

Energy Model of Networks-on-Chip and a Bus

Pascal T. Wolkotte*, Gerard J.M. Smit*, Nikolay Kavaldijev*, Jens E. Becker†, Jürgen Becker†

*University of Twente, Department of EEMCS, P.O. Box 217, 7500 AE Enschede, The Netherlands
P.T.Wolkotte@utwente.nl

†University of Karlsruhe, Institute for Information Processing Technology (ITIV) D76128 Karlsruhe, Germany

Abstract—A Network-on-Chip (NoC) is an energy-efficient on-chip communication architecture for Multi-Processor System-on-Chip (MPSoC) architectures. In earlier papers we proposed two Network-on-Chip architectures based on packet-switching and circuit-switching. In this paper we derive an energy model for both NoC architectures to predict their energy consumption per transported bit. Both architectures are also compared with a traditional bus architecture. The energy model is primarily needed to find a near optimal run-time mapping (from an energy point of view) of inter-process communication to NoC links.

I. INTRODUCTION

In the Smart chipS for Smart Surroundings (4S) project [1] we propose a heterogeneous Multi-Processor System-on-Chip (MPSoC) architecture with run-time software and tools. The MPSoC architecture consists of a heterogeneous set of processing tiles interconnected by a Network-on-Chip (NoC) as depicted in Figure 1. The size of a processing tile is assumed to be less than 5 mm^2 in $0.13 \mu\text{m}$ technology. By exploiting the available parallelism of the processing tiles they can run at a relatively low frequency (below 500 MHz) to achieve enough performance. The architecture including the run-time software can replace inflexible ASICs for future mobile systems.

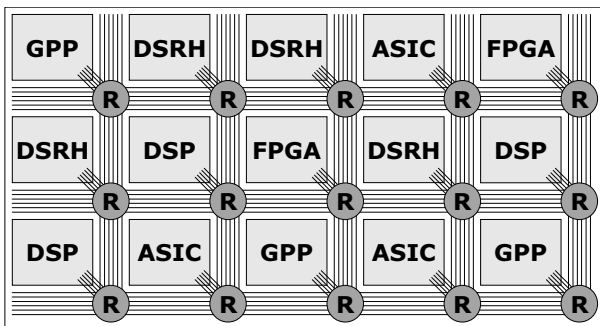


Fig. 1. An example of a heterogeneous System-on-Chip (SoC) with a Network-on-Chip (NoC). DSRH = Domain Specific Reconfigurable Hardware

Mobile systems are typically battery powered and have to support a wide range of applications so they have to be flexible as well as energy-efficient. We consider a set of streaming applications that run for a considerable period (seconds and more): e.g. wireless baseband processing (DAB, DRM, DVB), multi-media processing (MPEG-2, MPEG-4). To map these applications on a parallel architecture like a MPSoC we assume the application is represented as communicating parallel processes. One possible representation is a Kahn based process graph model [2], which is a directed graph with nodes

representing sequential processes and edges representing FIFO communication between processes.

The MPSoC architecture of the 4S-project is controlled by a central operating system called OSYRES [3], that runs on one of the GPPs of the MPSoC. The main task of OSYRES is to manage the system resources. It tries to satisfy Quality of Service (QoS) requirements, to optimize the resources usage and to minimize the energy consumption.

To reduce the energy consumption of the overall application we map the processes on the processing tile that can execute it most efficiently. This spatial mapping of processes is performed at run-time by the spatial mapping tool (SMIT) [4]. OSYRES determines when the spatial mapping tool is called. Due to the mapping of processes to processing tiles on the MPSoC communication is introduced, because data has to be moved to the successive processing tiles.

Traditionally communication between processing tiles is based on a shared bus. But for larger MPSoC with many processing tiles it is expected that the bus will become a bottleneck from both a performance, scalability and energy point of view [5]. Therefore, we propose a multi-hop Network-on-Chip, where the network consists of a set of routers interconnected by links.

