

# Power Management Controller for Online Power Saving in Network-on-Chips

Stephanie Friederich, Marco Neber, Jürgen Becker  
Karlsruhe Institute of Technology  
Karlsruhe, Germany  
Email: stephanie.friederich@kit.edu

**Abstract**—Growing chip integration density and increasing frequencies lead to tremendous leakage power and henceforth to chip heat problems. Power management is one possibility to reduce the power consumption and get the temperature problem under control. Current technology mainly focuses on power-gating techniques on basis of multi-core systems but leaving the network perspective out of scope. We provide a holistic concept, bringing together power-gating and frequency scaling techniques for network-on-chips. Following this, network static power consumption could be minimized without affecting the system performance. We present a light-weight power management controller for network-on-chips with online monitoring to optimize the power consumption of network resources. Our work comprises a hardware simulation model for design space exploration of varying technology specific parameters and an FPGS based prototype for verification. The power saving potential heavily depend on the network communication load. Our power controller adds only 2.1% of resources while 28.11% of the total power could be saved with clock gating.

**Index Terms**—Embedded Multicore/Many-core Systems-on-Chip; Network-on-chip; Low-power design;

## I. INTRODUCTION

*Network-on-chips (NoCs)* provide a scalable communication infrastructure for architectures with huge numbers of processing elements. To preclude that communication becomes the bottleneck in large multiprocessor designs, powerful network routers with large throughput are necessary. Consequently the proportion of network resources of the overall resource consumption is not negligible any longer.

Since the start of CPU development, processor frequency increased from few Megahertz to around three Gigahertz. However, in the last few years, core frequency growth has slowed. One of the reasons is a highly significant heat emission which comes in conjunction with increasing frequencies. Chip temperature can increase so much that it could cause actual physical damage to the chip. Because of the tremendous generated heat, simple cooling mechanisms are no longer strong enough to carry away the heat. To prevent the hardware from destruction, hardware monitors are included in the chips and switch off the device before it comes to actual damage. For taking corrective actions even before the heat increases too much, *dynamic voltage and frequency scaling (DVFS)* methods are applied [1]. The majority of DVFS methods focuses on the processor perspective and the communication is not considered. However, there are first developments to extend these methods also to the network layer [2], [3], [4].

With shrinking technology sizes, especially beyond 90nm, the amount of static power dissipation reaches a critical proportion of the total power. Accessory, the increasing static power dissipation causes high temperatures even in idle circuits and hence the static leakage is not negligible any more. The rising complexity of on-chip communication structures, is reaching a state where dedicated power management becomes beneficial. Moreover, with increasing network sizes it will not be possible to have one globally synchronous clock to supply the entire network, due to phase adaptations.

We propose a reduction of static power consumption of network resources without limiting the functionality of communication. To keep the balance between energy-saving and performance loss, a power management controlling unit was designed in this work. Therefore the existing network router was extended by a controller unit which evaluates monitoring information and controls the power shutdown mechanisms. In addition the controller gives recommendations about the optimal frequency in case of low network load.

The rest of the paper is organized as follows: In section II some selected previous works related to power-gating and low power networks are briefly discussed. Section III presents the fundamentals of on-chip power consumption and the original network architecture. In section IV the power management concept is described. The following section V provides simulation results to support the presented concept. In section VI we conclude the paper.

## II. RELATED WORK

There are several methods for power saving for *network-on-chip (NoC)* systems. They can be divided into two groups, focusing either on dynamic power reduction, or on reducing the static power. Reduction of dynamic power can be achieved by task or thread manipulation [5]. The reduction of static power is always linked with power- or clock-gating mechanisms or DVFS methods and is described in the following section. A general overview of low power designs is given in [6]. In [7] a power-gating scheme for virtual channels in on-chip networks is described, which uses an adaptive method to dynamically adjust the number of active VCs based on the on-chip traffic characteristics. But the proposed method does not comprise guaranteed service connection within the network. Also, other router resources than buffers will not be power-gated if they are inactive.

The following works are inspired by dark silicon research [8] and propose a different kind of power-gating granularity, where multiple physical networks run in parallel and each layer can be power-gated [9], [10]. The sub-networks can be power gated without compromising the connectivity of the primary network. Although the total power consumption of the multi-layer network can be reduced compared to single-layer networks, the resource consumption is not considered at all. Bokhari et al. [11] take the idea of multi-layer networks and describes the delay which is inserted when switching between the different layers. All these works have in common that they suffer a disadvantage in regards of resource overhead due to the design of sub-networks. To keep the resource overhead to a minimum due to power management, is one of the key features of this work.

