. In this way, communications inside the same cluster/set happen in traditional wires-based mode while inter-set communications happen via radio. In this manner, for a 64 core NoC the maximum hops count is equal to two. Such architecture is demonstrated to achieve a 2.5x performance increase and saving of 2x in terms of power for a 256-core system when compared to other competing architectures like networks with RF-Interconnect [19, 20] and WCube [51, 110, 99, 30].

4.3.2 Small-World Network Based WiNoCs

Both the schemes discussed in the previous sub-section, while improve performances in terms of latency, on the other hand a huge overhead in terms of required silicon area is introduced. This is due to the fact that such architectures require an high number of transceivers. In fact, WiNoCs transceiver and antenna are predominant in terms of consumed area if compared to a wireline router.

For facing with such overhead novel architectures inspired by complex network theory in conjunction with the on-chip wireless links were introduced in [41]. Networks with the small-world property [68] have a very short average path length, defined as the number of hops between any pair of nodes. The average path length of small-world graphs is bounded by a polynomial in $\log(N)$, where N is the number of nodes, making them particularly interesting for efficient communication with minimal resources.

Unfortunately, WiNoC proposed in [41] requires aggressive technologies since it uses carbon nanotube as antenna. Other interesting WiNoCs implementing the small-world property while proposing CMOS compatible solutions, can be found in [28, 104] and [79] named *HmWNoC* and *mSWNoC* respectively. Such architectures will be described in the following part of this subsection.

HmWNoC

Hierarchical mm-Wave NoC architecture (*HmWNoC*), as depicted in Fig. 4.3, it consists in a two layer network made of multiple small clusters of neigh-

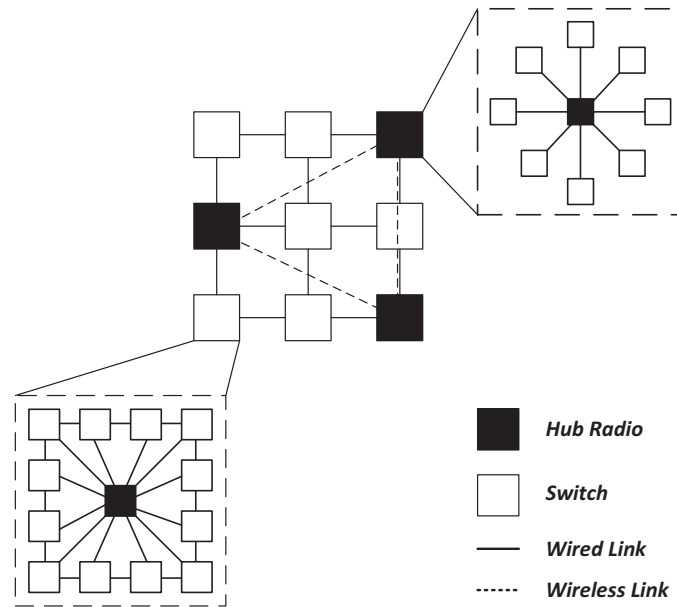


Figure 4.3: HmWNoC architecture.

boring switches called subnets interconnected by means of an upper level network. Each subnet and the upper level network can have a proper own topology. In a such architecture, *PEs* which belong to the same subnet are connected using traditional wires. Conversely, if a source and destination belong to different subnets, a communication can happen using a Hubs. Such element is essentially a big router which connects not only all router inside a specific subnet, but also several other Hubs. A traditional wire based Hub could be further augmented with a transceiver to enable wireless communication. In this case an Hub become a device named Radio Hub (RHub) or wireless interface (WI). Since only few Hubs become radio-hubs, the overhead introduced by implementing transceiver results mitigate. In a such hierarchical topology, routing algorithm works as follows. If a processing elements need to communicate with a core present in the same subnet, the communication happens as in traditional NoC. If a communication involves the upper layer, two situation could happen: 1) source and destination hubs are very close. In this case communication is wireline. 2) Source and destination are placed very far. In this case packets will be transmitted wireless.




Figure 4.4: mSWNoC architecture.

mSWNoC

The main limitation of *HmWNoC* is due to the presence of big Hubs. In fact, since the latter enable both inter-subnet and intra-subnet communications, these should have a huge number of ports which is translated to an additional overhead in terms of silicon area and dissipated power. mm-Wave Small-World NoCs (*mSWNoC*) try to mitigate such overhead by introducing a topology in which the wireline links between switches are established following a power-law distribution [79, 100] as reported in the Eqn. 4.1,

$$P(i, j) = \frac{l_{ij}^{-\alpha} f_{ij}}{\sum_{\forall i} \sum_{\forall j} l_{ij}^{-\alpha} f_{ij}} \quad (4.1)$$

where $P(i, j)$ is the probability of establishing a link between two switches i and j separated by a Euclidean distance of l_{ij} which is proportional to the distance raised to a finite power α . Further, f_{ij} is the frequency of traffic interaction between the switches i and j . As results of such distribution, more frequently communicating switches (according to the application) have a higher probability of having a direct link. Further, after that wireline links have been established, an optimal number of wireless interfaces will be placed to enable a long range communication. The main drawbacks for msWNoC architectures is the adopted routing algorithm. Since an irregular network is established, a congestion-aware adaptive layered shortest path routing (A-LASH) has been proposed [54, 79] to achieve an efficient, deadlock-free, dis-




Figure 4.5: Normalize energy (a) and bandwidth (b) in function of the number of WIs [28].

tributed and adaptive routing policy.

Placement of WIs

For both *HmWNoCs* and *mSWNoCs* the number of wireless interfaces (WIs) which comprise antenna and transceivers (transmitter + receiver) and their placement is could be computed by mean of an off-line computed optimization algorithm [18, 28]. The *WIs* introduce hardware overhead, and hence, their number should be limited without significantly compromising the overall performance. In particular in [28] the authors perform an optimization by defining two metrics. The former is defined, for a N Hubs network, as follows:

$$\mu = \sum h_{ij} \cdot \frac{f_{ij}}{(N - N^2) \cdot \sum f_{ij}} \quad (4.2)$$

where h_{ij} is the distance (in hops) between a generic source i and a destination j , and f_{ij} is the frequency of traffic interaction, as defined for the Eqn. 4.1. The second metric needed to optimize the cost is given by:

$$Cost(WIs) = A + P + L \quad (4.3)$$

where A , P , and L are normalized area, power, and wireless channel access delay overheads, respectively, arising from the WIs. Introducing metric defined by equations 4.2 and 4.3 multi-objective optimization will finally provide optimal number of WIs. As shown in Fig. 4.5(a)a-b, the optimal number




Figure 4.6: The zigzag antenna.

of WIs for a 64 core *mSWNoCs*. As can be noted the optimal number of interfaces in terms of performance an energy consumption is 12.

Once that optimal number of WIs has been established, their placement is crucial for optimum performance gain as it establishes high-speed, low-energy interconnects on the network.

To perform such optimization another metric μ could be introduced as follows:

$$\mu = \sum_{ij} h_{ij} f_{ij} \quad (4.4)$$

As results, WIs placement strongly depends on how the application is mapped on the NoC.

4.3.3 Physical Layer Management

In WiNoCs architectures new concepts that come from radio frequency are introduced. In fact, not only topological consideration are important. Other aspects that require a further investigation are, for example, the choose of the antenna, the radio frequency band and the transceiver architecture. In this subsection we will give some of these notions.

- **Integrated Antennas:** A possible classification of WiNoCs can be made on the basis of the portion of electromagnetic spectrum used for data transmission such as UWB [109] (few GHz), mm-wave [29, 32, 27, 26] (tens of GHz), sub-THz [51] (hundreds of GHz), and THz [41, 7] NoC. In mm-wave WiNoCs, zigzag antenna (Fig. 4.6) is considered as the best candidate solution for on-chip antenna [30]. A zigzag antenna for the mm-wave, can be designed and characterized with yet consolidated techniques and knowledge such as the use of field solvers. Furthermore, the use of regular topologies, like 2D meshes, allows the

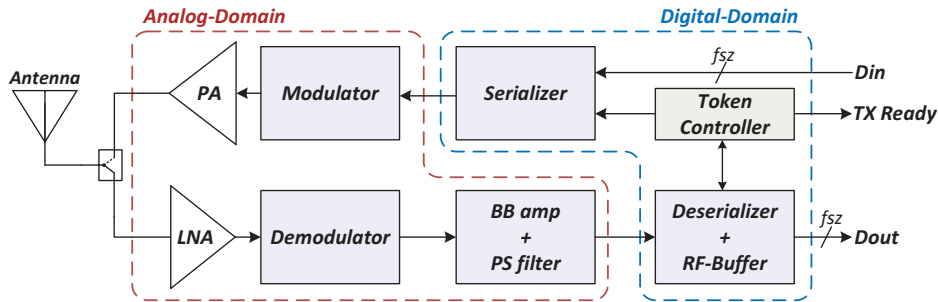


Figure 4.7: OOK Single Channel WiNoC transceiver.

exploitation of symmetries that simplify their characterization. For frequencies on the range of THz, authors in [41, 6] use carbon nanotube or graphene based antennas. Graphene-based antenna assures working frequency in the Terahertz band while utilizing less chip area for antennas as compared to the metallic counterparts.

- Modulation Scheme:** In this context, the most used modulation technique is the Amplitude Shift Keying or On Off Keying (ASK-OOK) [29, 32, 30]. Although, for a given bit error rate (BER), the ASK-OOK modulation requires a higher transmitting power than that required by other modulation techniques (*e.g.*, the Quadrature Amplitude Modulation (QAM) [22]), and has a poor spectral efficiency, its hardware implementation is simple (low area overhead as compared to QAM) and tailored to be applied in the on-chip context.
- Medium Access Mechanism (MAC):** Since wireless medium is shared among multiple transceivers, an access control mechanism must be implemented. In particular, a token based mechanism is used thanks to its implementation simplicity [92, 56]. The token is circulated as a wireless flit over the wireless medium in a round-robin sequence between the WIs.
- Transceiver architecture:**

in Fig. 4.7 is depicted a generic WiNoC OOK transceiver. Essentially it is based on two separated devices: the transmitter and the receiver. Both components share a common antenna by means of an RF switch. The transmitter is responsible to adapt a base band digital signal to the wireless medium. As shown in picture, a token controller is present

Table 4.3: Emerging NoCs comparison [15].

	3D NoCs	Photonic NoCS	WiNoCs
Design Requirements	Multiple layers with active devices	Silicon photonic components	On-chip metal or CNT/Graphene-based antennas
Performances	Low power and delay in vertical directions	Very high speed and Power independent from data-rate	One Hop capability, low power for high distances
Reliability	Vertical via failure	Temperature sensitivity of photonic components	Noisy wireless channel
Challenge	Heat dissipation due to higher power density	Integration of on-chip photonic components	Low power transceivers with smaller area occupation

to assure that the wireless channel is not busy at the moment of a transmission. If the channel is free incoming flit will be converted in a serial fashion by mean of the serializer. Finally, an OOK modulator converts data in an higher frequency signal that will be delivered to the antenna via a power amplifier (PA). The structure of the receiver is the opposite of the transmitter. Radio frequency signals will be converted in a baseband stream of data with a demodulator. A deserializer converts a serial stream in to a flit.

Many transceivers designed for WiNoC applications can be found in [106]. In particular, transceiver in [106] has an area of occupies an active area of 0.077 mm^2 while can reach an energy efficiency of bit-energy efficiency of 1.2 pJ/bit for a data rate of 16 Gbps . The just mentioned transceiver works in the range of mm-Wave on a single channel. As explained before, a multichannel approach or an efficient modulation scheme results in a complex transceiver [82] which is translated in a higher area. A trade-off between area and achievable bandwidth is thus, established. For this reason, a novel architecture named *DWiNoC* based on antenna with higher directivity has been proposed [62]. In such architecture, thanks to the fact that an antenna performances is heavily spatial, concurrent transmission could be achieved based on the reciprocal position of sources and destination are possible. This last property will further investigated and exploited by a technique proposed in this thesis for reducing transmitting power (Chapter 6).

4.4 Comparative Analysis

In this chapter three different emerging technologies have been introduced. In Tab. 4.3, a comparative analysis is reported. As can be noted, Photonic NoC seems to be the best candidate in terms of speed and power performances. Unfortunately, the difficulty to integrate photonic components in existing CMOS chips makes its choice unpractical. 3D NoCs, on the other hand, is an already available solution but is predicted that it does not satisfy performance requirements of future CMPs ,especially due to the difficulty of dissipating power consumption (Only top and bottom layer can be used to dissipate a temperature hotspot). For the reasons mentioned above, we believe that WiNoC solution, thanks to fact that is a CMOS-compatible solution, is the best candidate in the short and mid-period. It should be pointed that other interesting solutions could created by considering hybrid structures. An example could be found in in which the authors [91] exploit both advantages of 3D NoC and WINoCs to create a multilayer NoC with inductive couplers which can resolve the problem of fabrication due to vertical vias proper of 3D technology. Since power consumption of WiNoCs still remain a problem even if such hybrid structures are used, with our research we aim to face the problem of the transceiver power dissipation by introducing three different techniques. In particular, a former technique try to tune on-line transmitting power based on the position of source and destination 5 while another technique, exposed in chapter 6, reduces power by exploiting the directivity of antennas. Finally a smart transceiver was designed (chapter 6) to reduce power when a transceiver is not a recipient of a message. With the obtained energy savings, as reported in chapters 5, 6 and 7, we believe that WiNoCs architectures augmented with techniques presented in this thesis is the best solution in terms of energy efficiency and achievable low latency.

Chapter 5

Tunable Transmitting Power in mm-Wave WiNoC Architectures

As mentioned in the previous section, scalability issues in traditional NoCs have been solved by the introduction of WiNoC architectures. An open research point in this context remains in terms of power consumption. The major contribution is due to the radio transmitter front-end connected to the antenna. For instance, in [105] the transmitter is responsible for about 65% of the overall transceiver power consumption, while in [25] such contribution is more than 74%. Previous work in the context of WiNoCs are based on an architecture of the transmitter in which the transmitting power is kept constant (regardless the distance of the destination node), and able to guarantee a given reliability level (in terms of bit error rate, BER) in the worst case.

In this dissertation we will show a mechanism for improving the energy efficiency of the transmitters in WiNoC architectures. The basic idea is allowing the transmitter to run-time set its transmitting power based on the reliability requirements and the destination node of the current communication. We provide a systematic approach that, under a reliability constraint (given in terms of maximum BER) and for each antenna, allows to determine the optimal transmitting power for each destination node. The optimal transmitting power is off-line computed by using an accurate 3D field solver for a limited number of measurements. The obtained power figures are then

used for configuring the proposed variable gain controller which is responsible for driving the power amplifier connected to the transmitting antenna. The proposed technique is general and is agnostic as respect to the underlying WiNoC architecture. The proposed mechanism, have been explored for different WiNoC architectures, described in the previous section, such as iWise64 [32], McWiNoC [110], and HmWNoC [28]. Results show the effectiveness of the proposed technique in improving the energy efficiency (with energy savings up to 50%) without any impact on performance and with a negligible overhead in terms of silicon area. In addition, we show that, by exploiting the new degree of freedom provided by the application of the proposed mechanism during the mapping process, it is possible to further improve the energy metrics. It should be pointed out that, although the dynamic tuning of the transmitting power is a well known technique in the context of general wireless networks, for the best of our knowledge, this is the first work in which it is applied in the on-chip context. In this chapter we show the feasibility of using a runtime tunable transmitting power technique in the wireless interfaces of WiNoC architectures. We show the achievable advantages in terms of energy saving with a negligible impact on area figures and without affecting the overall communication performance and reliability metrics.

5.1 Adaptive Transmitting Power Transceiver

This section presents the proposed adaptive transmitting power transceiver which adaptively determines the optimal transmitting power, based on the packed destination address, under reliability constraints expressed in terms of maximum allowed communication bit error rate (BER).

5.1.1 Variable Gain Amplifier Controller

Traditional transceivers in WiNoC architectures use the same transmitting power regardless of the distance (location) of the destination node. In fact, the transmitting power is set for the worst case under a reliability (*i.e.*, maximum BER) constraint. We propose to runtime select the minimum transmitting power based on the physical location of the destination node of the current communication. Of course, the selected minimum transmitting

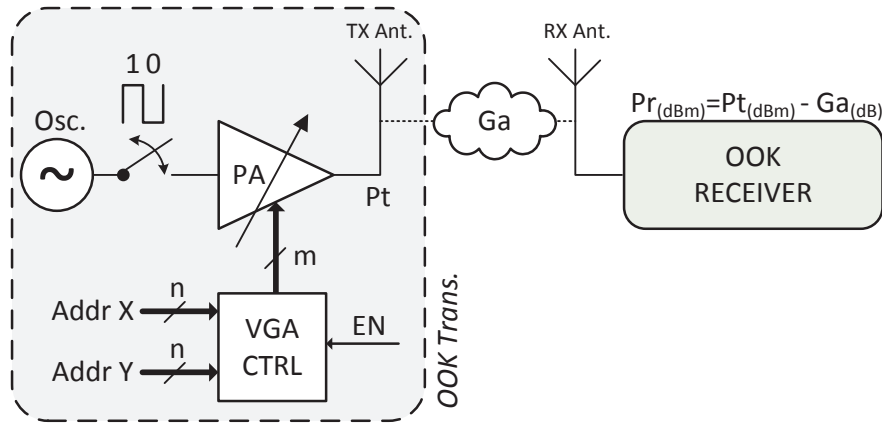


Figure 5.1: Scheme of the proposed adaptive transmitting power transceiver.

power must be high enough to meet the communication reliability constraints in terms of BER.

The general scheme of the proposed adaptive transmitting power transceiver is shown in Fig. 5.1. As compared to a traditional transceiver, it makes use of a tunable power amplifier (PA) controlled by a variable gain amplifier (VGA) controller. In the rest of the chapter, we consider the architecture of the PA presented in [25] which allows several transmitting power steps. Although dynamically tuning the transmitting power is an established technique in the context of radio communications (*e.g.*, mobile phones, wireless sensors network, *etc.*), its implementation requires sophisticated controller policies hardly replicable in the WiNoC domain. Thus, the proposed VGA controller uses the destination address of the packet for accessing a look-up table containing the configuration words used for configuring the PA. For a given destination, the associated configuration word enables the PA to use the minimum transmitting power to reach that destination by ensuring a specific reliability level in term of BER. Such optimal transmitting power is computed offline as it will be discussed in the next subsection.