In this paper we will derive a simple energy model of two Network-on-Chip architectures. This is primarily needed for the spatial mapping tool. Using this model the tool can find a near-optimal mapping (from an energy point of view) of inter-process communication to NoC links. Therefore a first-order estimation of the energy consumption is needed and sufficient. A complicated energy model would hamper the spatial mapping tool. A second motivation of deriving an energy model is that we can compare different NoC options (see also section V). We compare the energy consumption of a solution based on a packet-switched wormhole router with virtual channels, a circuit-switched router with a separate best effort network and a traditional bus.

One of the first power modeling tool was Orion, a cycle-accurate network power-performance simulator, that was proposed in [6]. The capacitance of each network component is derived based on architectural parameters, and activities at each cycle trigger calculations of network power.

The rest of this paper is organized as follows. The evaluated network routers are briefly described in section II. The energy consumption of the logic can be determined as described in section III. This power estimation of the logic does not include the long wires of the links between the routers or

wires required in a bus architecture. For the long wires of the communication architecture we use an analytical model of a wire, which is described in section IV. In section V we compare the derived energy models of the Network-on-Chip architectures with a traditional bus. In section VI we conclude the paper.

II. COMMUNICATION ARCHITECTURES

For the NoC we defined two networks (packet-switched and circuit-switched) that can both handle guaranteed throughput (GT) traffic and best-effort (BE) traffic. The guaranteed throughput traffic is defined as data streams that have a guaranteed throughput and a bounded latency. The best-effort traffic is defined as traffic where neither throughput nor latency is guaranteed. The BE traffic handles traffic like configuration data, interrupts, status messages etc.

A. Packet-Switched Network

The packet switching router implements wormhole switching with virtual channel flow control. The advantage of wormhole routing is the packet-size independent buffer-size. The virtual channels are used to decrease the chance of blocking and enables the routing of guaranteed throughput traffic.

The packet-switched router described by Kavaldjiev [7] has five input and five output ports and four virtual channels (VCs) per port. The flits (atomic unit) of a packet are labeled with their virtual channel number and they are buffered in four flit deep queues at the input ports. Per port four queues are available - one queue per virtual channel.

The outputs of the queues are not multiplexed per port, but directly connected to the crossbar. This is used to ease the arbitration compared to a standard wormhole router with virtual channels. The crossbar is asymmetric and has 20 inputs, one input for every queue, and five outputs that are directly connected to the router output ports.

The access to the crossbar is arbitrated by 5 round-robin arbiters - one arbiter per crossbar output. This arbitration is sufficient since a conflict can only arise when more than one queue contains flits destined to a same output port. Due to the predictable round-robin arbitration the router is able to handle guaranteed throughput traffic.

The following approach is used to implement a simple and fair VC allocator for best-effort traffic. All of the packets competing for a same output VC are tagged by the sender with a unique identifier, id. Each router has a global counter that counts permanently and whose value is distributed to all inputs. When an output VC is freed the next packet that takes it is the one whose id equals the current counter value. It may take several cycles until the counter value becomes equal to the id of one of the waiting packets. The uniqueness of the id guarantees conflict free arbitration. Since, at any time, the counter value is generally random, fairness is provided.

B. Circuit-Switched Network

The second network is a guaranteed throughput circuit-switched router [8] in combination with a separate best-effort

network [9]. By using dedicated techniques for both types of traffic (BE and GT) we can reduce the total area and power consumption.

For the moment the circuit-switched router has five bi-directional ports where one port is connected to a processing tile and four ports via a bi-directional link (16 bit wide per direction) to their neighboring circuit-switched routers. The bi-directional link between two routers consists of uni-directional lanes (e.g. four lanes in each direction). Each lane can be used by a unique data-stream. More than one lane per link increases the flexibility as in time division multiplexed systems. Furthermore, a link is not directly blocked because of one stream occupies the complete link as in the circuit-switched NoC of Wiklund [10]. Four lanes of four bits per link have been chosen to reduce the number of wires between routers, but it requires serialization of the 16 bit data items of the processing tiles. The serialization is handled by the data-converter that connects the (16 bit) tile interface to the small (4 bit) lanes.