Another method to achieve power reduction by reducing the dynamic power is DVFS. Lee et al. [12] presented a frequency and voltage scaling method based on the network workload. Another system-level DVFS technique is presented in [1] where higher level caches are comprised as well as network resources. Bogdan et al. [2], [4] present a control approach for NoC designs with multiple voltage and frequency islands. Their approaches focus on an analytical solution of the power management problem, which results in a high resource utilization and long decision times. As opposed to this, our method pursues a minimal size hardware implementation for the power management, with fast response to changing communication load since the complete logic is implemented in hardware. In addition the centralized approach in [2] does not scale well, in contrast to our approach.

Matsutani et al. present run-time power gating for network routers in several publications [13], [14]. But besides a not negligible resource overhead, additional computation overhead is necessary inside the router to produce the information required to generate a wake-up request signal. Our work in comparison does not require sophisticated routing algorithms and the resource overhead is kept to a minimum.

The number of switches between on and off states should be limited. First since each wake-up process results in power penalty and second because of high wake-up latencies [15], [16]. Both values highly depend on the applied technology and the depth of the sleep state [17]. Since there is no real chip available of our NoC architecture, the power measurements for evaluation of the design are based on simulation models. A basic performance and power evaluation model for NoC based interconnection networks is described in [18].

### III. BACKGROUND

#### A. Methodologies for Power Saving

The power consumption of integrated circuits can be divided into static and dynamic power. The dynamic power in turn is subdivided into internal (or leakage) and switching power. There are methods either to reduce the static power or to reduce the dynamic power. This paper addresses both power

reduction methods, since it combines clock scaling and power-gating techniques.

The switching power depends on the loading of output pins. Internal power results from the power dissipation of the internal short-circuit current when a signal changes from 0 to 1 or from 1 to 0. The switching activity can be determined from design simulations. If no simulation results are available, worst case values are assumed. In scope of this work we are using a standard-cell library from TSMC for power simulations.

The dynamic power of a gate is defined as follows:

$$P = \alpha * C_l * V_{dd}^2 * f \quad (1)$$

With:

- $\alpha$  referred to switching activity. The average number of 0→1 transitions in one clock cycle
- $C_l$  is the load capacitance
- $V_{dd}$  is the power supply voltage
- $f$  is the clock frequency

The transistor switch speed is proportional to the transistor voltage. If we take into account that the frequency can be raised proportionally to transistor switch speed only, the dynamic power dissipation is proportional to  $V^3$ . For example, if we double the frequency, the dynamic power dissipation will be eight times higher. Thus reducing the frequency can save a huge amount of dynamic power. Therefore we extended the power management controller by the opportunity to reduce the frequency if the routers are sparsely utilized.

The main part of the static power is caused by source/drain leakage current  $I_{sd,leak}$  and can be defined as follows:

$$P_{leakage} = I_{sd,leak} * V \quad (2)$$

Starting with the 90nm node and below, the leakage power becomes very significant. The difference between the supply voltage and the threshold voltage is decreasing and due to this the transistor delay increases exponentially. Lowering the threshold voltage results in an increased leakage power.

To gain low power circuits, the parameters of equation 1 and 2 should be minimized. Since the switching activity is given through the communication and the load capacitance is set by the used technology we can only minimize either the frequency or the voltage. These methods are made up of *Power Gating*, *Power Scaling*, *Clock Gating* and, *Clock Scaling* [7], [12]. Power gating adds sleep transistors to the design to separate a circuit's power supply. According to that, most of the static leakage is suppressed. In this work the reduction of leakage power of unused network components is the main goal. In regards of no significant reduction of system performance, it is important that on the one hand the network can be power gated during run time, depending on the actual network workload and on the other hand the frequency can be scaled up or down depending on the load of the router.

#### B. Top Level Architecture

The underlying NoC architecture we are using for this work is briefly described in the following. A detailed description of




Figure 1. Block diagram of the network router with a variable number of ports and an including monitoring infrastructure [19].

the network architecture can be found in [19]. The design is organized in a 2D-array mesh of tiles to increase scalability. Other topologies are also possible, but should not be further discussed in this work. Each router in this design has bidirectional connections to its neighbor routers and a connection to a local cluster. A network adapter builds the interface between the router and the local bus of the cluster.

The cluster architecture integrates a variable number of LEON3 cores, all connected to the local bus [20]. Furthermore, each cluster includes a L2 cache and a tile local memory. The distributed memory architecture reduces the network accesses, since tile local memory is sufficient for most of the computations. This arrangement has implications for the degree of network capacity utilization, since network accesses can be reduced. Nevertheless the simulation results in section V are based on synthetic generated network traffic, to prove that the concept is also valid for other network architectures and memory distributions with worst case communication scenarios.

### C. The NoC Router

The router design which forms the basis for this work was given by Heisswolf [21]. The router supports both connectionless *best effort* (BE) communication as well as connection-oriented *guaranteed service* (GS) communication. Both transmission strategies (BE and GS) share the router resources dynamically. As a routing strategy XY routing is used in this work.