5.1.2 Determining the Minimal Transmitting Power under a BER Constraint

The required transmitting power depends on many factors, including, the type of modulation, the transceiver noise figure, and the attenuation introduced by the wireless medium. Let us consider Fig. 5.2 which shows a




Figure 5.2: Friis transmission equation: geometrical orientation of transmitting and receiving antennas. As indicated, considering a spherical coordinate system, ϕ is the azimuthal angle in the XY plane, where the X axis is 0° and Y axis is 90° . θ is the elevation angle where the Z-axis is 0° , and the XY plane is 90° .

transmitting antenna with an output power P_t and a relative angle respect the receiving antenna (θ_t, ϕ_t) , and a receiving antenna, located at distance R , with a relative angle respect the transmitting antenna (θ_r, ϕ_r) . The fraction of the transmitting power that reaches the terminal of the receiving antenna, P_r , can be computed by the Friis transmission equation, Eqn. (5.1), valid when $R > 2D^2/\lambda$, where D is the maximum dimension of antenna (axial length in our case) and λ is the wavelength.

$$G_a = \frac{P_r}{P_t} = e_t e_r \frac{\lambda^2 D_t(\theta_t, \phi_t) D_e(\theta_r, \phi_r)}{(4\pi R)^2} \cdot (1 - |\Gamma_t|)(1 - |\Gamma_r|) |\hat{\rho}_t \cdot \hat{\rho}_r| \quad (5.1)$$

where:

- e_t and e_r are the efficiencies of the transmitting and receiving antenna, respectively. These parameters mainly represent the signal losses in the silicon substrate.
- D_t and D_r are the directivities of the transmitting and receiving antenna, respectively. They quantify how much better the antenna can transmit to or receive from a specific direction.
- λ is the effective wavelength. For an IC substrate, it is estimated by using the material properties of the top IC layers (silicon dioxide $\epsilon_r = 3.9$) [45].

is a parameter that describes how much of an electromagnetic wave is reflected

- $|\Gamma|$ is the reflection coefficient which quantifies the portion of the transmitting/receiving power that is reflected by an impedance discontinuity in the transmission medium (ideally $|\Gamma| = 0$).
- $|\hat{\rho}_t \cdot \hat{\rho}_r|$ takes into account the polarization status of the emitted EM wave (ideally, it is equal to one).

Eqn. (5.1) highlights the parameters which determine the gain G_a . It represents a first order model of the wireless channel which is valid for free-space communications. Although second order effects, including, wave reflections due to metal structures, and multi-path effects, are not modelled by Eqn. (5.1), Friis equation is a good starting point for understanding which parameters affect the attenuation. [108] presents a detailed study on the the propagation of radio waves in an on-chip context and confirms the effect of the directivity and the distance in the Friis formula. Since, as discussed above, the attenuation cannot be easily estimated by means of mathematical models, in this work the computation of G_a is carried out by means of Eqn. (5.2).

$$G_a = \frac{P_r}{P_t} = \frac{|S_{12}|}{(1 - |S_{11}|)(1 - |S_{22}|)}, \quad (5.2)$$

where, S_{11} , S_{12} , and S_{22} are the scattering parameters. Such parameters are not predicted but obtained by using accurate field solver simulation tools [38] or direct measurements from realized prototypes by means of a network analyzer.

Using Eqn. (5.2) it is possible to estimate the signal attenuation due to the wireless medium. Since the communication reliability is related to the energy per bit, E_b , spent to reach the receiver's antenna, we can determine the power required by the transmitter for each value of attenuation G_a . In particular, for the ASK-OOK modulation the bit error rate can be computed as:

$$BER = Q\left(\sqrt{\frac{E_b}{N_0}}\right), \quad (5.3)$$

where N_0 is the transceiver noise spectral density and the Q function is the tail probability of the standard normal distribution defined by Eqn. (5.4).

$$Q(x) = \frac{1}{\sqrt{2\pi}} \int_x^\infty e^{-\frac{y^2}{2}} dy. \quad (5.4)$$

Since $E_b = P_r/R_b$, where P_r is the power received at the terminal of the receiver antenna while R_b is the data rate, we can compute the required transmitting power for a given data rate and BER requirement and for a given transceiver's thermal noise as:

$$P_r = E_b \cdot R_b = [Q^{-1}(BER)]^2 N_0 R_b, \quad (5.5)$$

where Q^{-1} is the inverse of the Q function.

Thus, the minimum transmitting power to reach a certain receiver guaranteeing a maximum BER can be computed as:

$$P_t(dBm) = P_r(dBm) - G_a(dB), \quad (5.6)$$

where $P_r(dBm)$ is given by Eqn. (5.5) while $G_a(dB)$ is computed by using a field solver with the Friis formula when power is expressed in dBm.¹

5.1.3 Overall Flow

Now we present the basic steps needed for determining the optimal transmitting power for each node pairs. For the sake of clarity, we consider the case in which radio hubs are arranged on a mesh topology. This makes more simple the characterization of the antennas as symmetries can be exploited.

1. Computing the attenuation map. For each pair <transmitting antenna, receiving antenna>, extract the scattering parameters S_{11} and S_{22} and compute the gain by means of Eqn. (5.2). In this research work we used an accurate field solver simulator for estimating the scattering parameters S_{11} and S_{22} , however, in case of the availability of a test-chip, they can be directly measured by means of a network analyzer [67].
2. Computing the Power map. For each pair <transmitting antenna i , receiving antenna j >, based on the required transmission data rate and the maximum allowed BER, use Eqns. (5.5-5.6) for computing the minimum transmitting power that met the BER constraint. Let us indicate this transmitting power value with $PM(i, j)$.
3. Determining the power steps. Let n be the number of desired power steps and PM_{min} , PM_{max} the minimum and maximum value of PM ,

¹The absolute power, P , can be expressed in dBm by $P_{dBm} = 10 \cdot \log(P \cdot 10^3)$

respectively. The set of power steps $PS = \{ps_1, ps_2, \dots, ps_n\}$ is defined by dividing the interval $[PM_{min}, PM_{max}]$ in n equally spaced levels for which the i -th power step is:

$$ps_i = PM_{min} + (i - 1) \times \frac{PM_{max} - PM_{min}}{n - 1}.$$

4. Configuring the VGA controller. Upload the look-up table in each VGA controller as follows. Let LUT_i be the look-up table of the VGA controller into radio hub i . $LUT_i(j)$ encodes the power step to be used to transmit to radio hub j . Such power step is selected as the minimum $ps \in PS$ such that $PM(i, j) \leq ps$.

In the next section we assess the effectiveness of the proposed technique in terms of communication energy reduction.

5.1.4 The Mapping Problem

Several works in literature have shown the effectiveness of mapping techniques for improving different metrics including performance and energy consumption. The role played by the mapping becomes even more important in the context of WiNoCs due to the possibility of exploiting an additional degrees of freedom, including, the association between the radio hub and the cluster of concentrated cores, the directionality of the antenna, the number of radio channels to be used, etc. In the research presented in this thesis, we explore one of the aforementioned new mapping dimensions, namely, the association between the radio hub and the cluster of concentrated cores. Specifically, the mapping problem, shown in Fig. 5.3, is formulated as follows.

Let $NG = G(R, RR, L_R, L_{RR})$ be the *network graph* where R is the set of routers, RR is the set of radio routers, L_R is the set of links connecting the routers in R , and L_{RR} is the set of links connecting the radio routers in RR with the routers in R . We assume that all the links $l_R \in L_R$ and $l_{RR} \in L_{RR}$ have the same bandwidth capacity cap and the same energy consumption per bit e_l . Let PE the set of processing elements. We assume direct networks for each there is a router for each processing element.

Let $e_r(rr_s, rr_d)$, with $rr_s, rr_d \in RR$, be the *radio transmission energy function* which provides the minimum transmission energy per bit for a radio communication from radio router rr_s to radio router rr_d .

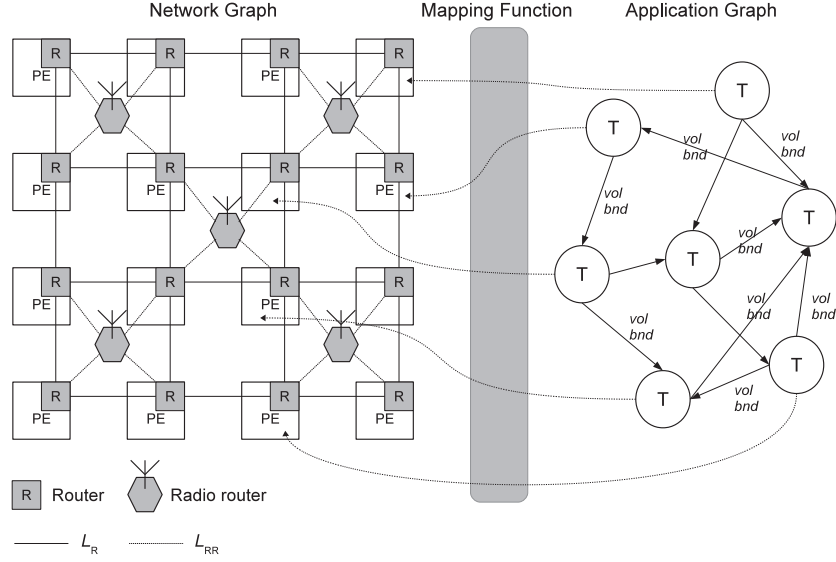


Figure 5.3: The mapping process.

Let $AG = G(T, C)$ be the *application graph* where T is the set of tasks and C is the set of communications among tasks. Let $bnd(c)$ and $vol(c)$ be the *communication bandwidth* (in bit/sec) and the *communication volume* (in bit) of communication $c \in C$, respectively.

Based on the above definitions, the mapping problem can be formulated as follows. Find a *mapping function*, $map : T \rightarrow PE$, such that the communication energy is minimized and the bandwidth constraints are met. The communication energy, E , is the the product between the communication volume and the total energy per bit spent on links and radio transmissions over all the communications. It is computed as follows:

$$\begin{aligned}
 E = & \sum_{\substack{c=(t_s, t_d) \in C \\ pe_s = map(t_s) \\ pe_d = map(t_d)}} vol(c) [|LT(pe_s, pe_d)| e_l + \\
 & + \sum_{(rr_s, rr_d) \in RRP(pe_s, pe_d)} e_r(rr_s, rr_d)],
 \end{aligned} \tag{5.7}$$

where $LT(pe_s, pe_d)$ returns the set of links traversed for the communication between pe_s and pe_d , and $RRP(pe_s, pe_d)$ returns the set of radio router pairs (transmitter, receiver) involved in the communication between pe_s and pe_d .

The bandwidth constraints refer to the fact that the aggregated band-

width on links cannot exceed their capacity. That is:

$$\sum_{\substack{c=(t_s,t_d)\in C \\ pe_s=map(t_s) \\ pe_d=map(t_d)}} bnd(c) \times PT(pe_s, pe_d, l) \leq cap \quad \forall l \in L_R,$$

where $PT(pe_s, pe_d, l)$ is the pass-through function which returns 1 if l belongs to the routing path for the communication between pe_s and pe_d and 0 otherwise. That is, $PT(pe_s, pe_d, l) = 1$ if $l \in LT(pe_s, pe_d)$.

Differently from the traditional mapping techniques proposed in literature [], here, the mapping selection depends also by the location of the radio routers which is accounted by the radio transmission energy function e_r in Eqn. (5.7). Such additional degree of freedom results in new opportunities for energy optimization as it will be shown in the experiments section.

5.2 Experiments

In this section we present the results of experiments in which a WiNoC architecture implemented into a 20 mm × 20 mm silicon die is considered. A zigzag antenna has been accurately modeled and characterized with Ansoft HFSS [1] (High Frequency Structural Simulator). HFSS is a leading commercial finite element method (FEM) field solver which simulates 3D structures and produces S-parameters and radiation patterns. We considered an high resistivity $\rho = 5 \text{ K}\Omega\text{cm}$ SOI with a substrate thickness of 350 μm and 30 μm for the oxide (SiO_2). The antennas are situated at an elevation of 2 μm from the substrate, compatibly with the guidelines reported in [84] for reducing the interference with others metal structures ([84] demonstrates that the interference due to other metallic structures is negligible by following such rules). The zigzag antenna has a thickness of 2 μm and an axial length of $2 \times 340 \mu\text{m}$ for operating at around 60 GHz. The same setup has been used in [66].

From HFSS simulation we obtain the scattering parameters (S_{11} and S_{12}) used for computing the Friis formula and then for calculating the attenuation introduced by the wireless medium. In particular, S_{11} is also used for determining the antenna bandwidth as discussed in the following subsection.




Figure 5.4: S_{11} parameter of the zigzag antenna. The bandwidth is the range of frequencies below -10 dB.

5.2.1 Bandwidth and Radiation Pattern

Fig. 5.4 shows the S_{11} parameter which quantifies the portion of transmitting power reflected to the power amplifier due to impedance mismatch (50Ω). Based on a thumb rule [10], it can be assumed that the antenna impedance matches with the transceiver when, at the operating frequency, the S_{11} is less than -10 dB. We used S_{11} for defining the antenna bandwidth because outside of the range of frequencies for which $S_{11} < -10$ dB, the antenna not only does not work properly as transducer but it could affect the physical integrity of the final stage of the PA.

Thus, looking at Fig. 5.4, a bandwidth of about 16 GHz is enough for providing a data rate upper bound of 8 Gbps with ASK-OOK modulation. Let us indicate with B_W such bandwidth, the antenna relative bandwidth is:

$$B_r = \frac{B_W}{f_c} = \frac{16 \text{ GHz}}{59 \text{ GHz}} = 0.27$$

where f_c is the resonance frequency. Such information is useful for determining at which resonance frequency we should design the antenna for obtaining data rates higher than 8 Gbps, or if we are interested in having more bandwidth for a frequency division multiplexing (FDM). For instance, for 4 channels with a data rate of 8 Gbps, we can design an antenna with a resonance frequency of at least:

$$f_c = \frac{B_W}{B_r} = \frac{4 \times 16 \text{ GHz}}{0.27} = 237 \text{ GHz}$$

which is obtainable by properly scaling the dimensions of the antenna (mainly the axial length). Another important result from simulation is the normal-




Figure 5.5: Radiation pattern for a zigzag antenna at the horizon ($\phi = 90^\circ$, continuous line) and at the elevation of maximum radiation ($\phi = 35^\circ$, dashed line). $\theta = 0^\circ$ is the direction parallel to the antenna's main axis while $\theta = 90^\circ$ is the orthogonal direction. According to Fig. 5.2, we assume the antenna situated upon the XY plane.

ized radiation pattern shown in Fig. 5.5. The radiation pattern is a polar representation of the directivity represented by the term D in Eqn. (5.1). As it can be observed, the best performance is obtained when the antenna transmits or receives along the direction of its main axis. With this information we can have an idea of the attenuation in a particular direction Eqn. (5.1) as it will be shown in the next subsections.

5.2.2 Attenuation Maps

Let us consider a mesh based WiNoC formed by a set of T tiles and a radio hub for each tile. We analyze the attenuation of the signal transmitted by an antenna in a tile $t \in T$ as perceived by the other antennas located at tiles $T \setminus \{t\}$. In the experiments we considered $|T| = 16$ in which the distance between two antennas in the same axis is 2.5 mm.

Fig. 5.6 shows the attenuation G_a for a transmitting antenna located on tile t_0 , t_1 , t_4 , and t_5 . The other attenuation maps (*i.e.*, the attenuations when the transmitting antenna is located in other tiles) can be found by

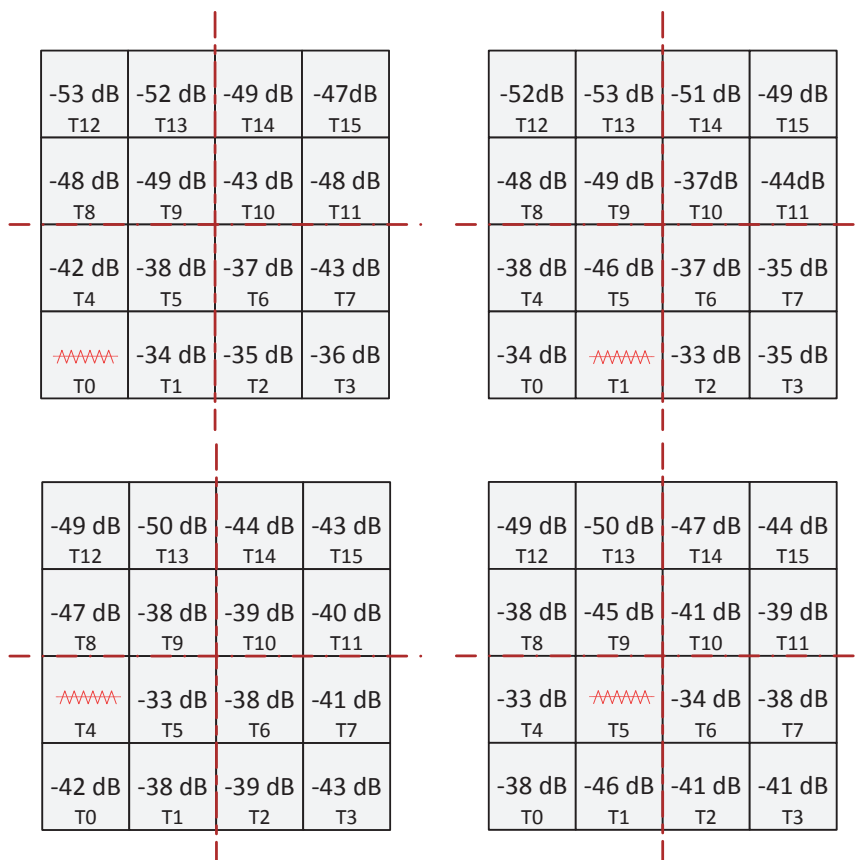


Figure 5.6: HFSS Simulation results: attenuation map (G_a) for the tiles t_0 , t_1 , t_4 and t_5 . The others map can be obtained considering the structure's symmetries.

symmetry. In fact, the antenna exhibits very different behavior when it is placed in different locations within the die [45]. Thus, the measures should be performed by considering all the possible positions for the transmitting and receiving antenna. Thanks to the symmetrical structure of mesh-based topologies, only four measures are needed in our case. For instance, the attenuation observed by a receiving antenna at tile t_{13} when the transmitting antenna is on tile t_{12} , $G_a(t_{12}, t_{13})$, is the same as observed by the receiving antenna located on tile t_1 when the transmitting antenna is on tile t_0 , $G_a(t_0, t_1)$. Similarly, we have $G_a(t_{15}, t_{14}) = G_a(t_0, t_1)$, $G_a(t_3, t_2) = G_a(t_0, t_1)$, and so on. In addition, $G_a(t_x, t_y) = G_a(t_y, t_x)$ for each $t_x, t_y \in T$.