To minimize energy consumption the circuit switching has fully separated data and control paths. Because in circuit-switching a data-packet cannot include routing information, we cannot serve best-effort traffic. The best-effort traffic is handled via a separate ring network [9] that can transport packets (16 bit data, 16 bit address) to all the processing tiles and circuit-switched routers. Via the configuration interface of the circuit-switched router a single best-effort packet can configure 1 lane. On average we can transport the reconfiguration data in less than 1 ms over the BE configuration network. This is fast enough, because the configuration of the crossbar will not change frequently due to the long-life guaranteed throughput data streams between processing tiles.

III. POWER MEASUREMENTS NETWORK ROUTERS

Benchmarking a NoC router is not a trivial task, because as far as we know no general method has been defined for on-chip networks. In this paper the power estimation of the logic is performed by modeling the design in VHDL. The synthesized VHDL-design is then annotated via a set of test-scenarios. We can estimate the power consumption per scenario using Synopsys Power Compiler [11] and the annotated design.

We expect that the power consumption of a single router is at least dependent on four parameters:

- 1) The average load of every data stream. This varies between 0% and 100% of the available bandwidth of a single lane/link.
- 2) The amount of bit-flips in the data stream. This varies from no bit-flips (ie. transmitting constant values) to continuous bit-flips.
- 3) The number of concurrent data streams through the router, which in our case has a maximum equal to the number of lanes (20).
- 4) The amount of control overhead in the router (e.g. buffers, arbitration)

#	Number of streams	Comment
1	0	The router is idle
2	1	Stream from and to other router
3	1	Stream from other router to processing tile
4	1	Stream processing tile to other router
5	2	Combination of 3 and 4
6	3	Combination of 2, 3 and 4
7	5	Combination of 5 and three times 2
8	10	Two times the number of streams of 7
9	15	Three times the number of streams of 7
10	20	All the lanes / virtual channels are occupied

TABLE I
SCENARIO DEFINITIONS

A. Used Traffic Patterns

To test the parameter sensitivity of our router we defined a set test-scenarios for traffic patterns. This set has three levels for the number of bit-flips:

- Best case (no bit-flips, transmitting only zeros)
- Worst case (continuous bit-flips)
- Typical case (random data with 50% bit-flips).

Furthermore, to vary the amount of traffic which concurrently traverse the router we defined ten scenarios. The scenarios have a variable number of concurrent data-streams with an variable load between 0% and 100%. The ten scenarios are listed in Table I.

The first scenario is a situation where no-data traverse the router during the time of the simulation. This will give the static offset in the dynamic power consumption. The other scenarios will simulate one or more concurrent data-streams. These scenarios are used to calculate the average energy consumption per bit [pJ/bit] to traverse one single router.

B. packet-switched network

For the packet-switched solution all the 10 scenarios are applied. In each scenario the data-streams use the guaranteed throughput protocol of the router. For each data-stream a header is send through the router, which reserves a virtual channel. After the reservation, the power consumption of the router is measured over 20 kB of data that is offered to the router in a variable time-interval. The variable interval is used to change the average load of the link. For every scenario, load and the amount of bit-flips we measured the power consumption per MHz [$\mu W/MHz$]. The left graph of Figure 2 depicts the dynamic power consumption depending on the offered load for typical data.

C. Circuit-switched network

For the circuit-switched solution the same 10 scenarios are applied. For each data-stream a configuration command is sent to the router, which configures the crossbar. After the configuration the power consumption is measured with the same method as used for the packet-switched network. The power consumption of the extra required best-effort network is measured with a separate testbench [9]. Due to the large