The routers are connected with each other via two unidirectional links which are divided into a parameterizable number of *virtual channels* (VCs). A single router unit is shown in figure 1. The router input and output links, which are usually placed in cardinal directions to the neighboring router are grouped together in this figure.

The power saving principle described in this paper was designed for a packet switching network with wormhole flow control. But it can be also applied to other networks easily. The router must provide monitoring information about the

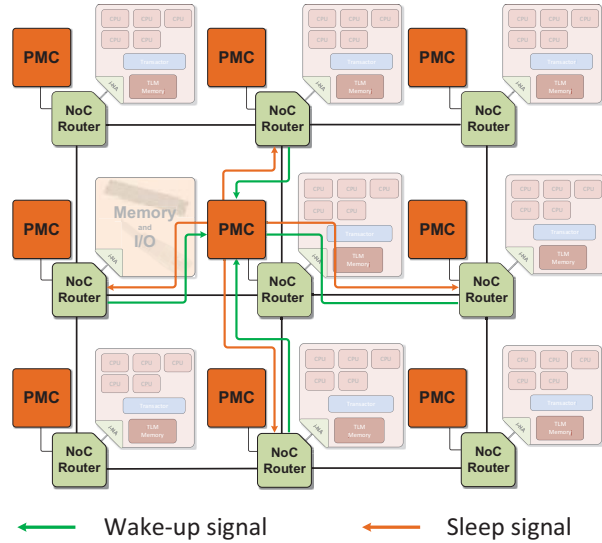


Figure 2. Schematic of a 3x3 mesh network, including one power management controller (PMC) for each router. The wake-up and sleep signals are only shown for the PMC in the middle, all others are connected accordingly.

current link utilization during run-time. In this work the monitoring infrastructure was already included in the basic router design and have been described in [19].

## IV. IMPLEMENTATION

The goal of the power management unit for on-chip networks is to decide whether it is reasonable to switch off a complete router or to scale down the frequency. Therefore a *power management unit* (PMU), composed of the *power management controller* (PMC) and the *Load Detection Unit* (LDU), have been developed. For efficient power gating functionality the PMC needs to observe the current utilization of the on-chip routers. Each router is already supplied with an implemented monitoring unit which is utilized by the PMC. The following section describes the hardware implementation in detail. The design implementation is written in *SystemVerilog* which offers the possibility to use the same files on the one hand for a *field programmable gate array* (FPGA) based prototype. On the other hand the design files can be utilized for an *application-specific integrated circuit* (ASIC) synthesis and the associated power estimation.

To come to a fast decision for every router while keeping the design scalable, there is one PMC connected to every router. The resulting resource overhead is proportionally quite small in contrast to the router design and is further evaluated in section V. Figure 2 shows a schematic of a 3x3 mesh NoC based on the architecture described in section III-B, where all routers have connection to their direct neighbors, the local cluster and to one PMC. To improve clarity, the wake-up and sleep signals are only given for the router in the center of figure 2. In the actual implementation all PMC have this two signals to all adjacent routers.




Figure 3. Block diagram of the power management unit, consists of the load detection unit inside the router, a clock generator, and the power management controller

The power management unit consists of the following three components. First the load detection unit which analyzes the monitoring information. Second the transmission control unit which extends the existing control unit so that no flits are transmitted as long as the following router is power-gated. And third the power management controller which decides whether the corresponding router can be power-gated or if the clock frequency should be scaled. An overview of the router including power management is given in figure 3. Of course the router includes still the components shown in figure 1, but for simplification they are not shown here again. In the following, the implementation of the load detection unit and the power management controller is described in detail.

#### A. Load Detection Unit (LDU)

The output of the monitoring unit is continuously evaluated by the LDU. While a router is active, the monitoring unit output values are greater than zero. The LDU checks in every single clock cycle if its input value is greater than zero. If it is equal zero for a complete given observation time ( $T_{idle\_detect}$ ), the router is marked as idle and hence the router could be switched off completely. Otherwise the busy\_out signal of the detector is set high to signalize a busy router device. The LDU is integrated into the router. The advantage of this approach is that, while the router is shut down the monitoring unit is also powered-off. While the appropriate PMC must be always active even when the router is shut down to watch out for newly arriving flits. Since the power management unit cannot be power gated, the goal is to keep it as small as possible in regards of resource and power consumption.

#### B. Power Management Controller (PMC)

The applied network provides two different communication types, *best effort (BE)* and *guaranteed service (GS)* communication. In case of BE all flits of one package are




Figure 4. Simplified state diagram of the power management controller for power gating and clock scaling

send right after each other. Hence the reserved channel will be utilized the whole time. In case of GS communication a channel will be reserved for a longer period and might not be utilized for any length of time.