As it can be observed from Fig. 5.6, the attenuation introduced by the wireless medium does not depend only by the relative distance between the radio hubs but it depends also by their relative orientation. For instance, $G_a(t_0, t_3) < G_a(t_0, t_4)$ although the distance between t_0 and t_3 is three times higher than the distance between t_0 and t_4 . This can be explained observing the radiation pattern in Fig. 5.5 in which the performance of the antenna increases as it transmits to or receives from its main axis direction.

In conclusion, the attenuation map is used for computing the maximum and minimum transmitting power for guaranteeing a certain reliability level. For the sake of example, let us consider a maximum BER of 3×10^{-14} and a data rate of 8 Gbps. From Eqn. (5.5), the power received by the receiving antenna must be -54 dBm. From the attenuation maps shown in Fig. 5.6, the maximum attenuation is -53 dBm. Thus, the transmitting power (which is maximum as this is the worst case) is computed by Eqn. (5.6) as $P_{t,max} = -54 - (-53) = -1$ dBm, that in linear scale is $P_{t,max} = 794 \mu\text{W}$. Similarly we can compute the minimum transmitting power. The minimum attenuation is -33 dBm, thus $P_{t,min} = -54 - (-33) = -21$ dBm, that in linear scale is $P_{t,min} = 8 \mu\text{W}$.

5.2.3 VGA Controller Analysis

Let us consider the architecture of the transceiver proposed in [25], also used in [32]. Such transceiver provides different transmitting power steps but, neither [25] nor [32] define the control policy for setting the appropriate power step. For the transceiver we estimate a power consumption of 7 mW to 23 mW for the minimum and maximum transmitting power, re-

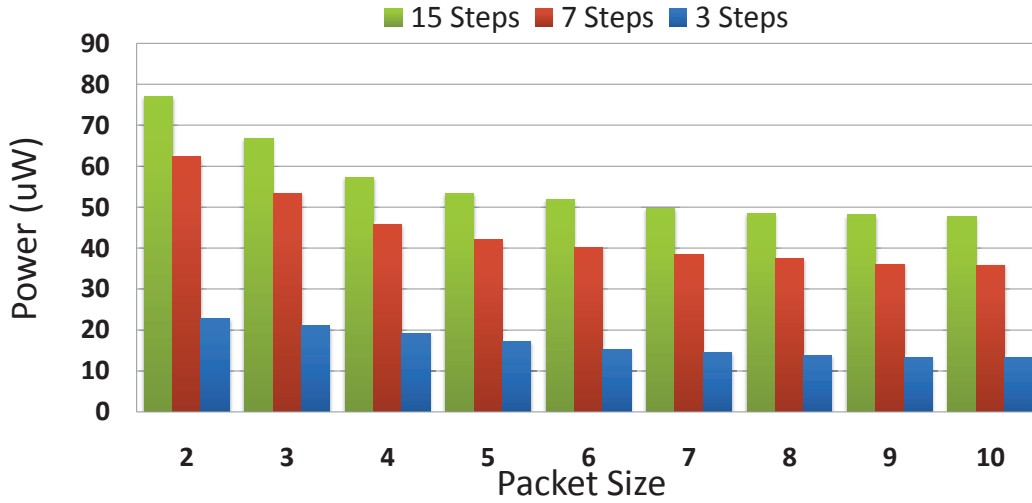


Figure 5.7: Average power dissipated by the VGA controller for different power steps and different packet sizes.

spectively. They corresponding to an energy per bit ranging from 0.42 pJ/bit to 1.4 pJ/bit.

With regard to the logic of VGA controller, it has been synthesized and evaluated by using Synopsys Design Compiler considering different number of admissible power steps (3, 7 and 15 power steps). Considering the gate-level implementation of the controller, the power analysis has been performed considering various test benches varying the size of packets. In fact, as packet size increases, the toggle rate of the VGA controller decreases as it is active only for the header flit of the packet. Fig. 5.7 shows the average power dissipation of the VGA controller for different packet size considering a 28 nm CMOS standard cell library from TSMC operating at 1 GHz. As it can be observed, for a 10-flit packet, the average power dissipation of the VGA controller is as low as 21 μ W for the 3-step implementation, and about 50 μ W for the 15-step implementation.

Fig. 5.8 shows the area and timing overhead due to the VGA controller for different number of power steps. With regard to the area overhead, it ranges from 50 μm^2 to 90 μm^2 for the implementations with 3 and 15 power steps, respectively. Timing results are shown in terms of FO4. In order to determine the set-up time for configuring the proper transmitters bias level (power step), the Digital to Analogue Converter (DAC) into the power control circuitry has been considered as lumped load for the generated gate-




Figure 5.8: VGA controller synthesis results: area and delay overhead.

level net-list and used as constraint during the synthesis phase.

In order to assess how the introduction of the VGA controller impacts the overall delay metrics of the radio hub, we consider the pipeline structure of the radio hub shown in Fig. 5.9. The radio hub is derived by the baseline router [24, 57] augmented with the proposed VGA controller. The transceiver is attached to its local port. For each pipeline stage, Fig. 5.9 reports timing information related to the critical path. As it can be observed, the VGA controller works in parallel while incoming flits are transferred to the serializer before the radio transmission. Thus, in terms of latency, the use of the

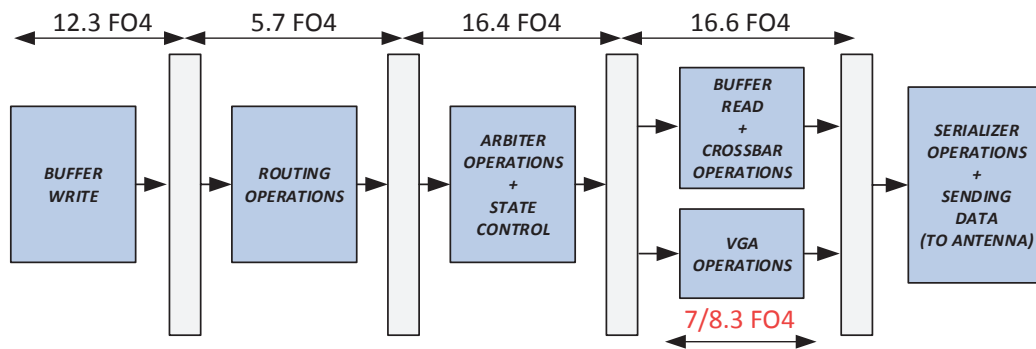


Figure 5.9: Pipeline of a conventional radio hub.




Figure 5.10: Area (a) and power (b) breakdown of the radio hub.

proposed technique does not affect the pipeline depth of the radio hub. In terms of clock frequency, the delay introduced by the VGA controller does not impact the critical path of the slowest stage (*i.e.*, buffer read and crossbar). For instance, the 15-step implementation of the VGA controller (*i.e.*, the slowest one among the three considered in this research), exhibits a delay of 8 FO4 which is far below the 16.6 FO4 delay exhibited by the buffer read and crossbar operations in the same stage. Finally, Fig. 5.10 shows the area and power breakdown of the radio hub. As it can be observed the VGA controller accounts for a negligible fraction of the overall area and power budget which is less than 0.05%.

5.2.4 Total Energy Saving in Mesh-Topology Based WiNoCs

The effectiveness of the proposed technique is affected by the number of power steps provided by the VGA controller. For quantifying such impact, we apply the proposed technique on two different mesh topology based WiNoC architectures proposed in literature, namely, iWise [32] and McWiNoC [32]. Specifically, we compare the following NoC architectures:

1. Wire-line: A traditional 8×8 concentrated mesh, with clusters formed by 4 cores.
2. McWiNoC: The architecture described in [110] for a 8×8 mesh with 4 cores associated with each radio hub. This kind of architecture uses TDM multiplexing for the wireless medium. The entire bandwidth can be allocate for each communications due to the particular structure of the architecture.

3. Proposed McWiNoC: Like McWiNoC but augmented with the proposed VGA controller.
4. iWise64: The architecture described in [32] in which the entire bandwidth is divided in four different channels.
5. Proposed iWise64: Like iWise64 but augmented with the proposed VGA controller.

Power data presented in the previous subsection have been used for back-annotating a cycle accurate NoC simulator based on Noxim [35] which has been extended for simulating WiNoC architectures.

Assuming the wire-line NoC as baseline, Fig. 5.11 shows the overall communication energy saving for different SPLASH-2 benchmarks when the proposed VGA controller is applied to iWise and McWiNoC. In particular, we considered four versions of the VGA controller, namely, 3-, 7-, 15-, and INF-step, which refer to the considered number of power steps. Please notice that, the INF-step version is a theoretical case (*i.e.*, it represents an upper-bound in terms of energy saving) in which the transmission energy is tuned in a continuous, rather than discrete, fashion. As it can be observed, on average, iWise and McWiNoC are 22% and 12% more energy efficient than the traditional wire-line NoC. By using the proposed approach, the average energy saving increases, on average, by 50% and 46% for iWise and McWiNoC, respectively. As expected, the number of power steps impacts the energy saving but no relevant improvements are observed with more than 7 power steps. For this reason, in the rest of the experiments, we assume a VGA controller with 7 power steps if not otherwise specified.

5.2.5 Total Energy Saving in Small-World Network Based WiNoCs

In order to explore the impact of the proposed scheme in mm-wave small-world topology based WiNoCs (HmWNoC), we apply the proposed scheme to the HmWNoC architecture presented in [28]. Such HmWNoC is a two levels hierarchical network where the top-level is a mesh topology whereas the lower-level sub networks are star-ring networks. Since the upper network




Figure 5.11: Energy saving over a traditional wire-line NoC when the proposed VGA controller is applied on a iWise 64 architecture (a) and on a McWiNoC architecture (b).

is a mesh, the set-up used for obtaining the attenuation maps (*cf.*, Sec. 5.2.2), can be easily reused.

Fig. 5.12 shows the effectiveness of the proposed technique when it is applied to a HmWNoC architecture. We analyze different network configurations in which the number of radio hubs is made to vary and in which different network sizes are considered. Specifically, we analyze three different network sizes with 256, 576, and 1024 nodes (cores) and in which the number of radio hubs is made to vary from 1 to 16, 6 to 24, and from 8 to 32, respectively.

In terms of energy saving, as the number of radio hubs increases, the energy saving increases due to the fact that there are more opportunity for wireless communications in which the proposed technique gives its contribution in terms of energy saving. Further, as the network size increases, the energy saving becomes more sensitive to the number of radio hubs. Such behavior can be explained observing that, for a given number of radio hubs, as the network size increases, the size of the subnetworks increases as well, and the fraction of communications which involve the use of the wireless medium decreases. Thus, since the proposed technique affects only the wireless communications, its impact on energy figures decreases.

To make this trend more clear, Fig. 5.13 plots on the x-axis the *radio hub density* and on the y-axis the energy saving for each of the considered network configurations. With the term radio hub density we refer to the ratio between the number of radio hubs and the number of nodes (cores) of the network. As it can be observed, as the network size increases, the effectiveness of the proposed technique increases for the same radio hub density. The energy saving gap between the different network configurations increases as the radio hub density increases. Specifically, for large network sizes (1024 nodes), below a radio hub density of 1%, the effectiveness of the proposed technique is low (less than 10% of energy saving). Above such threshold, the energy saving rapidly increases. A similar behaviour is observed for the other network sizes although with different thresholds. It should be pointed out that, the analysis has been carried out under uniform traffic. The results for the other traffic scenarios have not been reported for the sake of brevity and because they brought to the same conclusions.




Figure 5.12: Percentage of energy saving when the proposed technique is applied to HmWNoC architectures with different size (number of nodes) and different number of radio hubs.




Figure 5.13: Energy saving vs. radio hub density when the proposed technique is applied to HmWNoC architectures with different sizes under uniform traffic.

5.2.6 Application Mapping

The way in which tasks are mapped into the NoC has a tremendous impact on performance and power metrics [83]. In fact, the possibility of tuning the transmitting power based on the location of the destination node can be seen as a new degree of freedom in the mapping problem which results in new opportunities for energy optimization. In this subsection we assess the improvement in energy saving when the GAMAP mapping technique [74] is used in conjunction with the proposed technique. We selected a subset of benchmarks from the SPLASH-2 benchmarks suite that have been simulated with Graphite Multicore Simulator [58] for extracting the communication patterns. Such communication patterns have been then used for determining the communication graphs which form the input for the considered mapping technique.

Fig. 5.14 shows the percentage communication energy saving (considering the wireline NoC as baseline) when the mapping is optimized. In particular, for both iWise and McWiNoC we analyzed three configurations as follows. 1) The proposed technique is not applied and a random mapping is used, 2)




Figure 5.14: Impact of the mapping on energy consumption. Energy saving over a traditional wire-line NoC when the proposed VGA controller is applied on a iWise 64 architecture (a) and on a McWiNoC architecture (b).

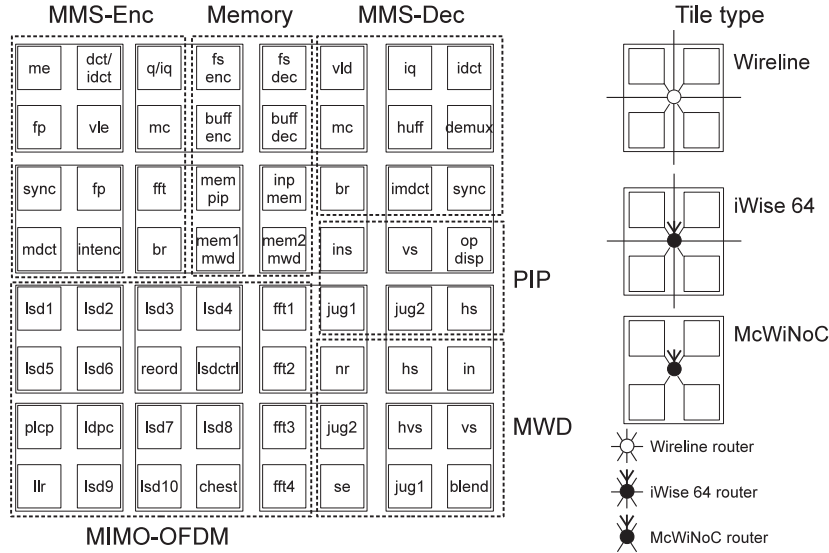


Figure 5.15: Heterogeneous system composed by a multimedia sub-system, a MIMO-OFDM receiver, a PIP and a MWD module.

The proposed technique is applied and a random mapping is used, and 3) The proposed technique is applied and the application mapping is optimized. The energy consumption in the case in which the random mapping is used is measured by averaging the energy consumption over 1,000 random mappings. As it can be observed, on average, the optimization of the mapping in conjunction with the proposed technique improves the energy efficiency by 72% and 62% for iWise and McWiNoC, respectively.

5.2.7 Case Study

Finally, as a case study, we consider a complex heterogeneous system shown in Fig. 5.15. The system is composed by a generic MultiMedia System which includes a H.263 video encoder, a H.263 video decoder, a MP3 audio encoder and a MP3 audio decoder [48], a MIMO-OFDM receiver [103], a Picture-In-Picture application (PiP) [49] and a Multi-Window Display application (MWD) [97]. We have mapped the application on both iWise and McWiNoC and assessed the energy saving when the proposed technique is used.

Fig. 5.16 shows the normalized energy consumption of the different architectures as compared to the wireline NoC. As it can be observed, the application of the proposed technique results in interesting energy saving up

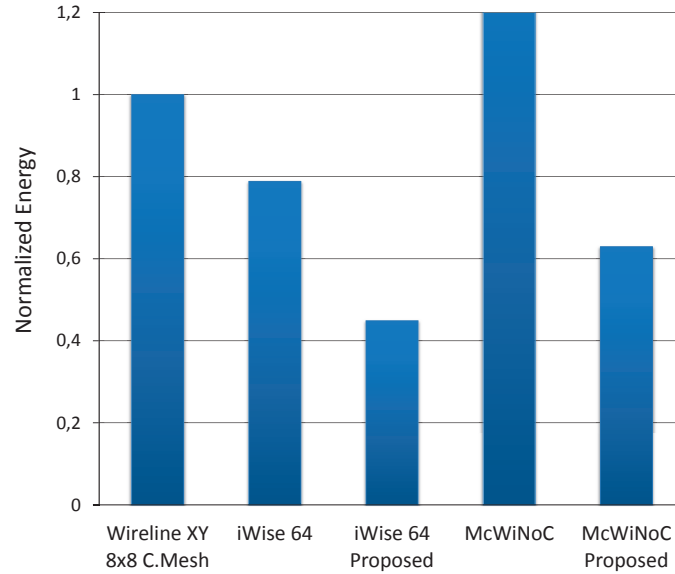


Figure 5.16: Normalized energy consumption for iWise64 and McWiNoC when the proposed technique is applied.

to 50% and 48% when applied to iWise64 and McWiNoC, respectively.

5.3 Conclusions

Emerging communication technologies like wireless NoC (WiNoC) are considered as a viable solution for facing the scalability and the energy consumption issues in many-core system architectures. Unfortunately, the transceiver of the radio hub in a WiNoC accounts for a significant fraction of the overall communication energy budget. In this chapter we have presented a reliability aware runtime tunable transmitting power technique for improving the energy efficiency of the transceiver in WiNoC architectures. The proposed technique is general and can be applied to any WiNoC architecture. The latter has been applied to three known WiNoC architectures, namely, iWise64 [32], McWiNoC [110], and HmWNoC [28]. The experimental results have shown important energy saving up to 60% without any impact on performance metrics. The hardware overhead, in terms of silicon area, introduced by the proposed technique is negligible as compared to the area of the transceiver (approx four order of magnitude less than the transceiver). We believe that the introduction of the proposed technique opens interest-

ing scenarios in several directions. For instance, as it will be seen in the next section energy reduction strategies might take into account the specific radiation patterns considering thus the orientation of the antennas as an additional degree of freedom for application specific optimization purposes.

Chapter 6

Exploiting Antenna Directivity in WiNoCs

As explained in the previous chapter, in a WiNoC the transmitting power could be tuned in order to save energy. From experiments (Sec. 5.2.2), it is quite clear that wireless attenuation depends not only by distance but, considering a specific antenna, also by direction. In fact, it is well known also from antennas theory that the behavior of the antenna strongly depends on the direction from/to which the signal is received/transmitted. Such behavior is described by a fundamental antenna parameter, namely, *antenna directivity*, which describes the variation of the transmitting/receiving signal intensity for different observation angles. The directivity effects, widely studied in the context of free space communications, have been recently investigated in the context of intra-chip communications [108]. Nevertheless, in the current WiNoC literature, the radiation pattern of the antenna is considered isotropic, that is, it is assumed that the antenna exhibits the same behavior irrespective of the transmitting/receiving directions. In fact, in the context of WiNoCs there are no works in literature that take into account the directivity effects, and the antenna orientation is left out from the set of design parameters to be explored.