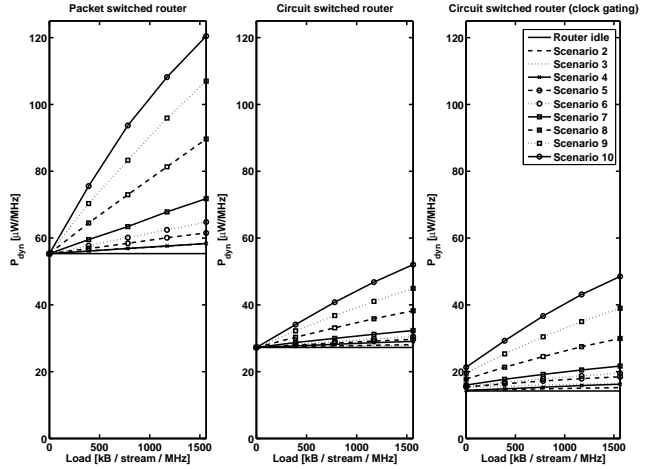


Fig. 2. Energy consumption of routers for typical data (random data with 50% bit-flips)

amount of registers the power consumption was not very load and data dependent. The power consumption of this small extra router varied between 8.4 and 12.3 $\mu W/MHz$. In this paper we use the measurement of the guaranteed throughput traffic and added the worst-case power consumption of the best-effort network to find the worst-case power consumption of the combination.

The middle graph of Figure 2 depicts the dynamic power consumption of the circuit-switched network + best-effort router depending on the offered load for typical data. We noticed a relative high offset in the dynamic power consumption. This could be reduced by including clock-gating to switch-off the inactive lanes. This resulted in the right graph of Figure 2, where the remaining offset is mainly determined by the best-effort network.

IV. POWER ANALYSIS WIRE

For the power figures of a wire we include the drivers and repeaters that are required in a link between two routing structures. In [12] the power of a link between two routers is given by:

$$P_{link} = (P_{drivers} + P_{repeaters} + P_{wire}) \cdot N_{wires} \quad (1)$$

where N_{wires} is equal to the number of parallel wires of the link. Each power factor can be defined as the sum of dynamic and leakage power. In this paper we only focus on the dynamic energy consumption. Via simulation of wires with a length less than 10 mm we discovered that for frequencies less than 100 MHz the repeaters can be safely ignored.

In [13] it is shown that the dynamic power consumption of a link (wire including the driver) is equal to:

$$P_{link_{dyn}} = \{ \alpha (s (c_p + c_0) + c \cdot l_{wire}) V_{DD}^2 f_{clk} \} \cdot N_{wires} \quad (2)$$

Where α is equal to the switching factor (or activity factor), l_{wire} the length of the wire in mm and c_0, c_p, c and s are

determined by the process, wire pitch and wire dimensions. We use $c_0 = 1.7[fF]$, $c_p = 3.5[fF]$, $c = 240[fF/mm]$ and $s = 151$, which are the values given for $0.13\mu m$ technology by [13]. For the voltage we use a $V_{DD} = 1V$, which is also used for the power estimation of the logic blocks.

The activity factor is data (the amount of bit-toggles) and load dependent. In a typical data-stream we have a 50% chance for a data change from 0 to 1 or visa versa. Therefore, for typical data-streams the activity factor is then only related to the load on the link $\alpha = 0.5 \cdot L_{link}$, where L_{link} is the average load of the link, with $0 \leq L_{link} \leq 1$. With $\alpha = 0.5 \cdot L_{link}$, $f_{clk} = 1MHz$ and $V_{DD} = 1V$ we get the power consumption in $[\mu W/MHz]$:

$$\begin{aligned} P_{link_{dyn}} &= 0.5 \cdot (s(c_0 + c_s) + c \cdot l_{wire}) \cdot N_{wires} \cdot L_{link} \\ &= (0.39 + 0.12 \cdot l_{wire}) \cdot N_{wires} \cdot L_{link} \end{aligned} \quad (3)$$

In the next sections we use the energy that is required to transport a single bit over a wire. In these cases the N_{wires} and L_{link} are both equal to 1.

V. COMPARING COMMUNICATION ARCHITECTURES

In this section we compare a bus based system with the two described networks.

A. Energy Consumption Model Packet-Switched Router

In Figure 2a we see a high offset in the dynamic power consumption of $55.34 \mu W/MHz$. Above the offset an almost linear dependency between the load of the streams, number of streams and the power consumption of the router is visible. From this linear dependency (slope of the lines) we can derive the amount of energy required for a single bit to pass the router. This is equal to 0.9776 pJ/bit.