The PMC has to signal the power state of its router to all adjacent devices to prevent the loss of data packets. The PMC sets a *hold signal* immediately when it detects an idle router, to block further outgoing transmissions by neighboring routers and its local tile. The *hold signal* interrupts the reservation process in the transmission control of the output virtual channels (VC). Hence flits are hold back till the neighboring router is activated again.

On the other hand, as soon as the arbitration module determines a new packet arrival, a *transmission request signal* is send to the router in the direction the packet should be forwarded, even before the actual arbitration process is finished. Consequently, power gated routers are started with their wake-up process, as early as possible to minimize the performance penalty due to power saving mechanisms. Isolation cells are integrated at the input of the router, to generate a known logic value during shutdown of the supply power.

The following describes the three main states of the power management controller.

**Power Down:** If the *Load Detection Unit (LDU)* signals that the router is inactive, the power management controller monitors for *transmission request signal* of adjacent routers, to avoid a false shut down while a transmission is already queued. Also the LDU checks whether there are channels reserved for GS communication. Only if both are cases are false, the PMC can take the next state for isolating the circuits of the router. Before the router is finally shut down on the next rising clock edge, the PMC checks again for incoming requests. If a request arises at this point in time (*Pre request* state in figure 4), the shutdown is canceled and the router stays active. Instead of isolating the router from the supply voltage, it is also possible to disconnect the clock source. Consequently, the

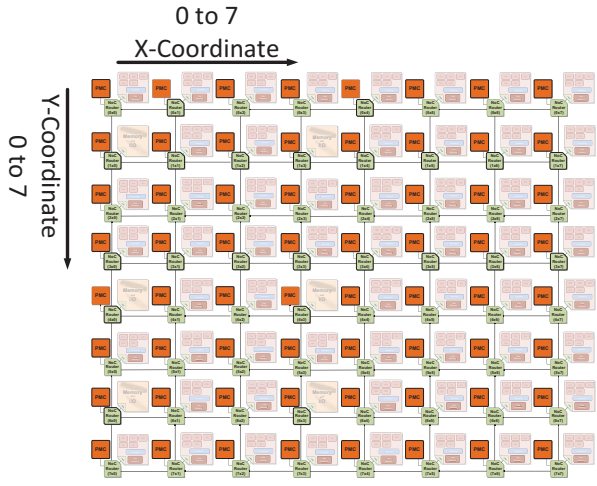


Figure 5. 8x8 mesh network architecture used for performance simulation

leakage power will not be reduced, but on the other hand the wake up delay will be much shorter.

**Wake-up Process:** In case of routing a flit to a power gated router, the device needs to be woken up. So the initiating source router sends a *transmission request signal* to the sleeping routers *power management controller (PMC)*. This causes the PMC to switch the routers sleep transistor and connect it to the systems supply. In the next step the PMC resets the virtual channel reservation table at the output port of the requesting router to release the blocked VCs. Afterwards the transmission is executed as usual. The time a router requires to wake up and to start its usual operation is defined as  $T_{wake-up}$  and is crucial for the added latency by the power gating functionality.  $T_{wake-up}$  depends on the used transistor technology. Since the technology is not known at this point in time, the delay time  $T_{wake-up}$  can be parametrized for simulation, to test different values against each other.

**Frequency Scaling:** In case of low loads or if channels are reserved by *guaranteed service (GS)* connections, the router cannot be power gated. To reduce power also in these cases, the PMC comprises a clock scaling unit.

The observe state is the initial and default state of the power management controller, see figure 4. Depending on the degree of router utilization, the controller either suggest a frequency scaling factor (scaling the frequency up or down) or power gates the complete router in case of no utilization. The router utilization upper and lower *threshold limit value (TLV)* can be chosen during design time. Depending on these variables, different scaling factors are calculated.

The realization of multiple clock domains in a design requires a clock domain crossing between the different regions. In our approach the buffers at each input port of the router have been replaced by asynchronous FIFOs. Thus neighboring routers can run in different clock domains without any further design changes. As a result of this, clock synchronization becomes redundant.

<b>Topology</b>	8 x 8 mesh
<b>Router Architecture</b>	5 port router (4 neighbors + 1 local tile)
<b>Input Buffer</b>	4 buffers per port (due to 4 virtual channels), each 4 flits deep
<b>Link Width</b>	128 bit data, 9 bit control
<b>Processor</b>	Gaisler Leon3
<b>Tile</b>	5 processors per tile (cluster)
<b>Frequency</b>	25 MHz (FPGA prototype) 1 GHz (ASIC implementation)

Table I  
 NETWORK ARCHITECTURE DETAILS FOR PERFORMANCE AND POWER SIMULATION

## V. RESULTS AND DISCUSSION

For performance and power evaluation we utilized the architecture described in section III-B and III-C and extended it by the power management. Figure 5 shows the architecture used for simulation. The detailed parameters of the network architecture utilized for evaluation are given in table I. The proposed router design and the power management unit are implemented in *SystemVerilog*. The synthesis was done on the one hand for an FPGA based prototype using one Xilinx Virtex7 2000T FPGA. On the other hand the same HDL files are used for an ASIC synthesis and power estimation with Synopsys *DesignCompiler*. Therefore the TSMC library tcn45gsbwp\_120a is applied.