In this chapter, we will see the impact of antennas orientation on energy metrics in WiNoC architectures. Based on such analysis, we formulate the problem of finding the antennas orientation in such a way to minimize the total communication energy in the following two cases. The case in which the information about the applications that will be mapped on the WiNoC and

their communication characteristics are not known, and the case in which they are known at design time. We refer to the first case as *general purpose* and the second one as *application specific*. Further, we also formulate the problem of finding the antennas orientation in such a way to minimize the transmitting power for the worst case. This latter problem is important in the case in which the WiNoC does not implement technique presented in the previous chapter. In such case, transmitting power is used for any communicating pairs irrespective of their position into the WiNoC. As it will be seen, experiments, carried out on state-of-the-art WiNoC architecture, such as HmWNoC [28] show that important energy saving, up to 80% can be obtained by properly set the orientation of the antennas.

6.1 Antenna Directivity Optimization

6.1.1 Antenna Directivity

As described in the previous chapter by Eqn. (5.1, page 66), the attenuation introduced by the wireless medium, strongly depends on the directivity function. Based on this, the orientation of the antennas in a WiNoC represents an interesting parameter to be explored.

The energy consumed by a wireless communication between a given transmitter and receiver pair, depends on their reciprocal location and orientations. Specifically, the energy consumed for wirelessly transmitting a bit of information from transmitter i to receiver j is:

$$E_{ij}^{tx} = \frac{P_{t_{ij}}}{\eta R_b}, \quad (6.1)$$

where η is the transmitter efficiency. Considering Eqns. (5.6) and (5.1), the Eqn. (6.1) can be written as:

$$E_{ij}^{tx} = \frac{P_r / D_t(\Phi_{ij}) D_r(\Phi_{ji}) \left(\frac{\lambda}{2\pi R_{ij}} \right)^2}{\eta R_b}. \quad (6.2)$$

Let ψ_i and ψ_j be a rotation of antenna i and antenna j as respect to the die plane, respectively. Thus, Eqn. (6.2) normalized by the constant terms can be written as:

$$\overline{E_{ij}^{tx}} = \frac{R_{ij}^2}{D_t(\Phi_{ij} - \psi_i) D_r(\Phi_{ji} - \psi_j)}. \quad (6.3)$$

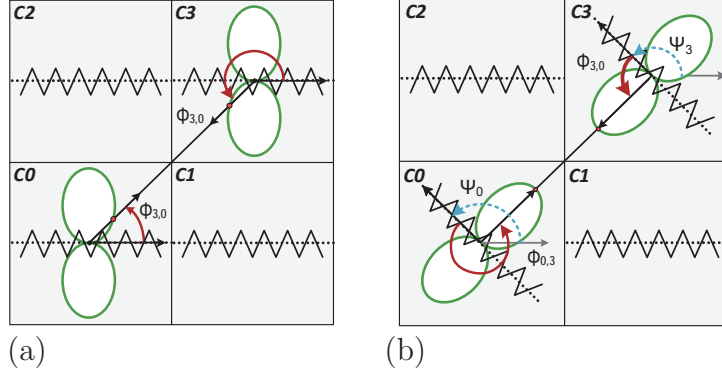


Figure 6.1: Reciprocal antennas orientation and directivity functions. The directivity between C0 and C3 improves from configuration (a) to configuration (b).

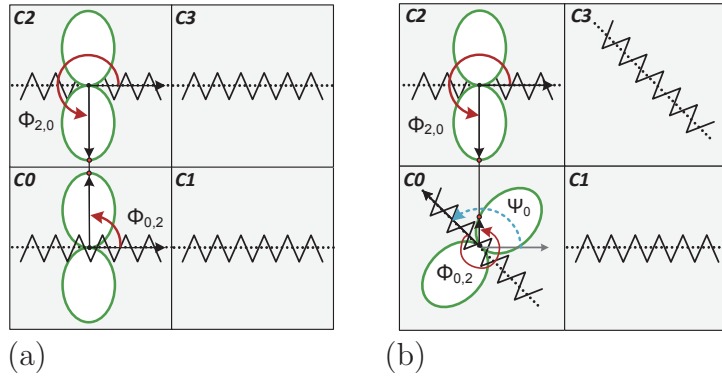


Figure 6.2: Reciprocal antennas orientation and directivity functions. Configuration (a) improves communication energy between tiles C0 and C2 whereas configuration (b) improves communication energy between tiles C0 and C3.

Thus, for minimizing the communication energy from i to j , it needs to determine a rotation ψ_i and a rotation ψ_j such that Eqn. (6.3) is minimized.

For the sake of example, let us consider Fig. 6.1(a) in which four antennas and their reciprocal orientations in the die plane are shown. The energy consumption of the communication between tiles C0 and C3 can be minimized by maximizing the directivity of their respective antennas. This is obtained by rotating the antennas in tiles C0 and C3 by the angles $\overline{\psi_0}$ and $\overline{\psi_3}$, respectively, as shown in Fig. 6.1(b).

It should be pointed out that, selecting a certain antenna orientation for improving the energy efficiency of a given transmitter/receiver pair, might negatively affects the energy figures of other transmitter/receiver pairs. For

instance, the directivity of the antennas in tiles C0 and C2 before the rotation is maximized [see Fig. 6.2(a)] and thus, their communication energy is minimized. However, after the rotation of the antennas in tiles C0 and C3 [see Fig. 6.2(b)], although the communication energy between C0 and C3 improves, that between C0 and C2 worsens.

6.1.2 Formulation of the Problem

Based on the above considerations, this subsection formulates the problem of minimizing the wireless communication energy by means of antennas orientation optimization. Specifically, three scenarios, namely, *application specific*, *general purpose*, and *worst case* scenarios will be analyzed.

Application Specific Scenario

In the *application specific scenario* it is assumed that communication traffic information are available at design time. Let V_{ij} be the traffic volume (in bits) from radio hub i to radio hub j . Let $E_{ij}(\Psi)$ be the energy consumption for transmitting one bit from radio hub i to radio hub j when the antennas are oriented based the orientation vector Ψ . The i -th component of Ψ represents the orientation of the antenna of radio hub i . $E_{ij}(\Psi)$ is computed by Eqn. (6.3). The total normalized wireless communication energy can be computed as:

$$E_{tot}^{(as)}(\Psi) = \sum_{i=0}^N \sum_{j=0}^N V_{ij} \times E_{ij}(\Psi). \quad (6.4)$$

Thus, the problem of minimizing the wireless communication energy, by means of antennas orientation optimization, for the application specific scenario can be formulated as finding the antennas orientation vector Ψ which minimizes $E_{tot}^{(as)}(\Psi)$.

General Purpose Scenario

In the *general purpose scenario* it is assumed that communication traffic information are not available at design time. Based on this, the same traffic volume for each communicating pair is assumed. The total wireless commu-

nication energy per bit can be computed as:

$$E_{tot}^{(gp)}(\Psi) = \sum_{i=0}^N \sum_{j=0}^N E_{ij}(\Psi). \quad (6.5)$$

Thus, the problem of minimizing the wireless communication energy, by means of antennas orientation optimization, for the general purpose scenario can be formulated as finding the antennas orientation vector Ψ which minimizes $E_{tot}^{(gp)}(\Psi)$.

Worst Case Scenario

Please notice that, in the above two scenarios (application specific and general purpose), it is assumed that the transceiver implements a transmitting power calibration technique (see the previous chapter or the technique presented in [60]) which allows the transmitting radio hub to use the minimum transmitting power to reach the destination guaranteeing a certain bit error rate.

In WiNoCs in which the power amplifier (PA) in the transceivers does not implement any transmitting power modulation mechanism, the PA is configured to use the maximum transmitting power irrespective of the recipient of the transmission. We refer to this case as *worst case* scenario. In this case, the total wireless energy consumption per bit is determined by the maximum $E_{ij}(\Psi)$:

$$E_{tot}^{(wc)}(\Psi) = \max_{i,j=1,\dots,N} E_{ij}(\Psi). \quad (6.6)$$

Thus, the problem of minimizing the wireless communication energy, by means of antennas orientation optimization, for the worst case scenario can be formulated as finding the antennas orientation vector Ψ which minimizes $E_{tot}^{(wc)}(\Psi)$.

6.1.3 General Design Flow

The Friis equation used in the problem formulation is not suitable for computing the actual attenuation of the wireless medium. In fact, it models only first order effects which are however enough for early design space exploration. Thus, the use of accurate field solver simulators or direct measurements on real prototypes are needed for implement the overall optimization flow. The basic steps which form the design flow can be summarized as follows.

1. Compute the radiation pattern of the antenna by means of a field solver simulator or by test-chip measurements. The radiation pattern represents the term D in Eqn. (6.2).
2. Explore the antennas orientations design space for determining the optimal antennas orientations which minimize the total communication energy (*cf.*, subsection 6.1.2).
3. Configure the antennas with the orientations found in the previous step and compute the actual attenuation by means of a field solver simulators for determining the transmitting power for each antennas pair. Let $P_{t_{ij}}$ be the transmitting power for communication from antenna i to antenna j .
4. For the general purpose and application specific scenarios, use the $P_{t_{ij}}$ computed in the previous step for configuring the variable gain amplifier controller [59, 60]. For the worst case scenario, set the transmitting power of every transmitter to $\max P_{t_{ij}}$.

In the next section, such design flow is applied for designing energy efficient WiNoC configurations for the different scenarios, under different traffic patterns and different parameters.

6.2 Experimental Results

In this section we explore the design space spanned by the orientations of the antennas in a WiNoC architecture with the goal of improving its energy efficiency.

6.2.1 Simulation Methodology

Given the non-linear, high dimensional, multi-modal, and non-smooth nature of $E_{ij}(\Psi)$, the optimization problems [Eqns. (6.7)–(6.9)], defined in the

Table 6.1: HFSS setup parameters.

Parameter	Value
Chip Size	20 mm × 20mm
Technology	28 nm SOI
Silicon Resistivity	$\rho = 5 \text{ K}\Omega\text{cm}$
Substrate Thickness	350 μm
Oxide (SiO_2) Thickness	30 μm
Antenna Elevation	2 μm
Antenna Thickness	2 μm
Antenna Axial Length	$2 \times 340 \mu\text{m}$
Operation frequency	60 GHz
Absolute Bandwidth	16 GHz

previous section, have been solved by means of simulated annealing.

$$\min_{\Psi} \sum_{i=0}^N \sum_{j=0}^N V_{ij} \times E_{ij}(\Psi) \quad (6.7)$$

$$\min_{\Psi} \sum_{i=0}^N \sum_{j=0}^N E_{ij}(\Psi) \quad (6.8)$$

$$\min_{\Psi} \max_{i,j=1,\dots,N} E_{ij}(\Psi) \quad (6.9)$$

For each of the above scenarios, namely, application specific [AS, Eqn. (6.7)], general purpose [GP, Eqn. (6.8)], and worst case [WC, Eqn. (6.9)], the optimal set of antennas orientation, namely, $\Psi_{opt}^{(AS)}$, $\Psi_{opt}^{(GP)}$, and $\Psi_{opt}^{(WC)}$, are simulated by means of an accurate field solver simulator for obtaining the scattering parameters. Then, scattering parameters are used by Eqn. (5.2) for computing the transmitting power for each transmit-receive antenna pair. Such transmitting power data are then used for back-annotating a cycle accurate WiNoC simulator [36] for determining the total energy figures under different traffic scenarios.

In all the experiments, as done for the technique described in the previous chapter, we consider a zigzag antenna modelled and characterized with Ansoft HFSS [1] (High Frequency Structural Simulator). For comodity, Tab. 6.1 resumes the simulation parameters used in all the experiments.

Fig. 6.3 shows the antenna directivity, by means of its radiation pattern, considering the direction of maximum radiation under the substrate ($\phi = 100^\circ$).




Figure 6.3: Radiation pattern for a zigzag antenna at the elevation of maximum radiation ($\phi = 100^\circ$). $\theta = 0^\circ$ is the direction orthogonal to the antenna's main axis. According to Fig. 5.2, we assume the antenna situated upon the XY plane (coplanar with silicon die).

6.2.2 Energy Saving Analysis

Let us know analyze the energy savings in the application specific (AS), general purpose (GP), and worst case (WC) scenarios. As communication traffic patters, we used a set of representative applications of SPLASH-2 and PARSEC benchmarks suites. Such benchmarks have been executed on Graphite Multi-core Simulator [58] and the communication topology graphs and communication volumes information have been extracted. In all the experiments, the baseline WiNoC architecture is msWiNoC [28] in which all the antennas have the same orientation. In the case of AS and GP, it is assumed that the power amplifier in the transceivers is equipped with the reconfigurable variable gain amplifier (R-VGA) module [59] with seven power steps. The estimated transmitting power ranges from $8 \mu\text{W}$ (-21 dBm) to $794 \mu\text{W}$ (-1 dBm), that in terms of energy per bit correspond to 0.42 pJ/bit and 1.4 pJ/bit, respectively. Based on this, we have selected seven equally spaced power steps into such range. That is, the i -th power step corresponds to a transmitting power of $8 + (i - 1) * 786/6 \mu\text{W}$.

Fig. 6.4 shows the percentage energy saving when the antennas are optimally oriented based on the solutions of optimization problems [Eqns. (6.7)–

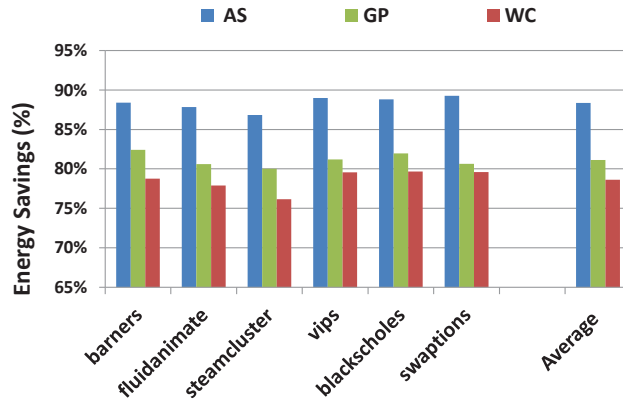


Figure 6.4: Energy saving obtained for ms-WiNoC with 256 nodes under different traffic scenarios.




Figure 6.5: Energy saving for different number of radio hubs.

(6.9)] considering four antenna orientations steps. On average, up to 89%, 82%, and 78% energy saving is observed for AS, GP, and WC scenario, respectively.

Fig. 6.5 shows the percentage energy saving when the number of radio hubs is made to vary. As expected, the energy saving increases as the number of radio hubs increases due to the fact that more communications make use of the radio medium. However, no relevant improvement is observed when the number of radio hub is greater than eight. Such trend is related to the network size that in our experiments consists of 256 communicating cores. In fact, for such a medium network size, eight radio hubs are enough

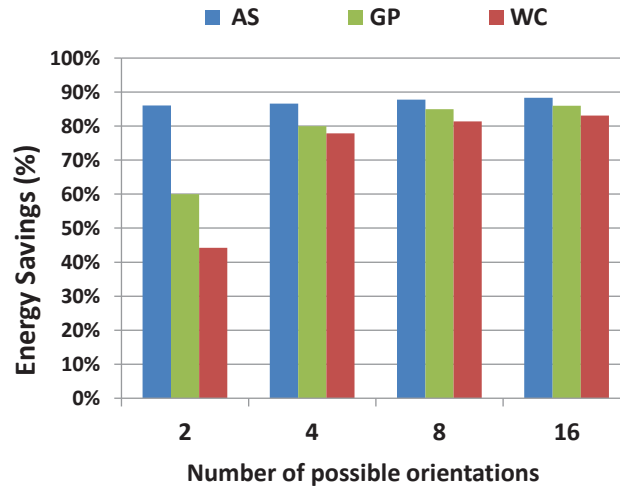


Figure 6.6: Energy saving for various number of possible orientations of the antenna (WiNoC configured with 16 radio hubs).

for drastically reducing the average hop count. Above eight radio hubs, the short distances between them, makes more suitable performing the communication by means of the wired underlying NoC. Please remind that, wireless transmissions becomes effective in term of energy efficiency when the path length is greater than three hops [25].

We analyzed four cases in which 2, 4, 8, and 16 orientations are allowed. Such allowed orientations are those obtained by equally dividing the orientations from 0° to 180° into 2, 4, 8, and 16 angles, respectively. Fig. 6.6 shows the percentage energy savings in such cases. It is interesting to observe that AS is quite insensitive to the increase of the number of admissible antenna orientations. This behavior is explained by the fact that AS directs the antenna along the direction with the maximum traffic volume. For this reason, having more than two available orientations, does not affect the solution found for AS. On contrary, GP and WC are strongly sensitive to the number of available antenna orientations. As it can be observed, passing from 2 to 4 possible orientations, it results in an energy saving gap of about 20% and 45% for GP and WC, respectively. For instance, in the case of GP, the optimal orientation of the antennas is determined by assuming that the traffic volume between all the radio hub pairs is the same. Thus, for a generic antenna, a trade-off orientation is determined in such a way to satisfy all the directions.




Figure 6.7: Optimal antennas orientations for the application specific and worst case scenarios.

6.2.3 Case Study

As a real case study, we consider the same complex heterogeneous platform as shown in the previous chapter (Fig. 5.15, page 86). Such application was mapped in a HmWNoC [28] partitioned in 16 subnetworks where the upper level network is a 2D mesh topology augmented with three radio hubs. The number of radio hubs and their placement into the network has been derived by using the optimization procedure described in [28].

We explore the antennas orientation design space for both the application specific (AS) and worst case (WC) scenarios. We consider antennas can be oriented among four angles (0° , 45° , 90° , 135°). We assume a tunable power amplifier supporting seven power steps with a transmitting energy per bit ranging from 0.42 pJ/bit to 1.4 pJ/bit. Fig. 6.7 shows the optimal orientation of the three antennas for the two scenarios. With regard to the AS scenario, due to the presence of memory elements close to C14, which represents an hot-spot region of the network, there is a relevant traffic volume between such region and both C0 and C7. Based on this, as can be observed from Fig. 6.7 (AS), both antennas in C7 and C0 are oriented in such a way to reduce energy when communicate with C14. With regard to the WC scenario, since no traffic information is used during the design space exploration, the optimal orientation of the antennas found, tries to minimize the worst case condition which is represented by the communication between the two far apart clusters, namely, C14 and C0. In fact, as it can be noticed, the antenna

Table 6.2: Transmitting energy per bit for each transmitting and receiving antennas pair.