The energy consumed by the router can be added to the energy consumption of the wire of equation 3. The dynamic energy (E_{ps} in [pJ/bit]) required to transport a bit between two processing tiles over a distance N_{hop} is equal to:

$$E_{ps} = 0.98 \cdot N_{hop} + (0.39 + 0.12 \cdot l_{wire}) \cdot (N_{hop} - 1) \quad (4)$$

B. Performance model circuit-switched router

In Figure 2b we see a relative lower offset in the dynamic power consumption of $27.3 \mu W/MHz$. Above the offset a linear dependency between the load of the streams, number of streams and the power consumption of the router is visible. From this linear dependency we can derive the amount of energy required for a single bit to pass the router. This is equal to 0.3722 pJ/bit, which in combination with equation 3 results the dynamic energy (E_{cs} in [pJ/bit]) to transport a bit between two processing tiles with a distance of N_{hop} :

$$E_{cs} = 0.37 \cdot N_{hop} + (0.39 + 0.12 \cdot l_{wire}) \cdot (N_{hop} - 1) \quad (5)$$

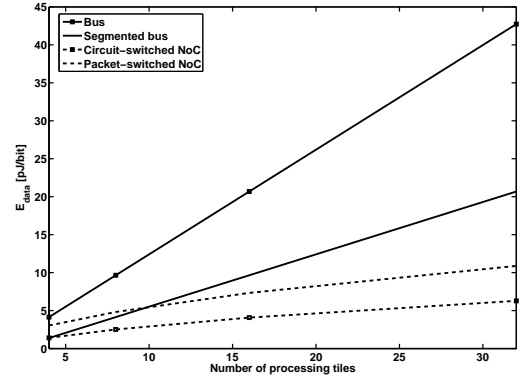


Fig. 3. Energy required for on-chip communication. ($l_{wire} = 2mm$)

C. Performance model bus

To derive the communication energy required in a (non-tristate) bus we use the analysis used in [14]. It is assumed that the bus system is organized as a regular grid of $N \times N$ processing tiles. In a single master bus system it is assumed that all slave-ports have to switch, which results that the data has to be transported over all wire segments. The minimum number of wire segments to connect all the N^2 processing tiles is equal to $N^2 - 1$.

The total amount of switching energy then equals:

$$E_{bus} = (N_{wires}/N_{data}) \cdot E_{wire}(N^2 - 1) \quad (6)$$

Where N_{wires}/N_{data} is the ratio between the number of data lines and the total number of wires (address, data, read, valid and accept flags) of the bus. For a 16 bit data and 16 bit address this ratio is equal to 2.19. Replacing E_{wire} with the energy per bit using equation 3 it results in the energy required to transport one single data-bit:

$$E_{bus} = 2.19 \cdot (0.39 + 0.12 \cdot l_{wire}) \cdot (N^2 - 1)[pJ/bit] \quad (7)$$

D. Comparison

In section V-A and V-B we derived the amount of energy to transport a single bit between processing tiles over a network-on-chip. This bit can be used as an address or data bit by the processing tiles. To make a fair comparison between the networks-on-chip and the bus we assume that 50% of the bits are used for address-bits. The energy required to transport this data bit is therefore twice the energy described by equations 4 and 5.

Using the equation 7 and the compensated equations 4 and 5 we compare the average dynamic energy required to transport a data bit between 2 processing tiles. We assume a regular grid of $N \times N$ processing tiles with a size of $4 mm^2$ each. This will result in a wire segment length (l_{wire}) equal to 2 mm. The average number of hops in a network-on-chip communication architecture depends on the distribution of the traffic. For uniform distributed traffic $\bar{N}_{hop} = \frac{2}{3}N$. More local oriented traffic will decrease the average number of hops.

Figure 3 depicts the average required energy per bit depending on the number of tiles in the MPSoC. For the bus we added an extra line, which models a segmented bus structure with 2 equally sized segments. It is assumed that this will half the number of wire segments that are used in a bus-transfer. The benefit of the Network-on-Chip is clearly visible for larger number of tiles.