The wake-up delay times for clock and power gated circuits depend on the one hand on the used design library and on the other hand on the depth of the sleep state. To keep flexibility and not to be assigned to one specific library, the wake-up delay time of the routers is implemented as a variable parameter. Hence the designer can test its architecture design with different delay times to assess the impact of the switching delay on the network performance.<sup>1</sup>

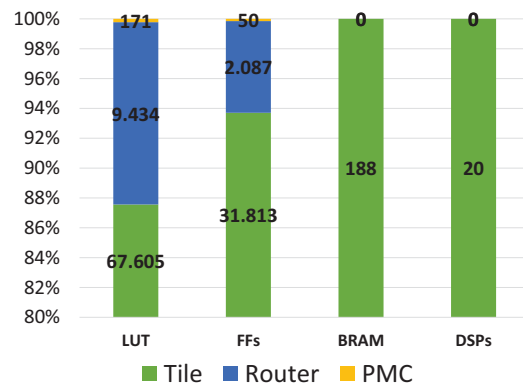


Figure 6. Resource consumption of a single tile, the original router, and the power management controller (PMC). The PMC consists of the load detection unit and the power management controller. The figure illustrates the synthesis results of a Xilinx Virtex7 2000T FPGA

<sup>1</sup>The results in this paper have been generated with a delay of 20 clock cycles.

(x,y)	0	1	2	3	4	5	6	7
0	21%	18%	16%	15%	15%	17%	22%	27%
1	19%	15%	13%	13%	14%	16%	20%	23%
2	18%	13%	11%	11%	12%	16%	18%	23%
3	18%	13%	12%	12%	12%	15%	19%	23%
4	18%	13%	13%	12%	13%	16%	18%	22%
5	19%	15%	13%	12%	12%	16%	18%	23%
6	21%	17%	16%	15%	16%	21%	22%	26%
7	30%	22%	19%	19%	17%	20%	23%	32%

Table II

PROPORTION OF TIME ROUTERS ARE POWER-GATED IN A 8x8 MESH, WITH SYNTHETIC UNIFORM RANDOM TRAFFIC AND AN INJECTION RATE OF 0.3 [FLITS/NODE/NS].

(x,y)	0	1	2	3	4	5	6	7
0	4%	4%	6%	6%	7%	8%	9%	11%
1	5%	5%	5%	6%	7%	7%	9%	9%
2	6%	6%	6%	6%	6%	7%	8%	8%
3	8%	7%	6%	6%	6%	7%	8%	8%
4	8%	8%	7%	7%	7%	6%	7%	8%
5	10%	9%	8%	7%	6%	6%	7%	7%
6	11%	10%	9%	7%	6%	6%	6%	7%
7	13%	11%	10%	8%	7%	7%	6%	7%

Table III

PROPORTION OF TIME ROUTERS ARE POWER-GATED IN A 8x8 MESH, WITH SYNTHETIC TRANSPOSE TRAFFIC AND AN INJECTION RATE OF 0.3 [FLITS/NODE/NS].

### A. Resource Consumption


The generated area overhead by adding the power management controller (comprising load detection unit and power management controller) to the router design is only 2.1%. Figure 6 shows the contrasting juxtaposition of resource consumption of a single cluster (tile), a network router without power management and the PMC. The design parameter of the tile and the router are given in table I. Resource consumptions for look up tables (LUT), Flip Flops (FFs), Block RAM (BRAM), and Digital Signal Processors block (DSPs) are listed in figure 6. The results are based on an FPGA synthesis for Xilinx Virtex7 2000T. The overhead for an ASIC synthesis has the same proportion outcome.

### B. Simulation Infrastructure

To evaluate the performance shift of the network by adding the power management, we are using a *hardware description language (HDL)* based *ModelSim* simulation of an 8x8 mesh network. The structure of the test architecture is shown in figure 5. A list of architecture details is given in table I. To reduce the simulation time, the processing elements are represent by simple models, hence only synthetic generated traffic is applied.