TX	RX	Transm. energy (pJ/bit)		
		BS	AS	WC
C14	C7	1.40	0.58	1.07
C14	C0	1.40	1.23	1.07
C7	C14	1.40	0.58	1.07
C7	C0	1.40	0.91	1.07
C0	C14	1.40	1.23	1.07
C0	C7	1.40	0.91	1.07
Total energy (J)		1.68×10^{-4}	1.23×10^{-5}	1.02×10^{-4}

in C14 fits the directivity of the antenna in C0 and vice-versa.

Tab. 6.2 reports, for each transmitting and receiving antennas pair, and for each considered scenario, the transmitting energy per bit. The table also shows a baseline scenario (BS) in which all the antennas have the same orientation (0°). Of course, for both the WC and BS scenarios, the transmitting energy is constant irrespective of the location of the transmitting and receiving antenna. The optimization of the antennas orientation in the WC scenario allows to reduce the communication energy by 39% as respect to the BS scenario. By considering the AS scenario, in which the transmitting power is tuned online, the communication energy reduces by 51%.

6.3 Conclusions

Several work in literature in the context of low power WiNoC architectures, do not take into account the impact of the antenna directivity on power metrics. In addition, they assume that antennas present an omni-directional radiation pattern which is far away from the reality (see, for instance, Fig. 6.3). In this chapter we have highlighted the need for antennas orientation design space exploration for improving the energy figures of WiNoC architectures. We have considered three main scenarios, namely, application specific (AS), general purpose (GP), and worst case (WC). In the AS scenario, communication information are exploited for optimizing the orientation of the antennas in such a way to maximize the overlap of the radiation patterns of the antennas which communicate more. The GP scenario, is derived from AS by assuming that all the radio hubs communicate each other with the same

probability. Finally, the WC scenario is when the transceiver does not implement any transmitting power on-line calibration scheme and, therefore, all the radio hubs communicate using the same transmitting power irrespective of their location in the chip. A state of the art small-world based WiNoC architecture has been used as reference WiNoC architecture in the experiments. The exploration of the antennas orientation design space under different traffic patterns resulted in important energy savings in all the three scenarios considered.

Chapter 7

Smart Transceiver for Wireless Network-on-Chip Architectures

In the previous two sections we proposed a solution to reduce power only in a part of the transceiver; the transmitter. Reducing power on the receiver and other radio-hub components remain instead an open issue. In [63, 61], a mechanism which selectively switches off those receivers not involved in any communication by means of a centralized controller is proposed. Unfortunately, the main inconvenient of such a mechanism is its centralized nature. Since, usually, WiNoC lies on a large SoC, signal traveling from/to the central controller must be buffered. This leads to an increase in terms of delay, wire congestion, and power consumption due to nets driven by large buffers. For these reasons, in [40, 64] large buffered have been replaced with *glines* [96]. Unfortunately, using *glines* results in a custom design approach, with a non conventional CMOS design flow.

Based on the above considerations, In the proposed research, we developed SiESTA (Smart Energy Saving Transceiver), a power managing technique tailored for WiNoC architectures and aimed at improving the energy efficiency of radio-hubs. SiESTA acts on different power-hungry elements of the radio-hub, including, input/output buffers from/to the attached tiles, antenna buffers, and transceivers by attacking both their static and dynamic power contributions. SiESTA exploits the fact that the radio channel is a shared resource that can be used by a single radio-hub per transmission. Based on this, the radio-hubs that are not recipients of the current wireless communication can be switched off for a number of clock cycles that can be

simply computed by processing the header flit of the current packet under the assumption that wormhole switching technique is used. During such *sleeping* period, and based on the status of buffers and reservation tables, both clock gating and power gating actions can be taken for the sake of energy efficiency improvement.

The proposed power management technique has been assessed on several WiNoC architectures featuring different numbers of radio-hubs and under different traffic scenarios (both synthetic and generated by the execution of real applications) and by varying several parameters including packet size, flit size, buffer size, and packet injection rate. The design space spanned by the aforementioned parameters has been then analyzed in terms of energy saving as compared to the correspondent WiNoCs that are not equipped with SiESTA. We found that, up to 80% of wireless communication energy saving and up to 25% of total communication energy saving can be obtained when SiESTA is used, without any performance degradation and with a negligible impact on the silicon area of the radio-hub.

7.1 SiESTA Power Reduction Strategy

7.1.1 Radio-Hub Architecture Overview

The SiESTA power management approach proposed in this work is based on switching on/off actions dynamically applied to different components of the WiNoC transceiver. Such actions are taken by the SiESTA control unit as a consequence of both the status of these components and the transmission events occurring in the network. Prior to describe the internals of the SiESTA approach, it is essential to get a picture of the main elements constituting the radio-hub wireless communication architecture and the involved data flows.

As explained in the Sec. 4.3.3, the radio-hub component plays a fundamental role in a WiNoC architecture, allowing single hop communication between tiles that are not adjacent. Each radio-hub has a set of NoC tiles locally connected via wired communication. In general, a radio-hub can support one or more wireless channels, but only radio-hubs with a common channel can exchange data, in accordance with [104]. In the following, we will assume that a simple packet-based token passing mechanism [30, 73] is used to manage access control when trying to transmit along a channel.




Figure 7.1: Wireless communication between two tiles using radio-hubs. Four classes of radio-hub buffers are involved: (A) $antenna_buffer_TX_{Ch_x}$, (B) $antenna_buffer_RX_{Ch_x}$, (C) $buffer_from_tile_i$, and (D) $buffer_to_tile_i$

There are two fundamental types of communications handled by a radio-hub component: wired Tile-Hub connections between radio-hub and tiles, and wireless Hub-Hub connections between different radio-hubs. Fig. 7.1 shows the four classes of buffers involved in a simple scenario with two tiles t_1 and t_2 , connected to hub_a and hub_b , respectively. Further, we assume that hub_a and hub_b support transmissions on a shared channel Ch_0 . In wired Tile-Hub communications, two kind of radio-hub buffers are involved, depending on the direction of the data flow: The input buffer (C) of the radio-hub, that receives data from a tile i , referred to as $buffer_from_tile_i$ and, on the opposite direction, the output buffer (D) of the radio-hub, that stores data to be sent to tile i , referred to as $buffer_to_tile_i$. Discussions about the low-level mechanisms involved in the communication protocol are beyond the scope of this research. [16] for details). It is sufficient to know that these wired Tile-Hub connections use identical communication phases as in wired Tile-Tile connections. With regard to wireless Hub-Hub communications, we will use the notation $antenna_buffer_RX_{Ch_x}$ and $antenna_buffer_TX_{Ch_x}$ to denote the buffers used for receiving and transmitting data, respectively, using the radio channel Ch_x (see buffers A and B in Fig. 7.1). Of course, usually more than

a single tile node is connected to each radio-hub and also several different radio-hubs can share a common channel, however, we will refer to the trivial example in Fig. 7.1 for the sake of clarity.

Let us now suppose that tile t_1 , connected to hub_a , wants to send data to tile t_2 , connected to hub_b , as shown by the dotted path in Fig. 7.1. Thus, a header flit comes from path along the t_1 wired connection and it is stored in the *buffer_from_tile₁*, that is, the input buffer associated to the radio-hub port connecting t_1 . When the hub_a analyzes the buffer, it finds the header flit and takes a routing decision. Please notice that for the purposes of this work we are not interested in analyzing which particular routing algorithm is used in the radio-hub. Without loss of generality, we can assume a simple “use-wireless-when-possible” policy, e.g., since the destination t_2 is connected to a different radio-hub hub_b , a wireless transmission will be issued to reach t_2 .

Then hub_a tries to reserve the output direction associated to the channel Ch_0 . If such direction is available, hub_a reserves the output inserting the header flit on the output *antenna_buffer_TX_{Ch₀}*: the buffer physically connected to the transmitting antenna associated to the channel Ch_0 . Then, a new transmission is started if the radio-hub is not yet currently busy and has the ownership of the medium access. Please notice that different policies can be used to control the access to the channel. In this work we will assume a *token-packet based* medium access control in which a token is shared between radio-hubs supporting a common channel. To enforce wormhole switching, the token grants the ownership of the channel for all the clock cycles required to complete a packet transmission. Radio-hubs having nothing to transmit simply forward the token to the next radio-hub following some logical token-ring order. Finally, after a certain number of delay cycles, which depend on the antenna data rate and flit size, the flit is received by the destination antenna and put into the *antenna_buffer_RX_{Ch₀}* buffer of hub_1 . As we will see in the next subsection, this delay plays a crucial role in the whole power management mechanism presented in this work.

7.1.2 Radio-hub data flows

Recalling the reference architecture depicted in Fig. 7.1, we can see how radio-hub elements in the WiNoC consist of hardware components (buffers and logic) that can be functionally associated with four main types of data

flow: (1) Wireless data transmission starting from the current radio-hub (2) Wireless data reception coming from another radio-hub (3) Flit dispatching from the current radio-hub towards a tile via wired Tile- Hub connection (4) Flit reception from a tile via wired Tile-Hub connection. The main idea behind the SiESTA power management scheme is to dynamically switch off the components involved in these four types of data flows exploiting information about the status of the same components and the wireless transmission events occurring the WiNoC.

Going further into details, let us consider the Fig. 7.2, which shows the radio-hub transceiver internals for both TX and RX data flows. Next to the analog part of the transceiver we can find the antenna buffers already described in the previous subsection. Buffers to/from tile are also present in rightmost part of the picture. Please notice that some architectural components, required in multiple instances, are represented using multiple shapes. For example, a radio-hub supporting more than a single channel would require multiple instances of antennas, modulation/demodulation hardware and other logic. In particular, a crossbar mechanism for both RX/TX data flow is required to connect antenna buffers to the buffer from/to tile, since multiple instances of antenna buffers (one for each channel) must be connected to multiple instances of buffers from/to tile (one for each connected tile). Further, two reservation tables (*antenna_to_tile_RT*, *antenna_from_tile_RT*) are required, because, once a flit header has reserved a channel, the same channel should be considered busy until a tail flit releases it. The elements controlled by the SiESTA power management have been highlighted with dashed boxes. Further, when a set of components can be logically considered as a single unit from the SiESTA power management perspective, a label has been put outside the dashed box.

7.1.3 Detecting the *Radio Event Sleep* status

SiESTA control logic implements a simple mechanism to detect conditions for performing safe power down decisions on different elements. Tab. 7.1 summarizes the conditions that the control logic of SiESTA checks on each radio-hub. At each cycle, the behavior of the power management is dictated by the evaluation of these conditions, meaning that no power reduction strategy is being applied to the components whose condition is not matched. A

fundamental role is played by a particular status that can be enabled by SiESTA for a given radio-hub and transmission channel, namely *Radio Event Sleep (RES)*. The idea behind *RES* is to exploit information gathered from channel events in order to determine a condition that can be used to safely switch off some architectural components of the radio-hub. For example, let us assume that a new transmission is started by a radio-hub hub_j on channel Ch_x . From the perspective of a different radio-hub hub_i , only two different situations can happen: (i) the radio-hub hub_i is the recipient, (ii) the hub_i is not the recipient. The SiESTA control logic on hub_i detects situations belonging to the case (ii), enabling the *RES* status of radio-hub hub_i for channel Ch_x . The meaning of such status is that hub_i will not be the target of any transmission on channel Ch_x for a certain amount of cycles. At the same time, from a TX perspective, this means also that hub_i will not be able to start any transmission on the channel for the same amount of cycles. Please notice that precise and deterministic number of cycles (N_{RES}) for holding the *RES* status can be computed, since we are assuming a packet transmission token policy in which the token is released only after the whole packet has been transmitted. In particular, the number of safe N_{RES} cycles can be computed as:

$$N_{RES} = T_{delay} \times packet\ size \times clock\ frequency \quad (7.1)$$

where *packet size* is expressed in number of flits and T_{delay} is the amount of time required to transmit a flit, computed as:

$$T_{delay} = \frac{flit\ size}{DR} \quad (7.2)$$

where DR is the data rate of the antenna in bit/s and *flit size* is the flit size in bits. In other words, the meaning of the above computation is that is not feasible to complete the packet transmission in less than N_{RES} cycles, assuming that antenna data rate, flit size, packet length and clock frequency. For example, assuming a clock frequency of 1 GHz, a flit size of 32 bits and a 16 Gbs mm-Wave antenna using On-Off keying (OOK) modulation, T_{delay} is 2 ns, corresponding to 2 cycles N_{RES} . This is a conservative assumption since we do not know *when* the transmission will actually end. For example, heavy traffic loads and congestion events could introduce further delay cycles.

Finally, it is important to point out that enabling the *RES* status on hub_i for a given channel does not necessarily mean that all hardware blocks related

to wireless data transceiver of that channel can be switched off. In fact, as explained in the following subsections, the proposed approach implements a selective strategy checking (cycle-by-cycle) which components can be safely disabled while the *RES* status is enabled.

7.1.4 TX Power Management

Let us first describe how SiESTA power management works by considering the TX data flow of the radio-hub, depicted in the upper part of Fig. 7.2. The components involved, highlighted with dot boxes, are *Analog TX*, *Serializer*, and *antenna_buffer_TX*. It is important to point out that the decision about switching off these components cannot be only based on the presence/absence of transmissions on a given channel. In fact, the absence of transmissions is only a necessary condition, not sufficient to perform a safe switch off decision. Starting from the right, the first component interested by the power management is the *antenna_buffer_TX* (for a given channel Ch_x). It is not sufficient to check whether the buffer is currently being used or not, since new data flits could arrive from any *buffer_from_tile*. Looking at the Ch_x reservation entry in the *antenna_from_tile_RT*, in fact, we can be sure that no flit will be sent (in the next cycle) to the *antenna_buffer_TX* of channel Ch_x . This condition is formally described in the first row of Tab. 7.1. Continuing along with the TX data flow, we have the *Analog TX* and *Serializer* components. A condition sufficient for switching off is that the *RES* status is enabled, since we trivially know that no request for transmissions will be issued for a given amount of cycles. More interesting is to analyze how to take decisions while *RES* status is *not* enabled, since we cannot safely ignore wireless events in the next cycles. However, a further condition can be checked to be sure of not requiring *Analog TX* and *Serializer* components, that is, checking that no transmission is going to be scheduled in the next cycle. This can be easily accomplished by exploiting the same condition already checked for *antenna_buffer_TX*. It is a conservative approach, but it safely works when the *RES* status is not enabled. Further, this is not a so uncommon situation, since it occurs everytime no transmission is queued in the channel. The above conditions are formally summarized in the second row of Tab. 7.1. A consideration should be made about *buffer_from_tile* buffers. While their status is an input for the SiESTA power management logic, they

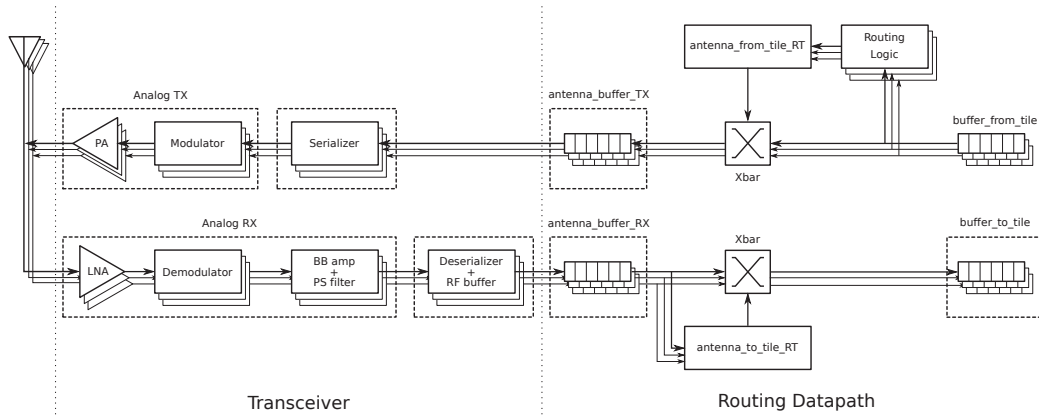


Figure 7.2: Internal modules of the radio-hub architecture: elements highlighted with dashed boxes are those subjected to SiESTA power management actions.

cannot be subjected to any powering off decision. Recalling Fig. 7.1, each of these buffers is associated to a specific port physically connected to a given radio-hub. But the events determining flit arrivals on these buffers are external to the radio-hub, i.e. outside the control of the radio-hub hardware, so none of these buffers can be safely switched off a priori.

7.1.5 RX Power Management

Let us describe now how the *RES* status for a given channel is used for the RX part of the SiESTA power management conditions seen in Tab. 7.1. Starting from the antenna on the left side of Fig. 7.2, the first components affected by SiESTA are those related to *Analog RX*, *Deserializer* and *RF buffer*. These components can be safely disabled while the *RES* is enabled, since no events on the channel can involve them. Note how, even if subjected to the same off condition, in Fig. 7.2 we explicitly distinguished the analog part (signal snooping, carrier sensing, demodulation, filtering etc...) from the digital deserialization/buffering one. Continuing along with the RX data flow, the next component is the *antenna_buffer_RX*. In this case, no decision can be taken using only the *RES* status: even if no transmission event will involve the radio-hub in the next cycles, the *antenna_buffer_RX* cannot be disabled whenever flits are present in the buffer. This mandatory requirement is formalized in the first part of the AND logic expression in the related row

Module	Off condition
<i>antenna_buffer_tx_{Ch_x}</i> Serializer, Analog TX	<i>antenna_buffer_tx_{Ch_x} is empty</i> AND <i>antenna_from_tile_RT[Ch_x] is not reserved</i>
Deserializer + RF buffer, Analog RX	<i>RES status is enabled</i> OR (<i>antenna_buffer_tx is empty</i> AND <i>antenna_from_tile_RT_i is not reserved</i>)
<i>antenna_buffer_rx</i> <i>buffer_to_tile_i</i>	<i>antenna_buffer_rx is empty</i> AND (<i>RES status is enabled</i> OR <i>Analog RX is inactive</i>)
	<i>buffer_to_tile_i is empty</i> AND <i>antenna_to_tile_RT_i is not reserved</i>

Table 7.1: Conditions evaluated by the control logic of SiESTA to determine when a specific module of the radio-hub can be switched-off.