VI. CONCLUSION

In this paper we presented two Network-on-Chip architectures that are compared with a traditional bus architecture. They are compared on the required energy to transport a bit over the communication architecture. For each architecture we derived a simple energy model that can be used for the spatial mapping tool to optimally map the on-chip communication streams. The energy model for all the architectures are relatively simple due to the derived first-order equations. This eases the computational requirements of the spatial mapping tool to calculate the communication costs.

The energy models showed a lower energy consumption per bit for the Network-on-Chip architectures. Especially for larger number of processing tiles the Network-on-Chip architectures consume less energy per bit. The circuit-switched network is the most energy efficient solution due to the small amount of control and buffering.

For the circuit-switched router a clock-gated implementation was also evaluated. The clock-gated design disabled the clock for in-active (not configured) lanes. The implementation showed a relative large decrease of the offset in dynamic power consumption. For the best-effort network and the packet-switched router clock-gating still has to be implemented.

ACKNOWLEDGEMENT

This research is conducted within the Smart Chips for Smart Surroundings project (IST-001908) supported by the Sixth Framework Programme of the European Community. The Software described in this document is furnished under a license from Synopsys (Northern Europe) Limited. Synopsys and Synopsys product names described herein are trademarks of Synopsys, Inc.

REFERENCES

- [1] <http://www.smart-chips.com>.
- [2] G. Kahn, "The semantics of a simple language for parallel programming," in *Information processing*, J. L. Rosenfeld, Ed. Stockholm, Sweden: North Holland, Amsterdam, Aug 1974, pp. 471–475.
- [3] "Osyres, operating framework for re-configurable embedded systems," <http://www.ti-wmc.nl>.
- [4] L. T. Smit, *et al.*, "Run-time mapping of applications to a heterogeneous reconfigurable tiled system on chip architecture," in *Proceedings of the International Conference on Field-Programmable Technology*, December 2004.
- [5] L. Benini and G. de Micheli, "Networks on chips: A new soc paradigm," *IEEE Computer*, vol. 35, no. 1, pp. 70–78, January 2002.
- [6] H. Wang, *et al.*, "Orion: A power-performance simulator for interconnection networks," in *Proceedings of MICRO 35*, Istanbul, Turkey, November 2002.
- [7] N. Kavalajiev, G. J. M. Smit, and P. G. Jansen, "A virtual channel router for on-chip networks," in *Proceedings of IEEE International SOC Conference*. IEEE Computer Society Press, September 2004, pp. 289–293.
- [8] P. T. Wolkotte, *et al.*, "An energy-efficient reconfigurable circuit-switched network-on-chip," in *Proceedings of the 12th Reconfigurable Architectures Workshop (RAW 2005)*, Denver, Colorado, USA, April 4-5 2005.
- [9] P. T. Wolkotte, G. J. Smit, and J. E. Becker, "Energy-efficient noc for best-effort communication," in *Accepted for 15th International Conference on Field Programmable Logic and Applications 2005 (FPL 2005)*, Tampere, Finland, August 24-28 2005.
- [10] D. Wiklund and D. Liu, "Socbus: Switched network on chip for hard real time systems," in *Proceedings of the International Parallel and Distributed Processing Symposium (IPDPS)*, Nice, France, April 2003.
- [11] <http://www.synopsys.com>.
- [12] A. Morgenshtein, *et al.*, "Comparative analysis of serial vs. parallel links in networks on chip," in *Proceedings SOC 2004 International Symposium on System-on-Chip*. Tampere, Finland: IEEE Computer Society Press, Los Alamitos, California, November 2004, ISBN 0-7803-8558-6.
- [13] K. Banerjee and A. Mehrotra, "A power-optimal repeater insertion methodology for global interconnects in nanometer designs," *IEEE Transactions on Electron Devices*, vol. 49, no. 11, pp. 2001–2007, November 2002.
- [14] "Private communication."