The proportion of power-gated routers depends particularly on the application. Having communication intensive applications, the amount of time a router can be power-gated will be much lower, than having computation intensive applications. To evaluate a worst case scenario, we are using synthetic generated uniform random traffic with varying injection rates as well as transpose traffic. Even these worst case consumptions still show that the approach by adding a small amount of resources can save power without major performance loss. The injection rates were measured in flits/node/ns. Table II shows the percentage of power-gated time of each router in a 8x8 mesh network with an injection rate of 0.3 flits/node/ns. Building the average over all injection rates between 0.1 and 0.6, in as many as 14.2% of the time a router is switched off.

Figure 7 compares the throughput and latency of a 8x8 mesh network without power management with a network of equal dimensions including the *power management controller (PMC)*, using uniform random traffic. The throughput (see figure 7(a)) is almost equal in both implementations. Due to the power on delay, packets which should be transmitted



(a) Throughput

(b) Average latency

Figure 7. Network performance, using uniform random traffic, of two and three dimensional mesh networks, each with 64 nodes and homogenous bandwidth of 128bit. Clock scaling not enabled.

to a power gated router, will be delayed. Hence the latency with low injection rates differ much more, since it is more likely that routers can be power gated. If the injection rates increases, it is less likely that routers can be power gated, due to high utilization.

Figure 8 shows the equivalent results for transpose traffic. The latency overhead including the power management is comparatively small for transpose traffic. The reason for this is the small percentage of power-gated routers as it is stated in table III. In transpose traffic, router (x,y) send packages to router (y,x). As a consequence the traffic is evenly distributed in the mesh network and routers on the margin are utilized in the same degree, as routers in the center of the mesh. To reduce the power consumption for such traffic patterns as well, where routers carry only low loads and hence cannot be power gated, the PMC suggest a clock scaling factor. A selection of fixed values are possible as frequency scaling factors. They can be chosen during design time by defining the *threshold limit value (TLV)* accordingly.

Instead of retarding packets if routers are power-gated, it would also be possible to use a different routing scheme than XY routing. For example Heisswolt et al. [22] describe a rerouting mechanism. But the more complex routing algorithms come with a resource overhead and increased latency if non minimal routes are used. Besides that, the destination router need to be waked up anyways and hence retarding packets cannot be fully prevented. For this reasons we abandoned the option of a different routing algorithm than XY routing.

Version	Area ( $\mu\text{m}^2$ )	Power (Min.)			Power (Max.)		
		Leakage power (mW)	Dynamic power (mW)	Total power (mW)	Leakage power (mW)	Dynamic power (mW)	Total power (mW)
Base Router	229944	2.30	10.81	13.1028	2.06	449.2474	451.3119
Router with PM	232565 (+1.14%)	2.13	7.29	9.4197 (-28.11%)	2.08	487.71	489.7825 (+8.52%)

Table IV  
 ASIC TSMC 45 nm POWER ESTIMATION RESULTS FOR A SINGLE ROUTER INCLUDING THE POWER MANAGEMENT  
 (FREQUENCY = 1 GHz; OPERATING VOLTAGE = 0.8V)

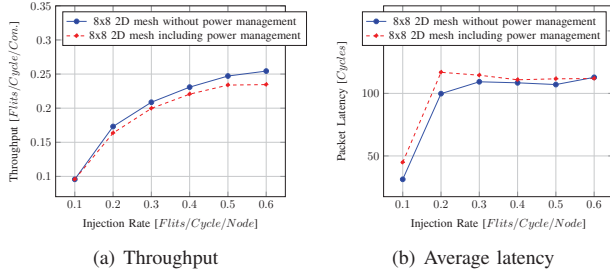


Figure 8. Network performance, using transpose traffic, of two and three dimensional mesh networks, each with 64 nodes and homogenous bandwidth of 128bit. Clock scaling not enabled.

(x,y)	0	1	2	3	4	5	6	7
0	77	69	56	55	57	60	66	71
1	66	58	51	48	53	55	61	67
2	54	56	45	43	52	55	58	64
3	55	58	45	48	51	57	60	59
4	52	52	50	49	49	57	57	59
5	62	60	54	48	52	54	63	64
6	63	57	56	52	61	55	60	66
7	82	78	69	55	59	63	70	82

Table V

POWER ON SWITCHING ACTIVITY OF ROUTERS IN 8X8 MESH, WITH SYNTHETIC UNIFORM RANDOM TRAFFIC AND AN INJECTION RATE OF 0.3 [FLITS/NODE/NS]. SIMULATION TIME: 315  $\mu\text{s}$ , 135,000 FLITS DISPATCHED IN TOTAL. [TOTAL NUMBER OF POWER-ON SWITCHES].

### C. Power Analysis

There are different methods available to measure the power dissipation of FPGAs. Jevtic et al. [23] describe a methodology to measure separate values of static, clock, interconnect, and logic power in FPGAs. However the power consumption of FPGAs differs considerably from ASIC power consumption [24]. Hence the FPGA prototype is only used for system functionality verification and not for power evaluations.