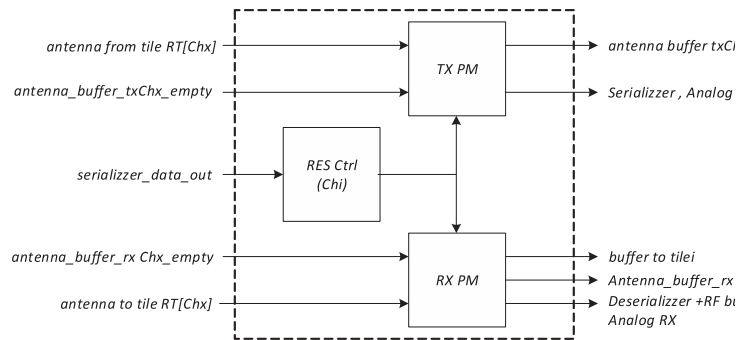


Figure 7.3: Control logic of the SiESTA power manager.

of Tab. 7.1. The second part of the condition formalizes the other two sub-conditions that can enable the switch-off condition: the first is being in a *RES* status, the second is having the *Analog RX* marked as *inactive*, that is, not processing any data from the antenna. With regard to *inactive*, as for the other *empty*, *not reserved*, *enabled* variables, it is the control logic of SiESTA that is responsible for collecting the necessary signals and determine their values. Continuing in Fig.7.2, the last components interested by SiESTA are the *buffer_to_tile* buffers. Each single *buffer_to_tile₁* can be selectively disabled if it is empty and not reserved by any wireless incoming transmission. As can be seen in Tab.7.1, this information can be easily gathered from the *antenna_to_tile_RT*, reading the appropriate entry on the reservation table. Please notice the duality between this last condition and the one introduced in the first row of Tab.7.1: in both cases the subject is a buffer connected to a crossbar, but two different reservation tables are taken into account depending on the direction of the data flow involved.

7.1.6 Hardware Implementation

In this subsection SiESTA is discussed from its hardware implementation viewpoint. The module implementing SiESTA is responsible for the genera-

tion of control signals driving the MOS transistors of the power-gating logic for each of the building blocks in the dashed rectangular frames in Fig. 7.2. The implementation of the power gating mechanism into the different components of the transceiver is out of the scope of our work. Interested readers can refer to [63, 61].

The three main building blocks which form SiESTA, namely, TX PM, RX PM, and RES Ctrl, are shown in Fig. 7.3. While TX PM and RX PM are responsible for generating the control signals for activating or deactivating the different components of the transmitter and receivers, respectively, RES Ctrl implements the logic for signaling the *RES* status (*cf.*, Sec. 7.1.3) which is used as input for both TX PM and RX PM. Basically, RES Ctrl is a counter which is triggered when the current radio-hub is not the recipient of the current wireless transmission. Thus, RES Ctrl is connected to the deserializer output in order to decode the destination address stored in the first bits of the incoming head flit.

TX PM gets as inputs, other than the *RES* status, the information coming from the reservation table for the specific channel ($antenna_from_tile_RT[Ch_x]$) and the empty signal associated to $antenna_buffer_tx_{Ch_x}$. Its outputs drive the power gating logic of the $antenna_buffer_tx_{Ch_x}$, the Serializer, and the Analog TX front-end. RX PM gets the same input of TX PM but related for the complementary modules: specifically, the information coming from the reservation table for the specific channel ($antenna_to_tile_RT[Ch_x]$) and the empty signal associated to $antenna_buffer_rx_{Ch_x}$. Its outputs drive the power gating logic of the $buffer_to_tile$, $antenna_buffer_rx$, and the receiver radio frequency front-end Deserializer + RF buffer and Analog RX.

It should be pointed out that, the logic for implementing SiESTA is very simple thus resulting in poor area utilization and very low power dissipation. In fact, looking again at its input/output behaviour summarized in Tab. 7.1, it can be implemented by a simple combinational structure (excluding the counter into the RES Ctrl). Further, the design is scalable, since the number of bits used by the counter into the RES Ctrl only depends on the maximum packet length. A gate level synthesis (65 nm technology) of SiESTA, designed for a 4 inputs radio-hub and 32 flits packet size, results in an area occupation as low as $102 \mu\text{m}^2$ and a power dissipation of 0.08 mW. Such area and power overheads constitute a negligible 0.04% and 0.21% of




Figure 7.4: Area (a) and Power (b) breakdown of the radio-hub implementing SiESTA power manager.

the total radio radio-hub area and power, respectively. With this regard, Fig. 7.4 shows a breakdown of the major contributions, confirming how the transceiver constitutes the most impacting component on the power/area of the radio-hub architecture seen in Fig. 7.2. Another important contribution is due the Routing Datapath, which comprises internal antenna buffers, buffers to/from tile, reservation table and the crossbars required to dispatch flits in both directions. In this implementation (a radio-hub with 4 ports), such contribution still remains below 25%. However, it should be pointed out that, an higher number of ports could result in bigger Routing Datapath contribution, due to the increased crossbar radix and number of buffers.

7.2 Experiments

In this section SiESTA is assessed on several WiNoC configurations and under different traffic scenarios. If not otherwise specified, energy saving figures refer to the total energy saving when SiESTA is used as compared to a baseline WiNoC with the same configuration and under the same input traffic.

7.2.1 Simulation Setup

The experiments have been carried out on an extended version of Noxim [16] augmented with the logic implementing SiESTA. A 256-node network has




Figure 7.5: RH-4,8,12, and 16 architectures considered for the simulations: grey tiles indicated how the radio-hubs have been mapped into network.

been considered and it has been configured into four different WiNoC architectures with 4, 8, 12, and 16 radio-hubs, respectively, as shown in Fig. 7.5. Adaptive Layered Shortest Path Routing (ALASH) [54] is used as the routing algorithm, allowing for messages to be routed along the shortest path between the source and destination while maintaining deadlock freedom.

Two main traffic scenarios, namely, uniform and hot-spot, have been used through the experiments. In uniform traffic, all nodes receive packets with the same probability. In hot-spot traffic, some nodes are marked as hot-spot nodes. The hot-spot nodes receive more traffic than the non hot-spot nodes in the network. Specifically, hot-spot traffic is defined along with an hot-spot probability, p . The destination of a packet is a specific hot-spot node with probability p and a non hot-spot node with the remaining probability $1 - np$, where n is the number of hot-spot nodes. In the experiments, when hot-spot traffic is considered, hot-spot nodes are those directly attached to the radio-hubs (*i.e.*, those with gray background in Fig. 7.5). The considered traffic scenarios allow the network to work in two well distinct zones characterized by a different utilization of the wireless medium. Specifically, we define the *wireless utilization* as the ratio between the number of communications




Figure 7.6: Percentage wireless utilization for different WiNoC architectures under uniform and hot-spot traffic scenarios.

which use, totally or in part, the wireless medium and the total number of communications. Fig. 7.6 shows the wireless utilization in the four considered network configurations under uniform and hot-spot traffic scenarios. For the hot-spot traffic we considered a hot-spot probability of 5% irrespective of the WiNoC configuration. As expected, the utilization of the wireless medium is more pronounced under hot-spot traffic. This is due to the fact that, since the hot-spot nodes correspond to those directly attached to the radio-hubs, the routing algorithm is biased toward selecting a wireless path that allows to reach the communication recipient in a single hop. Further, the wireless utilization increases as the number of radio-hubs increases. This is due to the fact that, as the number of radio-hubs increases, the probability for a packet of encountering a radio-hub along its way increases resulting in the increase of chances of using the wireless shortcut for reaching a far destination.

7.2.2 Effect of Packet Size and Packet Injection Rate

Packet injection rate (pir) and packet size are two parameters that strongly affect the energy and performance metrics of a NoC. In this subsection, we study their impact on energy saving when SiESTA is used. Specifically, we modulate the pir and the packet size in three different ways as follows:

- *fixed workload*: pir and packet size have been made to vary by keeping

the workload constant. That is, as the packet size increases, the pir is proportionally reduced by the same factor in such a way that the total traffic volume in the network remains constant. The aim of this experiment is to demonstrate the effectiveness and stability of SiESTA when different choices are made at the application level to encode the same amount of data.

- *fixed packet size with variable pir*: packet size is kept fixed while pir is made to vary. The aim of this experiment is to assess the effectiveness of SiESTA when the network load increases filtering the effect of the packet size which affects the duration of the sleep periods, as seen when computing N_{RES} cycles in Equation 7.2.
- *fixed pir variable with packet size*: pir is kept fixed while packet size is made to vary. In fact, the fraction of time in which some of the modules of the transceiver can be switched off depends on the packet size. The aim of this experiment is to assess the effect of packet size on the energy savings achieved by SiESTA.

Fig. 7.7 shows the energy saving for the *fixed workload* case, for both uniform and hot-spot traffic scenarios. As it can be observed, the energy saving increases with the number of radio-hubs. In fact, as the number of radio-hubs increases, the fraction of total energy due to the wireless infrastructure increases as well and, since SiESTA acts only on the wireless infrastructure, the energy saving increases as well. The energy saving for a given WiNoC architecture seems to be quite stable for the different pairs packet size/ pir . Indeed, using a fixed workload affects only the RES time, which depends upon the packet length (see Equation 7.2), but the overall effect is mitigated by the large number of wired communication components (*e.g.*, routers buffer), which are unaffected by the RES time of SiESTA. However, a different behaviour emerges when SiESTA is more stressed, as in the hot-spot traffic for RH-16, in which the energy savings for packets shorter than 12 flits are penalized by the smaller RES times.

Fig. 7.8 shows the energy saving for the *fixed packet size with variable pir* case, for both uniform and hot-spot traffic. Under uniform traffic, SiESTA is quite insensitive to the variation of pir . Conversely, in hot-spot traffic, the energy saving slightly decreases as pir increases with a 4.5% energy saving

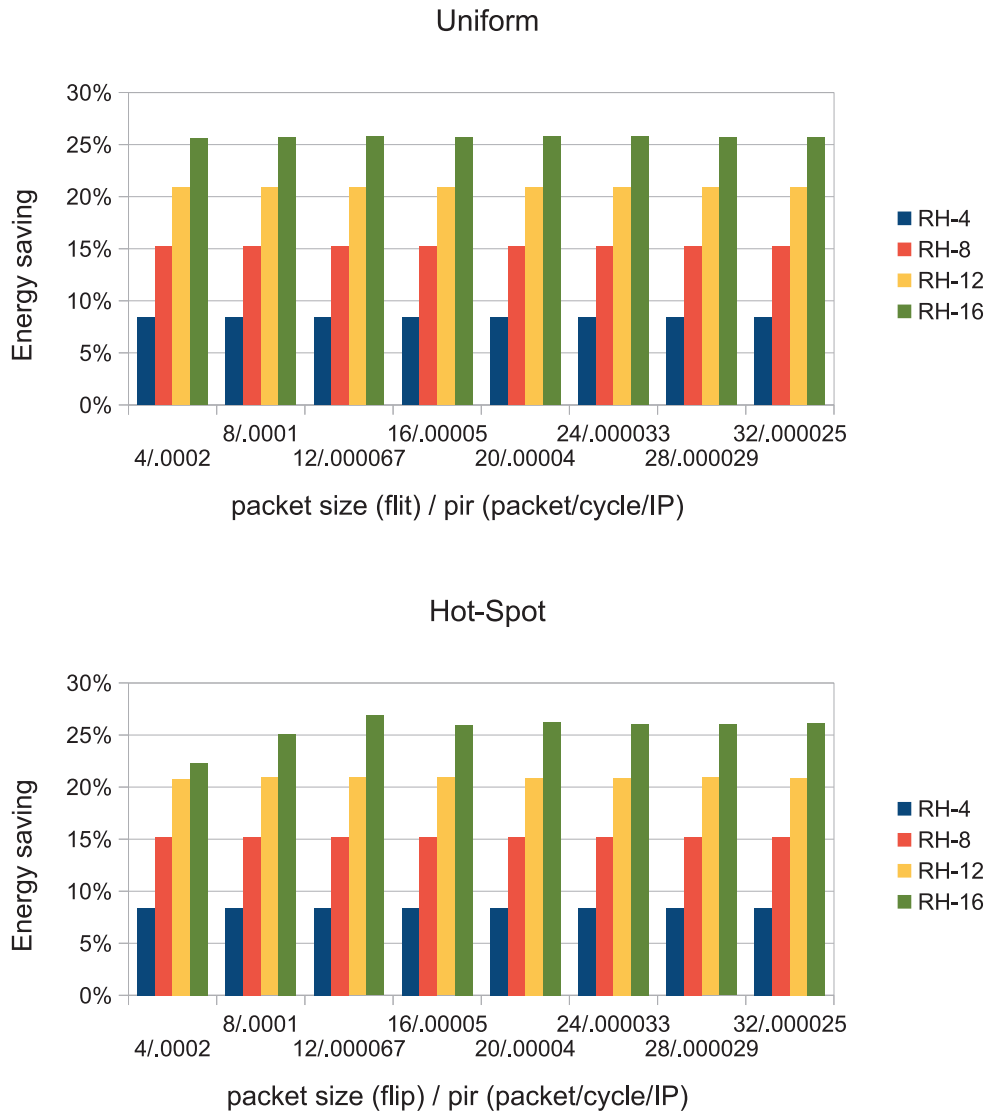


Figure 7.7: Energy saving for the *fixed workload* case under uniform and hot-spot traffic scenarios.




Figure 7.8: Energy saving for the *fixed packet with size variable pir* case under uniform and hot-spot traffic scenarios.

variation between the minimum and maximum pir . Indeed, given the nature of the hot-spot distribution, an increase of pir leads to an increment of wireless communications that is much higher than the one of non-wireless communications, thus determining an higher pressure on radio-hubs and less probability for SiESTA of satisfying off conditions reported in Tab. 7.1. Nevertheless, this decrease in SiESTA performance can be judged as acceptable, especially when considering that the pir has been made to vary across a 10x excursion in the value.

Finally, Fig. 7.9 shows the energy saving for the *fixed pir variable packet size* case, for both uniform and hot-spot traffic, confirming how SiESTA savings are stable when a fixed workload scenario is assumed, mainly affected by the number of radio-hubs of the network.

7.2.3 Effect of Flit Size

Like pir and packet size, the flit size has a relevant impact on performance and energy metrics. Its impact is more evident in WiNoC architectures as wireless communications are realized in a serial fashion and, then, their latency, as well as energy contribution, is affected by the flit size. Based on this, in this subsection we investigate on the effect of flit size on energy saving obtained by SiESTA.

Fig. 7.10 shows the energy saving obtained when SiESTA is used for the different WiNoC architectures under uniform and hot-spot traffic for different flit size. For both the traffic scenarios and for all the considered WiNoCs the general trend is a saddle function with its minimum on 32-bit flit size. To better understand this behavior it is useful to have a deeper sight of which are the actual contributions that realize the energy savings obtained with SiESTA.

In particular, Fig. 7.11 shows the power breakdown under hot-spot traffic for flit sizes of 16-, 32-, and 64-bit, where the left and right columns refer to the case without and with SiESTA, respectively. The point of inversion at flit size 32-bit can be explained by observing the power contribution of the antenna buffers (`antenna_buffer_pwr_s`) which has a minimum when the flit size is set to 32-bit. Thus, since the antenna buffers are among the major contributors to the SiESTA energy saving, their reduced impact negatively affects the effectiveness of SiESTA. The minimum value in correspondence




Figure 7.9: Energy saving for the *fixed pir variable packet size* case under uniform and hot-spot traffic scenarios.




Figure 7.10: Energy saving for the for different flit size under uniform and hot-spot traffic scenarios.




Figure 7.11: Power breakdown for different flit size under hot-spot traffic scenarios.

Table 7.2: Sensitivity analysis of buffers depth versus energy saving.

Buffer	Sensitivity degree (%)			
	RH-4	RH-8	RH-12	RH-16
router buffer	3.16	5.62	7.37	8.87
buffer to tile	1.22	2.23	3.06	3.70
buffer from tile	0.04	0.15	0.32	0.55
antenna buffer	4.99	8.69	11.49	12.94

of 32-bit flit size can be explained with the influence of two different and conflicting effects that the increase of flit size has on antenna buffer consumption: on one hand, bigger flit sizes make buffers bigger, thus leading to an overall increased antenna buffer power; on the other hand, the remaining router buffers (unaffected by SiESTA) also increase their contribution, thus limiting the fraction of overall power for which SiESTA can be effective. The 32 bit flit size is thus the point where the increased router buffer contribution is not counterbalanced by the antenna buffer power consumption increase.

7.2.4 Effect of Buffers Size

Looking again at Fig. 7.11 it can be noticed that buffers (*i.e.*, antenna buffers, buffers to/from tiles, and routers buffers) account for a significant fraction of the total energy consumption. In this subsection, we assess the effectiveness of SiESTA when buffers size is made to vary. Specifically, for router buffer, buffer from tiles, and buffer to tiles, we consider the buffer depth space $\{2, 4, 8\}$ flits and $\{16, 32, 64\}$ flits for antenna buffer (both rx and tx). Each of the 81 WiNoCs of the design space are then simulated with and without enabling SiESTA and the corresponding energy savings are collected. To quantitatively describe the impact of buffers' configuration on energy metrics, we have performed a sensitivity analysis [65] on buffers depth versus energy saving and Tab. 7.2 reports the correspondent sensitivity degrees. For each WiNoC architecture, antenna buffers and router buffers are those elements which mostly affect the energy efficiency of SiESTA. In particular, since SiESTA does not act on routers' buffers and since the sensitivity degree of antenna buffers increases as the number of radio-hubs increases, the effectiveness of SiESTA in improving the energy efficiency of a WiNoC increases with the number of radio-hubs.

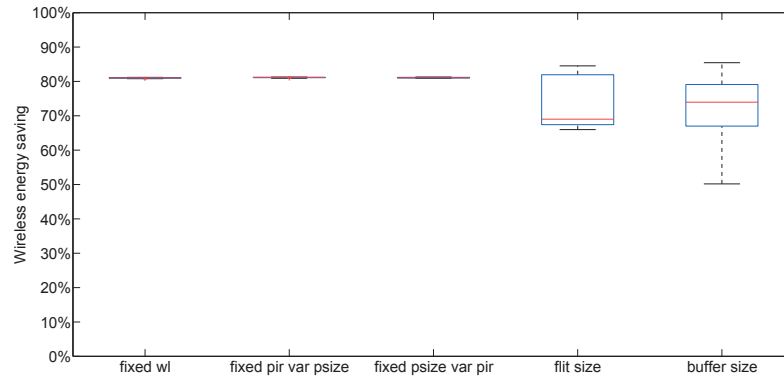


Figure 7.12: Boxplots representing the values distribution of wireless communication energy saving under uniform traffic.

7.2.5 Wireless Communication Energy Saving

In the previous subsections, we have assessed SiESTA by considering, as figure of merit, the total communication energy saving which includes both the contributions of wired communications and wireless communications. In this subsection, we want to focus on the energy contribution of wireless communication only. For each visited configuration of the design space considered in the previous subsections, Fig. 7.12 shows a boxplot of the wireless energy saving for the experimental scenarios described in the previous subsection under uniform traffic.

As can be observed, the effectiveness of SiESTA on radio-hub energy consumption is quite stable with an 80% energy saving for the fixed workload, fixed pir and fixed packet size scenarios. Larger excursions in values are observed in the two rightmost scenarios, where very power-impacting parameters such as flit size and buffer size are explored. However, these still represent good results, since the 75% of values, corresponding to the internal rectangle of the boxplot, are placed in the 70-80% region.

Overall, we can assess that SiESTA power management actions are very effective when considering the specific radio-hub energy consumption; however, as seen in the previous subsections, how these savings translates into actual global savings strictly depends upon the fraction that wireless energy consumption represents with respect to the total amount of energy of the whole network.




Figure 7.13: Wireless utilization for different WiNoC architectures and for different benchmarks.

7.2.6 Assessment under Real Traffic Scenarios

In this subsection, we assess SiESTA on real traffic scenarios. We use Graphite Multicore Simulator [58] for simulating a sub-set of both SPLASH-2 and PARSEC benchmarks and obtaining the detailed communication trace files that are subsequently used by Noxim for network simulation. In particular, we have selected four SPLASH-2 benchmarks, *i.e.*, FFT, RADIX, LU, and WATER [101], and four PARSEC benchmarks, *i.e.*, CANNEAL, DEDUP FLUIDANIMATE (FLUID) and VIPS [13]. Fig. 7.13 shows the wireless utilization for different WiNoC architectures. For a given WiNoC architecture, the utilization of the wireless sub-network shows a not so relevant variability which is confined in a $\pm 3\%$ band. On the other side, for a given benchmark, the wireless utilization is on average low (less than 10%) for WiNoCs with 4 and 8 radio-hubs, and it rapidly increase as the 8 radio-hubs threshold is exceeded.

The corresponding energy savings are shown in Fig. 7.14. The general trend mimics that of wireless utilization. On average, the use of SiESTA on WiNoCs with 4, 8, 12, and 16 radio-hubs results in 7%, 14%, 21%, and 24% energy saving, respectively.




Figure 7.14: Energy saving for different WiNoC architectures and for different benchmarks.

7.3 Conclusions

In this chapter we presented SiESTA, a novel power management technique specifically designed for achieving energy efficiency of radio-hubs in a Wireless Network-on-Chip. The proposed power management strategy has been assessed on different network architectures and under several traffic scenarios, obtaining energy saving results ranging from 7% up to 25% without any impact on performance metrics. A concrete implementation has been presented to demonstrate its negligible impact on the silicon area of the radio-hub.

Chapter 8

Conclusion

The on-chip communication network accounts for a significant fraction of the overall energy budget of a multi/many-core system. In particular, the crossbar and the links which connect the routers are the main responsible for the energy consumption of the NoC. While reducing the voltage swing of these energy hungry elements has a positive effect in terms of energy saving, on the other side, the communication reliability decreases due to the increase of the bit-error-rate (BER). Starting from the assumption that, in general, not all the communications in an application have the same reliability requirements, in this thesis we have presented methods and architectures for runtime tuning the voltage swing for signaling in crossbars and links traversed by the flits of a packet based on the communication reliability requirement of that particular communication. Experiments carried out both synthetic and real traffic scenarios have shown the effectiveness of the proposed technique in terms of energy saving. As compared to the state-of-the-art in link energy reduction through data encoding schemes, the proposed technique provides higher energy saving without impacting the performance metrics of the system. A possible extension of the proposed technique is annotating the communication graph, rather than with only two reliability levels, with the required BER for the specific communication. Such opportunity will be investigated in future developments. Further, as the proposed technique has been thought to be applied in the context of message-passing multicore architectures, future work will investigate its application on cache-coherent networked multicore architectures.

Even if the proposed technique is applied, mainly due to its multi-hop

nature, traditional NoCs will not satisfy the communication requirements for future System-on-Chip both in terms of energy and latency. Technologies like WiNoCs are thus considered as a viable solution for facing the scalability and the energy consumption issues in many-core system architectures. Unfortunately, the transceiver of the radio hub in a WiNoC accounts for a significant fraction of the overall communication energy budget. In this research, we have presented a reliability aware runtime tunable transmitting power technique for improving the energy efficiency of the transceiver in WiNoC architectures. The proposed technique is general and can be applied to any WiNoC architecture. In this research, it has been applied to three known WiNoC architectures, namely, iWise64, McWiNoC, and HmWNoC. The experimental results have shown important energy saving up to 60% without any impact on performance metrics. The hardware overhead, in terms of silicon area, introduced by the proposed technique is negligible as compared with the area of the transceiver (approximately four orders of magnitude less than the transceiver).

Another interesting open research point it concerns the antenna. In fact, works present in literature in the context of low power WiNoC architectures, do not take into account the impact of the antenna directivity on power metrics. In addition, they assume that antennas present an omnidirectional radiation pattern which is far away from the reality. After that the adaptive transmitting power has been introduced, we have highlighted the need for antennas orientation design space exploration for improving the energy figures of WiNoC architectures. We have considered three main scenarios, namely, application specific (AS), general purpose (GP), and worst case (WC). In the AS scenario, communication information are exploited for optimizing the orientation of the antennas in such a way to maximize the overlap of the radiation patterns of the antennas which communicate more. The GP scenario, is derived from AS by assuming that all the radio hubs communicate each other with the same probability. Finally, the WC scenario is when the transceiver does not implement any transmitting power on-line calibration scheme and, therefore, all the radio hubs communicate using the same transmitting power irrespective of their location in the chip. A state of the art small-world based WiNoC architecture has been used as reference WiNoC architecture in the experiments. The exploration of the antennas orientation

design space under different traffic patterns resulted in important energy savings in all the three scenarios considered.

Since the transmitter is not the only responsible of transceiver power consumption we presented SiESTA, a novel power management technique specifically designed for achieving energy efficiency of radio-hubs in a Wireless Network-on-Chip. The proposed power management strategy has been assessed on different network architectures and under several traffic scenarios, obtaining energy saving results ranging from 7% up to 25% without any impact on performance metrics. A concrete implementation has been presented to demonstrate its negligible impact on the silicon area of the radio-hub.

After the introduction of proposed techniques, a designer can optimize energy consumption of both transmitter and receiver of a WiNoC. As future development, it will be interesting to study the application of such technique in a unified framework, evaluating energy saving when proposed solutions are applied at the same time.

Appendices

Appendix A

Scientific Production

The contributions of this thesis have been adapted and/or published both in journals and conferences. The publications related to the work of this thesis are as follows.

Journal publications:

- Vincenzo Catania, Andrea Mineo, Salvatore Monteleone, Maurizio Palesi and Davide Patti. “*Cycle-Accurate Network on Chip Simulation with Noxim*”, ACM Transaction on Modeling and Computer Simulation. 9, 4, Article 39 (March 2015).
- Andrea Mineo, Maurizio Palesi, Giuseppe Ascia, Partha Pratim Pande and Vincenzo Catania. “*On-Chip Communication Energy Reduction through Reliability Aware Adaptive Voltage Swing Scaling*”, To appear in IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD).
- Andrea Mineo, Maurizio Palesi, Giuseppe Ascia and Vincenzo Catania. “*Exploiting antenna directivity in wireless NoC architectures*”, Microprocessors and Microsystems (MICPRO), 46, pages 59-66, 2016.
- Mineo Andrea, Maurizio Palesi, Giuseppe Ascia and Vincenzo Catania. “*Runtime Tunable Transmitting Power Technique in mm-Wave WiNoC Architectures*”, IEEE Transactions on Very Large Scale Integration Systems (TVLSI) 2015.
- Mohd Shahrizal Rusli, Andrea Mineo, Maurizio Palesi, Giuseppe Ascia, Vincenzo Catania, Ooi Chia Yee, M. N. Marsono. “*A Closed Loop*

Power Manager for Transmission power Control in Wireless Network-on-Chip Architectures”, Jurnal Teknologi, www.jurnalteknologi.utm.my, June 2015.

- Maurizio Palesi, Mario Collotta , Andrea Mineo and Vincenzo Catania. “*An Efficient Radio Access Control Mechanism for Wireless Network-On-Chip Architectures*”, Journal of Low Power Electronics and Applications, ISSN 2079-9268, www.mdpi.com/journal/jlpea. 27/11/2015, 5, 38-56
- Vincenzo Catania, Andrea Mineo, Salvatore Monteleone, Davide Patti. “*Distributed Topology Discovery in Self-Assembled Nano Network-On-Chip*”, Computers & Electrical Engineering, Elsevier, Volume 40, Issue 8, November 2014, Pages 292–306.

Conference publications:

- Vincenzo Catania, Andrea Mineo, Salvatore Monteleone, Maurizio Palesi and Davide Patti. “*Energy Efficient Transceiver in Wireless Network on Chip Architectures*”, Design, Automation & Test in Europe Conference (DATE 2016), 14-18 March 2016, Dresden, Germany.
 - Vincenzo Catania, Andrea Mineo, Salvatore Monteleone, Maurizio Palesi and Davide Patti. “*Improving the Energy Efficiency of Wireless Network on Chip Architectures through Online Selective Buffers and Receivers Shutdown*”, The 13th Annual IEEE Consumer Communications & Networking Conference (CCNC2016), 9-12 January 2016, Las Vegas, USA.
 - Vincenzo Catania, Andrea Mineo, Salvatore Monteleone, Maurizio Palesi and Davide Patti. “*Noxim: An Open, Extensible and Cycle-accurate Network on Chip Simulator*”, 26th IEEE International Conference on Application-specific Systems, Architectures and Processors (ASAP15) , 17-19 July, Toronto, Canada.
 - Andrea Mineo, Maurizio Palesi, Giuseppe Ascia, Vincenzo Catania. “*Exploiting Antenna Directivity in Wireless NoC Architectures*”, Design Automation Conference (DAC 2015), San Francisco, CA, June 7-11, 2015 - Accepted as poster publication.
-

- Andrea Mineo, Mohd Shahrizal Rusli, Maurizio Palesi, Giuseppe Ascia, Vincenzo Catania and M. N. Marsono. “*A Closed Loop Transmitting Power Self-Calibration Scheme for Energy Efficient WiNoC Architectures*”, Design, Automation & Test in Europe Conference (DATE 2015), 9-13 March 2015, Grenoble, France.
 - Vincenzo Catania, Andrea Mineo, Salvatore Monteleone, Davide Patti. “*A Low-resource and Scalable Strategy for Segment Partitioning of Many-core Nano Networks*”, Second ACM International Workshop on Manycore Embedded System (MES 14), In conjunction with the 41st Internatinal Symposium on Computer Architecture (ISCA 2014), June 14-18 Minneapolis, USA.
 - Mohd Shahrizal Rusli, Andrea Mineo, Maurizio Palesi, Vincenzo Catania and M.N. Morsono. “*A Closed Loop Control based Power Manager for WiNoC Architectures*”, Second ACM International Workshop on Manycore Embedded System (MES 14), In conjunction with the 41st Internatinal Symposium on Computer Architecture (ISCA 2014), June 14-18 Minneapolis, USA.
 - Davide Patti, Andrea Mineo, Salvatore Monteleone, and Vincenzo Catania. “*Topology Discovery in Deadlock Free Self-Assembled DNA Networks*”, 3rd Computer Science On-line Conference 2014, CSOC14.
 - Andrea Mineo, Maurizio Palesi, Giuseppe Ascia, Vincenzo Catania. “*An Adaptive Transmitting Power Technique for Energy Efficient mm-Wave Wireless NoCs*”, Design, Automation & Test in Europe Conference (DATE 2014), 24-28 March 2014, Dresden, Germany.
 - Vincenzo Catania, Andrea Mineo, Salvatore Monteleone, Davide Patti. “*A First Effort for a Distributed Segment-based Approach on Self-Assembled Nano Networks*”, 6th International Workshop on Network on Chip Architectures (NoCArc 2013). To be held in conjunction with the 46th Annual IEEE/ACM International Symposium on Microarchitecture, December 7 (or 8), 2013, Davis, California.
 - Marina Masi, Andrea Mineo, Maurizio Palesi, Giuseppe Ascia, Vincenzo Catania. “*Low Energy Mapping Techniques under Reliability*
-

and Bandwidth Constraints”, 1th IEEE/IFIP International Conference on Embedded and Ubiquitous Computing (EUC 2013), Zhangjiajie, China, November 13-15, 2013.

- Andrea Mineo, Maurizio Palesi, Giuseppe Ascia, Vincenzo Catania. “*Runtime Online Links Voltage Scaling for Low Energy Networks on Chip*”, EUROMICRO DSD/SEAA 2013, Santander, Spain, September 4-6, 2013.
 - Andrea Mineo, Maurizio Palesi, Giuseppe Ascia, Vincenzo Catania. “*NoC Links Energy Reduction through Link Voltage Scaling*”, 13th International Conference on Embedded Computer Systems: Architectures, Modeling, and Simulation (SAMOS XIII), Samos, Greece, July 15–18, 2013.
-

Bibliography

- [1] Ansoft HFSS.
- [2] NanGate 45nm open cell library.
- [3] ITRS 2011 edition – interconnect. International Technology Roadmap for Semiconductors, 2011.
- [4] ITRS 2011 edition - system drivers. International Technology Roadmap for Semiconductors, 2011.
- [5] ITRS 2012 update – rf and analog/mixed-signal technologies (rfams). International Technology Roadmap for Semiconductors, 2012.
- [6] S. Abadal, E. Alarcón, A. Cabellos-Aparicio, M. C. Lemme, and M. Nemirovsky. Graphene-enabled wireless communication for massive multi-core architectures. *Communications Magazine, IEEE*, 51(11):137–143, November 2013.
- [7] S. Abadal, M. Iannazzo, M. Nemirovsky, A. Cabellos-Sparicio, H. Lee, and E. Alarcon. On the area and energy scalability of wireless network-on-chip: A model-based benchmarked design space exploration. *Transactions on Networking*, PP(99), 2014.
- [8] B. E. S. Akgul, L. N. Chakrapani, P. Korkmaz, and K. V. Palem. Probabilistic CMOS technology: A survey and future directions. In *International Conference on Very Large Scale Integration*, pages 1–6, 2006.
- [9] AMD. FX 8-core processor.
- [10] C. Balanis. *Modern Antenna Handbook*. Wiley, 2008.

-
- [11] R. G. Beausoleil, P. J. Kuekes, G. S. Snider, S. Y. Wang, and R. S. Williams. Nanoelectronic and nanophotonic interconnect. *Proceedings of the IEEE*, 96(2):230–247, Feb 2008.
- [12] L. Benini and G. D. Micheli. Networks on chips: a new SoC paradigm. *IEEE Computer*, 35(1):70–78, Jan. 2002.
- [13] C. Bienia, S. Kumar, J. P. Singh, and K. Li. The parsec benchmark suite: Characterization and architectural implications. In *International Conference on Parallel Architectures and Compilation Techniques*, Oct. 2008.
- [14] Cadence Design Systems Inc. EDI system user guide. Technical report, Cadence, 2014.
- [15] L. P. Carloni, P. Pande, and Y. Xie. Networks-on-chip in emerging interconnect paradigms: Advantages and challenges. In *Networks-on-Chip, 2009. NoCS 2009. 3rd ACM/IEEE International Symposium on*, pages 93–102, 2009.
- [16] V. Catania, A. Mineo, S. Monteleone, M. Palesi, and D. Patti. Noxim: An open, extensible and cycle-accurate network on chip simulator. In *IEEE International Conference on Application-specific Systems, Architectures and Processors*, Toronto, Canada, July 2015.
- [17] V. Catania, A. Mineo, S. Monteleone, M. Palesi, and D. Patti. Cycle-accurate network on chip simulation with noxim. *ACM Trans. Model. Comput. Simul.*, 27(1):4:1–4:25, Aug. 2016.
- [18] K. Chang, S. Deb, A. Ganguly, X. Yu, S. P. Sah, P. P. Pande, B. Belzer, and D. Heo. Performance evaluation and design trade-offs for wireless network-on-chip architectures. *J. Emerg. Technol. Comput. Syst.*, 8(3):23:1–23:25, Aug. 2012.
- [19] M. Chang, J. Cong, A. Kaplan, M. Naik, G. Reinman, E. Socher, and S.-W. Tam. Cmp network-on-chip overlaid with multi-band rf-interconnect. In *High Performance Computer Architecture, 2008. HPCA 2008. IEEE 14th International Symposium on*, pages 191–202, 2008.
-

-
- [20] M. C. F. Chang, J. Cong, A. Kaplan, C. Liu, M. Naik, J. Premkumar, G. Reinman, E. Socher, and S.-W. Tam. Power reduction of cmp communication networks via rf-interconnects. In *Microarchitecture, 2008. MICRO-41. 2008 41st IEEE/ACM International Symposium on*, pages 376–387, 2008.
- [21] G. Chen, M. A. Anders, H. Kaul, S. K. Satpathy, S. K. Mathew, S. K. Hsu, A. Agarwal, R. K. Krishnamurthy, S. Borkar, and V. De. 16.1 A 340mV-to-0.9V 20.2Tb/s Source-synchronous Hybrid packet/circuit-switched 16x16 Network-on-Chip in 22nm Tri-Gate CMOS. In *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2014 IEEE International*, pages 276–277, Feb 2014.
- [22] L. Couch. *Digital and Analog Communication Systems*. Pearson international edition. Pearson/Prentice Hall, 2007.
- [23] R. Courtland. What intel’s xeon phi coprocessor means for the future of supercomputing. *IEEE Spectrum*, 2013.
- [24] W. J. Dally and B. Towles. *Principles and Practices of Interconnection Networks*. Morgan Kaufmann, San Francisco, CA, 2004.
- [25] D. Daly and A. Chandrakasan. An energy-efficient ook transceiver for wireless sensor networks. *Solid-State Circuits, IEEE Journal of*, 42(5):1003–1011, 2007.
- [26] S. Deb, K. Chang, M. Cosic, A. Ganguly, P. P. Pande, D. Heo, and B. Belzer. CMOS compatible many-core noc architectures with multi-channel millimeter-wave wireless links. In *Proceedings of the great lakes symposium on VLSI, GLSVLSI ’12*, pages 165–170, New York, NY, USA, 2012. ACM.
- [27] S. Deb, K. Chang, A. Ganguly, X. Yu, C. Teuscher, P. Pande, D. Heo, and B. Belzer. Design of an efficient noc architecture using millimeter-wave wireless links. In *Quality Electronic Design (ISQED), 2012 13th International Symposium on*, pages 165–172, 2012.
- [28] S. Deb, K. Chang, X. Yu, S. Sah, M. Cosic, A. Ganguly, P. Pande, B. Belzer, and D. Heo. Design of an energy-efficient cmos-compatible
-