There are two parameters which are important to evaluate the power saving potential. First the span of time, a router is powered off. And second the amount of switches between power off and on state, since each switch costs additional power. Table V shows the number of on switches of each router in a 8x8 mesh network with an injection rate of 0.3 flits/node/ns using uniform random traffic.

For power estimations at RTL level, we are using Synopsys PowerCompiler included in the Synopsys tool DesignCompiler. The tool computes the leakage and dynamic power of a router based on a given test bench scenario. For generating the power estimation results, the same test bench

as described in section V-B has been used. Table IV shows the ASIC area and power consumption of a reference router without power management, compared to a router including the PMC. Thereby, the power consumption is given on the one hand for no communication (Power (Min.) columns) and on the other hand for high communication load (Power (Max.) columns).

In case of no communication, the router is inactive all the time and can be power-gated. If so, the highest power savings are possible. In case of high communication load, the router is active all the time and cannot be switched off at any time. And resulting from that, the router including the power management consumes 8.52% additional power than the implementation without power management. The PowerCompiler measurements given in table IV, only comprise clock gating instead of power gating. Hence the leakage power is not reduced in this case. The maximum power saving potential using clock gating is 28.11%, in case of uniform random traffic. If instead of clock gating, power gating is applied even higher power savings could be achieved. Since the same concept of the power management for network resources can be used for clock gating as well as for power gating, no additional implementation effort is required.

Matsutani et al. [13] introduced a look-ahead sleep control to reduce the performance penalty. In contrast, our approach can achieve the same results without additional resources.

Compared to the network described in [7] our VC power gating technique covers BE and GS communication and is also capable of switching off complete routers. In case of *guaranteed service (GS)* connection, the utilization can be zero, but switching off a router is not possible since latency guarantees can be violated otherwise. In this case power will be saved by scaling down the frequency to a minimum. Finally, depending on the actual communication an optimal energy saving measure is applied.

In addition we have a working FPGA prototype of the NoC including power management and processing elements. Although the target platform is an ASIC, the FPGA prototype is used to verify the implementation with more realistic communication patterns.

## VI. CONCLUSION

In this paper, we presented a power management controller for on chip networks. Each router comprises its own controlling unit to realize a decentralized power optimization while

minimizing the performance loss. Depending on the degree of router utilization, the controller either suggest a frequency scaling factor or power gates the complete router in case of no utilization. The resource overhead of the PMC is kept to a minimum, while saving an average of 14.2% of the network power. A FPGA based implementation of the approach is available for evaluation as well as an ASIC design flow for power simulation.

#### ACKNOWLEDGMENT

This work was supported by the German Research Foundation (DFG) as part of the Transregional Collaborative Research Center "Invasive Computing" (SFB/TR 89).