- noc architecture with millimeter-wave wireless interconnects. *Computers, IEEE Transactions on*, 62(12):2382–2396, Dec 2013.
- [29] S. Deb, A. Ganguly, K. Chang, P. Pande, B. Beizer, and D. Heo. Enhancing performance of network-on-chip architectures with millimeter-wave wireless interconnects. In *Application-specific Systems Architectures and Processors (ASAP), 2010 21st IEEE International Conference on*, pages 73–80, 2010.
- [30] S. Deb, A. Ganguly, P. Pande, B. Belzer, and D. Heo. Wireless noc as interconnection backbone for multicore chips: Promises and challenges. *Emerging and Selected Topics in Circuits and Systems, IEEE Journal on*, 2(2):228–239, 2012.
- [31] Y. Z. Ding and M. O. Rabin. Hyper-encryption and everlasting security. In *Annual Symposium on Theoretical Aspects of Computer Science*, pages 1–26, 2002.
- [32] D. DiTomaso, A. Kodi, S. Kaya, and D. Matolak. iwise: Inter-router wireless scalable express channels for network-on-chips (nocs) architecture. In *High Performance Interconnects (HOTI), 2011 IEEE 19th Annual Symposium on*, pages 11–18, 2011.
- [33] A. Ejlali, B. M. Al-Hashimi, P. Rosinger, S. G. Miremadi, and L. Benini. Performability/energy tradeoff in error-control schemes for on-chip networks. *IEEE Transactions on Computer Aided Design of Integrated Circuits and Systems*, 18(1):1–14, Jan. 2010.
- [34] C. S. et al. Dsent - a tool connecting emerging photonics with electronics for opto-electronic networks-on-chip modeling. pages 201–210, May 2012.
- [35] F. Fazzino, M. Palesi, and D. Patti. Noxim: Network-on-Chip simulator. <http://noxim.sourceforge.net>.
- [36] F. Fazzino, M. Palesi, and D. Patti. winoxim: Wireless Network-on-Chip simulator. <https://code.google.com/p/winoxim/>.
-

-
- [37] A. Flores, J. L. Aragón, and M. E. Acacio. An energy consumption characterization of on-chip interconnection networks for tiled cmp architectures. *The Journal of Supercomputing*, 45(3):341–364, 2008.
- [38] B. Floyd, C.-M. Hung, and K. O. Intra-chip wireless interconnect for clock distribution implemented with integrated antennas, receivers, and transmitters. *Solid-State Circuits, IEEE Journal of*, 37(5):543–552, 2002.
- [39] H. Fuks. Non-deterministic density classification with diffusive probabilistic cellular automata. *Physical Review*, 66, July 2002.
- [40] S. H. Gade, H. K. Mondal, and S. Deb. A hardware and thermal analysis of dvfs in a multi-core system with hybrid wnoc architecture. In *VLSI Design (VLSID), 2015 28th International Conference on*, pages 117–122, Jan 2015.
- [41] A. Ganguly, K. Chang, S. Deb, P. Pande, B. Belzer, and C. Teuscher. Scalable hybrid wireless network-on-chip architectures for multicore systems. *Computers, IEEE Transactions on*, 60(10):1485–1502, 2011.
- [42] A. Ganguly, P. P. Pande, and B. Belzer. Crosstalk-aware channel coding schemes for energy efficient and reliable NOC interconnects. *IEEE Transactions on Very Large Scale Integration Systems*, 17(11):1626–1639, Nov. 2009.
- [43] P. Garrou, C. Bower, and P. Ramm. *Handbook of 3D Integration*. Wiley-VCH, 2008.
- [44] P. R. Gray, P. J. Hurst, S. H. Lewis, and R. G. Meyer. *Analysis and Design of Analog Integrated Circuits*. Wiley, 5th edition, 2009.
- [45] F. Gutierrez, S. Agarwal, K. Parrish, and T. Rappaport. On-chip integrated antenna structures in cmos for 60 ghz wpan systems. *Selected Areas in Communications, IEEE Journal on*, 27(8):1367–1378, 2009.
- [46] J. Han and M. Orshansky. Approximate computing: An emerging paradigm for energy-efficient design. In *Test Symposium (ETS), 2013 18th IEEE European*, pages 1–6, May 2013.
-

-
- [47] R. Ho, K. W. Mai, S. Member, and M. A. Horowitz. The future of wires. In *Proceedings of the IEEE*, pages 490–504, 2001.
- [48] J. Hu and R. Marculescu. Energy- and performance-aware mapping for regular NoC architectures. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 24(4):551–562, Apr. 2005.
- [49] E. G. T. Jaspers and P. H. N. de With. Chip-set for video display of multimedia information. *IEEE Transactions on Consumer Electronics*, 45(3):706–715, Aug. 1999.
- [50] J. Kim, K. Choi, and G. Loh. Exploiting new interconnect technologies in on-chip communication. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, 2(2):124–136, June 2012.
- [51] S.-B. Lee, S.-W. Tam, I. Pefkianakis, S. Lu, M. F. Chang, C. Guo, G. Reinman, C. Peng, M. Naik, L. Zhang, and J. Cong. A scalable micro wireless interconnect structure for CMPs. In *Proceedings of the 15th annual international conference on Mobile computing and networking, MobiCom '09*, pages 217–228, New York, NY, USA, 2009. ACM.
- [52] F. Li, C. Nicopoulos, T. Richardson, Y. Xie, V. Narayanan, and M. Kandemir. Design and management of 3d chip multiprocessors using network-in-memory. In *33rd International Symposium on Computer Architecture (ISCA '06)*, pages 130–141, 2006.
- [53] J. Lin, H.-T. Wu, Y. Su, L. Gao, A. Sugavanam, J. Brewer, and K. O. Communication using antennas fabricated in silicon integrated circuits. *Solid-State Circuits, IEEE Journal of*, 42(8):1678–1687, 2007.
- [54] O. Lysne, T. Skeie, S. A. Reinemo, and I. Theiss. Layered routing in irregular networks. *IEEE Transactions on Parallel and Distributed Systems*, 17(1):51–65, Jan 2006.
- [55] D. J. C. MacKay. Bayesian interpolation. *Neural Computing*, 4(3):415–447, May 1992.
- [56] N. Mansoor, P. J. S. Iruthayaraj, and A. Ganguly. Design methodology for a robust and energy-efficient millimeter-wave wireless network-on-
-

- chip. *IEEE Transactions on Multi-Scale Computing Systems*, 1(1):33–45, Jan 2015.
- [57] H. Matsutani, M. Koibuchi, H. Amano, and T. Yoshinaga. Prediction router: A low-latency on-chip router architecture with multiple predictors. *IEEE Transactions on Computers*, 60(6):783–799, June 2011.
- [58] J. E. Miller, H. Kasture, G. Kurian, C. G. III, N. Beckmann, C. Celio, J. Eastep, and A. Agarwal. Graphite: A distributed parallel simulator for multicores. In *IEEE International Symposium on High-Performance Computer Architecture*, Jan. 2010.
- [59] A. Mineo, M. Palesi, G. Ascia, and V. Catania. An adaptive transmitting power technique for energy efficient mm-wave wireless nocs. In *Design, Automation Test in Europe Conference (DATE14)*, Mar. 2014.
- [60] A. Mineo, M. S. Rusli, M. Palesi, G. Ascia, V. Catania, and M. N. Marsono. A closed loop transmitting power self-calibration scheme for energy efficient winoc architectures. In *Design, Automation Test in Europe Conference (DATE15)*, Mar. 2015.
- [61] Mondal, H. K. Mondal, and S. Deb. An energy efficient wireless network-on-chip using power-gated transceivers. In *System-on-Chip Conference (SOCC), 2014 27th IEEE International*, pages 243–248, Sept 2014.
- [62] H. Mondal, S. Gade, M. Shamim, S. Deb, and A. Ganguly. Interference-aware wireless network-on-chip architecture using directional antennas. *IEEE Transactions on Multi-Scale Computing Systems*, PP(99):1–1, 2016.
- [63] H. K. Mondal and S. Deb. Energy efficient on-chip wireless interconnects with sleepy transceivers. In *Design and Test Symposium (IDT), 2013 8th International*, pages 1–6, Dec 2013.
- [64] H. K. Mondal, N. S. Harsha, and S. Deb. An efficient hardware implementation of dvfs in multi-core system with wireless network-on-chip. In *VLSI (ISVLSI), 2014 IEEE Computer Society Annual Symposium on*, pages 184–189, July 2014.
-

-
- [65] D. C. Montgomery. *Design and Analysis of Experiments*. Wiley, 8 edition, 2012.
- [66] S. Montusclat, F. Gianesello, D. Gloria, and S. Tedjini. Silicon integrated antenna developments up to 80 ghz for millimeter wave wireless links. In *Wireless Technology, 2005. The European Conference on*, pages 237–240, 2005.
- [67] K. O, K. Kim, B. Floyd, J. Mehta, H. Yoon, C.-M. Hung, D. Bravo, T. Dickson, X. Guo, R. Li, N. Trichy, J. Caserta, I. Bomstad, W.R., J. Branch, D.-J. Yang, J. Bohorquez, E. Seok, L. Gao, A. Sugavanam, J. J. Lin, J. Chen, and J. Brewer. On-chip antennas in silicon ics and their application. *Electron Devices, IEEE Transactions on*, 52(7):1312–1323, 2005.
- [68] U. Ogras and R. Marculescu. "it's a small world after all": Noc performance optimization via long-range link insertion. *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, 14(7):693–706, 2006.
- [69] S. Orfanidis. Electromagnetic waves and antennasn, online @ <http://www.ece.rutgers.edu/orfanidi/ewa/>.
- [70] K. Palem and A. Lingamneni. What to do about the end of moore's law, probably! In *Design Automation Conference*, pages 924–929, 2012.
- [71] K. V. Palem. Energy aware computing through probabilistic switching: a study of limits. *IEEE Transactions on Computers*, 54(9):1123–1137, 2005.
- [72] M. Palesi, G. Ascia, F. Fazzino, and V. Catania. Data encoding schemes in networks on chip. *IEEE Transactions on Computer Aided Design of Integrated Circuits and Systems*, 30(5), May 2011.
- [73] M. Palesi, M. Collotta, A. Mineo, and V. Catania. An efficient radio access control mechanism for wireless network-on-chip architectures. *Journal of Low Power Electronics and Applications*, 5(2):38–56, 2015.
- [74] M. Palesi, R. Tornero, J. M. O. na, D. Panno, and V. Catania. Designing robust routing algorithms and mapping cores in networks-on-chip:
-

- A multi-objective evolutionary-based approach. *Journal of Universal Computer Science*, 18(7):937–969, 2012.
- [75] D. Pamunuwa, J. Öberg, L.-R. Zheng, A. J. Mikael Millberg, and H. Tenhunen. Layout, performance and power trade-offs in mesh-based network-on-chip architectures. In *International Conference on Very Large Scale Integration*, pages 362–367, 2003.
- [76] P. Pande, C. Grecu, M. Jones, A. Ivanov, and R. Saleh. Performance evaluation and design trade-offs for network-on-chip interconnect architectures. *IEEE Transactions on Computer*, 54(8):1025–1040, Aug. 2005.
- [77] P. P. Pande, A. Ganguly, H. Zhu, and C. Grecu. Energy reduction through crosstalk avoidance coding in networks on chip. *Journal of Systems Architecture*, 54:441–451, 2008.
- [78] D. Park, S. Eachempati, R. Das, A. K. Mishra, Y. Xie, N. Vijaykrishnan, and C. R. Das. Mira: A multi-layered on-chip interconnect router architecture. In *Computer Architecture, 2008. ISCA '08. 35th International Symposium on*, pages 251–261, June 2008.
- [79] W. Paul, K. Ryan, M. Jacob, X. Yu, P. P. Pande, G. Amlan, and H. Deukhyoun. Design space exploration for wireless nocs incorporating irregular network routing. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, 33(11):1732–1745, Nov 2014.
- [80] J.-M. Philippe, S. Pillement, and O. Sentieys. Area efficient temporal coding schemes for reducing crosstalk effects. *IEEE Proceedings of the 7th international symposium on Quality Electronic Design*, 2006.
- [81] J. M. Rabaey, A. Chandrakasan, and B. Nikolic. *Digital Integrated Circuits (2nd Edition)*. Prentice Hall, 2006.
- [82] B. Razavi. *RF Microelectronics*. Prentice Hall, second edition, 2012.
- [83] P. K. Sahu and S. Chattopadhyay. A survey on application mapping strategies for network-on-chip design. *Journal of Systems Architecture*, 59(1), 2013.
-

-
- [84] E. Seok and K. Kenneth. Design rules for improving predictability of on-chip antenna characteristics in the presence of other metal structures. In *Interconnect Technology Conference, 2005. Proceedings of the IEEE 2005 International*, pages 120–122, 2005.
- [85] A. Shacham, K. Bergman, and L. P. Carloni. Photonic networks-on-chip for future generations of chip multiprocessors. *IEEE Transactions on Computers*, 57(9):1246–1260, Sept 2008.
- [86] S. R. Sridhara and N. R. Shanbhag. Coding for system-on-chip networks: A unified framework. *IEEE Transactions on Very Large Scale Integration Systems*, 13(6):655–667, June 2005.
- [87] M. R. Stan and W. P. Burlison. Bus invert coding for low power I/O. *IEEE Transactions on Very Large Scale Integration Systems*, 3:49–58, Mar. 1995.
- [88] F. Steenhof, H. Duque, B. Nilsson, K. Goossens, and R. P. Llopis. Networks on chips for high-end consumer-electronics TV system architectures. In *Conference on Design, Automation and Test in Europe*, pages 148–153, 2006.
- [89] I. E. Sutherland, R. F. Sproull, and D. Harris. *Logical Effort: Designing Fast CMOS Circuits*. Morgan Kaufmann Publishers, 1999.
- [90] Synopsys Inc. IC compiler user guide: Implementation. Technical report, Synopsys, 2014.
- [91] Y. Take, H. Matsutani, D. Sasaki, M. Koibuchi, T. Kuroda, and H. Amano. 3D NoC with inductive-coupling links for building-block SiPs. *Computers, IEEE Transactions on*, 63(3):748–763, March 2014.
- [92] A. S. Tanenbaum. *Computer Networks*. Prentice Hall, fourth edition, 2003.
- [93] M. B. Taylor. Is dark silicon useful? harnessing the four horsemen of the coming dark silicon apocalypse. In *Proceedings of the 49th Annual Design Automation Conference (DAC12)*, pages 1131–1136, 2012.
- [94] Tiler. TILE-Gx8072 processor.
-

-
- [95] Y. F. Tsai, F. Wang, Y. Xie, N. Vijaykrishnan, and M. J. Irwin. Design space exploration for 3-d cache. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 16(4):444–455, April 2008.
- [96] K. Tushar, K. Amit, C. Patrick, E. Mattan, and L.-S. Peh. Noc with near-ideal express virtual channels using global-line communication. In *High Performance Interconnects, 2008. HOTI '08. 16th IEEE Symposium on*, pages 11–20, Aug 2008.
- [97] E. B. van der Tol and E. G. Jaspers. Mapping of MPEG-4 decoding on a flexible architecture platform. *Media Processors*, 4674:362–375, 2002.
- [98] S. R. Vangal, J. Howard, G. Ruhl, S. Dighe, H. Wilson, J. Tschanz, D. Finan, A. Singh, T. Jacob, S. Jain, V. Erraguntla, C. Roberts, Y. Hoskote, N. Borkar, and S. Borkar. An 80-tile sub-100-W TeraFLOPS processor in 65-nm CMOS. *IEEE Journal of Solid-State Circuits*, 43(1):29–41, Jan. 2008.
- [99] C. Wang, W.-H. Hu, and N. Bagherzadeh. A wireless network-on-chip design for multicore platforms. In *Parallel, Distributed and Network-Based Processing (PDP), 2011 19th Euromicro International Conference on*, pages 409–416, 2011.
- [100] P. Wettin, J. Murray, P. Pande, B. Shirazi, and A. Ganguly. Energy-efficient multicore chip design through cross-layer approach. In *Design, Automation Test in Europe Conference Exhibition (DATE), 2013*, pages 725–730, March 2013.
- [101] S. C. Woo, M. Ohara, E. Torrie, J. P. Singh, and A. Gupta. The SPLASH-2 programs: characterization and methodological considerations. In *International Symposium on Computer Architecture*, pages 24–36, June 1995.
- [102] Y. Xie, G. H. Loh, B. Black, and K. Bernstein. Design space exploration for 3d architectures. *J. Emerg. Technol. Comput. Syst.*, 2(2):65–103, Apr. 2006.
-

-
- [103] S.-R. Yoon, J. Lee, and S.-C. Park. Case study: Noc based next-generation wlan receiver design in transaction level. In *International Conference on Advanced Communication Technology*, pages 1125–1128, 2006.
- [104] X. Yu, J. Baylon, P. Wettin, D. Heo, P. P. Pande, and S. Mirabbasi. Architecture and design of multi-channel millimeter-wave wireless network-on-chip. *Design Test, IEEE*, PP(99):1–1, 2014.
- [105] X. Yu, S. Sah, S. Deb, P. Pande, B. Belzer, and D. Heo. A wideband body-enabled millimeter-wave transceiver for wireless network-on-chip. In *Circuits and Systems (MWSCAS), 2011 IEEE 54th International Midwest Symposium on*, pages 1–4, 2011.
- [106] X. Yu, S. P. Sah, Rashtian, H. Rashtian, S. Mirabbasi, P. P. Pande, and D. Heo. A 1.2-pj/bit 16-gb/s 60-ghz ook transmitter in 65-nm cmos for wireless network-on-chip. *Microwave Theory and Techniques, IEEE Transactions on*, 62(10):2357–2369, Oct 2014.
- [107] H. Zhang, V. George, and J. M. Rabaey. Low-swing on-chip signaling techniques: Effectiveness and robustness. *IEEE Transactions on Very Large Scale Integration Systems*, 8(3):264–272, June 2000.
- [108] Y. P. Zhang, Z. M. Chen, and M. Sun. Propagation mechanisms of radio waves over intra-chip channels with integrated antennas: Frequency-domain measurements and time-domain analysis. *Antennas and Propagation, IEEE Transactions on*, 55(10):2900–2906, Oct 2007.
- [109] D. Zhao and Y. Wang. Sd-mac: Design and synthesis of a hardware-efficient collision-free qos-aware mac protocol for wireless network-on-chip. *Computers, IEEE Transactions on*, 57(9):1230–1245, 2008.
- [110] D. Zhao, Y. Wang, J. Li, and T. Kikkawa. Design of multi-channel wireless noc to improve on-chip communication capacity. In *Networks on Chip (NoCS), 2011 Fifth IEEE/ACM International Symposium on*, pages 177–184, 2011.
-