#### REFERENCES

- [1] X. Chen, Z. Xu, H. Kim, P. V. Gratz, J. Hu, M. Kishinevsky, U. Ogras, and R. Ayoub, "Dynamic voltage and frequency scaling for shared resources in multicore processor designs," in *Proceedings of the 50th Annual Design Automation Conference*. ACM, 2013, p. 114.
- [2] P. Bogdan, R. Marculescu, S. Jain, and R. T. Gavila, "An optimal control approach to power management for multi-voltage and frequency islands multiprocessor platforms under highly variable workloads," in *Networks on Chip (NoCS), 2012 Sixth IEEE/ACM International Symposium on*. IEEE, 2012, pp. 35–42.
- [3] A.-M. Rahmani, M.-H. Haghbayan, A. Kanduri, A. Y. Weldezion, P. Liljeberg, J. Plosila, A. Jantsch, and H. Tenhunen, "Dynamic power management for many-core platforms in the dark silicon era: A multi-objective control approach," in *Low Power Electronics and Design (ISLPED), 2015 IEEE/ACM International Symposium on*. IEEE, 2015, pp. 219–224.
- [4] U. Y. Ogras, R. Marculescu, D. Marculescu, and E. G. Jung, "Design and management of voltage-frequency island partitioned networks-on-chip," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 17, no. 3, pp. 330–341, 2009.
- [5] Y. Ge, P. Malani, and Q. Qiu, "Distributed task migration for thermal management in many-core systems," in *Proceedings of the 47th Design Automation Conference*. ACM, 2010, pp. 579–584.
- [6] M. Keating, D. Flynn, R. Aitken, A. Gibbons, and K. Shi, *Low power methodology manual: for system-on-chip design*. Springer Publishing Company, Incorporated, 2007.
- [7] A. Mirhosseini, M. Sadrosadati, A. Fakhrzadehgan, M. Modarressi, and H. Sarbazi-Azad, "An energy-efficient virtual channel power-gating mechanism for on-chip networks," in *Proceedings of the 2015 Design, Automation & Test in Europe Conference & Exhibition*, 2015.
- [8] J. Henkel, H. Bukhari, S. Garg, M. U. K. Khan, H. Khdr, F. Kriebel, U. Ogras, S. Parameswaran, and M. Shafique, "Dark silicon: From computation to communication," in *Proceedings of the 9th International Symposium on Networks-on-Chip*. ACM, 2015, p. 23.
- [9] R. Das, S. Narayanasamy, S. K. Satpathy, and R. G. Dreslinski, "Catnap: Energy proportional multiple network-on-chip," *SIGARCH Comput. Archit. News*, vol. 41, no. 3, pp. 320–331, Jun. 2013. [Online]. Available: <http://doi.acm.org/10.1145/2508148.2485950>
- [10] A. Canidio, "A power gating methodology to aggressively reduce leakage power in networks-on-chip buffers," 2015.
- [11] H. Bokhari, H. Javaid, M. Shafique, J. Henkel, and S. Parameswaran, "darknoc: Designing energy-efficient network-on-chip with multi-vt cells for dark silicon," in *Design Automation Conference (DAC), 2014 51st ACM/EDAC/IEEE*. IEEE, 2014, pp. 1–6.
- [12] K. Lee, S. joong Lee, S. eun Kim, H. mi Choi, D. Kim, S. Kim, M. wuk Lee, and H. jun Yoo, "A 51mw 1.6ghz on-chip network for low-power heterogeneous soc platform," in *Solid-State Circuits Conference, 2004. Digest of Technical Papers. ISSCC. 2004 IEEE International*, 2004.
- [13] H. Matsutani, M. Koibuchi, D. Wang, and H. Amano, "Run-time power gating of on-chip routers using look-ahead routing," in *Proceedings of the 2008 Asia and South Pacific Design Automation Conference*. IEEE Computer Society Press, 2008, pp. 55–60.
- [14] H. Matsutani, M. Koibuchi, D. Ikebuchi, K. Usami, H. Nakamura, and H. Amano, "Ultra fine-grained run-time power gating of on-chip routers for cmps," in *2010 Fourth ACM/IEEE International Symposium on Networks-on-Chip (NOCS)*. IEEE, pp. 61–68.
- [15] V. D. Tovinakere, O. Sentieys, and S. Derrien, "Wakeup time and wakeup energy estimation in power-gated logic clusters," in *VLSI Design (VLSI Design), 2011 24th International Conference on*. IEEE, 2011, pp. 340–345.
- [16] —, "A semiempirical model for wakeup time estimation in power-gated logic clusters," in *Proceedings of the 49th Annual Design Automation Conference*. ACM, 2012, pp. 48–55.
- [17] K. Agarwal, K. Nowka, H. Deogun, and D. Sylvester, "Power gating with multiple sleep modes," in *Proceedings of the 7th International Symposium on Quality Electronic Design*. IEEE Computer Society, 2006, pp. 633–637.
- [18] N. Banerjee, P. Vellanki, and K. S. Chatha, "A power and performance model for network-on-chip architectures," in *Design, Automation and Test in Europe Conference and Exhibition, 2004. Proceedings*, vol. 2. IEEE, 2004, pp. 1250–1255.
- [19] J. Heisswolf, A. Zaib, A. Weichslgartner, M. Karle, M. Singh, T. Wild, J. Teich, A. Herkersdorf, and J. Becker, "The invasive network on chip—a multi-objective many-core communication infrastructure," in *Architecture of Computing Systems (ARCS), 2014 27th International Conference on*. VDE, 2014, pp. 1–8.
- [20] A. Gaisler and S. Göteborg, "Leon3 multiprocessing cpu core," *Aeroflex Gaisler, February*, 2010.
- [21] J. Heisswolf, "A scalable and adaptive network on chip for many-core architectures," Ph.D. dissertation, Karlsruhe Institute of Technology, 2014. [Online]. Available: <http://digbib.ubka.uni-karlsruhe.de/volltexte/1000045305>
- [22] J. Heisswolf, M. Singh, M. Kupper, R. Konig, and J. Becker, "Rerouting: Scalable noc self-optimization by distributed hardware-based connection reallocation," in *Reconfigurable Computing and FPGAs (ReConFig), 2013 International Conference on*. IEEE, 2013, pp. 1–8.
- [23] R. Jevtic and C. Carreras, "Power measurement methodology for fpga devices," *Instrumentation and Measurement, IEEE Transactions on*, vol. 60, no. 1, pp. 237–247, 2011.
- [24] I. Kuon and J. Rose, "Measuring the gap between fpgas and asics," *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, vol. 26, no. 2, pp. 203–215, 2007.