An Energy-Efficient Reconfigurable Circuit-Switched Network-on-Chip

Pascal T. Wolkotte, Gerard J.M. Smit, Gerard K. Rauwerda, Lodewijk T. Smit
University of Twente, Department of EEMCS
P.O. Box 217, 7500 AE Enschede, The Netherlands
P.T.Wolkotte@utwente.nl

Abstract

Network-on-Chip (NoC) is an energy-efficient on-chip communication architecture for multi-tile System-on-Chip (SoC) architectures. The SoC architecture, including its run-time software, can replace inflexible ASICs for future ambient systems. These ambient systems have to be flexible as well as energy-efficient. To find an energy-efficient solution for the communication network we analyze three wireless applications. Based on their communication requirements we observe that revisiting of the circuit switching techniques is beneficial. In this paper we propose a new energy-efficient reconfigurable circuit-switched Network-on-Chip. By physically separating the concurrent data streams we reduce the overall energy consumption. The circuit-switched router has been synthesized and analyzed for its power consumption in 0.13 μm technology. A 5-port circuit-switched router has an area of 0.05 mm^2 and runs at 1075 MHz. The proposed architecture consumes 3.5 times less energy compared to its packet-switched equivalent.

1. Introduction

In the Smart chipS for Smart Surroundings (4S) project [1] we propose a heterogeneous multi-tile System-on-Chip (SoC) architecture with run-time software and tools. The SoC architecture contains a heterogeneous set of processing tiles interconnected by a Network-on-Chip (NoC) as depicted in Fig. 1. The run-time software determines a near optimal mapping of applications to the heterogeneous architecture at run-time. The architecture including the run-time software can replace inflexible ASICs for future ambient systems.

These ambient systems have to support wide range of applications so they have to be flexible as well as energy-efficient. The designer has to partition the application into a Kahn like process graph model. In this model the application is represented as a graph with communicating functional processes (see for example Fig. 2). At run time, the individual processes of the application will be mapped on the tiles that can execute it most efficiently. The communication channels between processes are mapped on the NoC architecture.

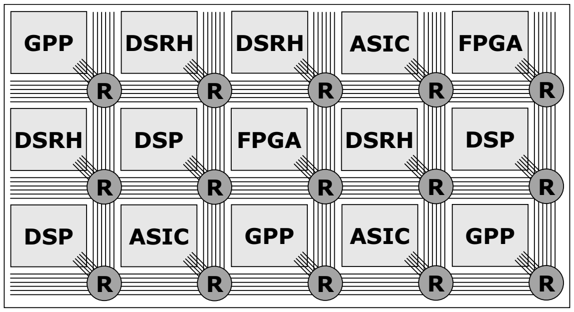

Figure 1. An example of a heterogeneous System-on-Chip (SoC) with a Network-on-Chip (NoC). DSRH = Domain Specific Reconfigurable Hardware

The multi-tile SoC architecture has many advantages: a) tiles of the same type can be duplicated when the number of transistors grow in the next technology step, b) replication of tiles eases the verification process, c) tiles do not grow in complexity with a new technology, d) relative small tiles allow extensive optimization, e) computational performance scales about linearly with the number of tiles, f) unused tiles can be switched off to reduce the energy consumption of the chip, g) locality of reference is exploited, h) it is possible to have individual clock domains per tile, and i) for reconfigurable tiles it is possible to do partial dynamic reconfiguration on a per tile basis.

Current wireless applications are based on a large set of quick evolving 3G/4G wireless standards. At design-time, the reconfigurability of the chip enables adaptation of the application in case of changes in the standards. At run-time, the SoC can be reconfigured for adaptation of the algorithms/parameters due to changes in the reception quality [2]. Furthermore, the reconfigurable SoC can share its resources among several standards (e.g. WLAN in combi-

0-7695-2312-9/05/$20.00 (c) 2005 IEEE

nation with UMTS). The reconfigurability and parallelism of the SoC provide the necessary options for a so called multi-mode transceiver system.

1.1. SoC Architecture

Our System-on-Chip consists of a heterogeneous set of processors connected via a Network-on-Chip as depicted in Fig. 1. The network consists of a set of routers interconnected by links. In this paper we assume a regular two dimensional mesh topology of the routers. Every router is connected with its four neighboring routers via bidirectional point-to-point links and with a single processor tile via the tile interface. The SoC system is organized as a centralized system: one node, called Central Coordination Node (CCN), performs system coordination functions.

The main task of the CCN is to manage the system resources. It performs run-time mapping of the newly arrived applications to suitable processing tiles and inter-processing communications to a concatenation of network links [3]. It also tries to satisfy Quality of Service (QoS) requirements, to optimize the resources usage and to minimize the energy consumption. The CCN does not perform run-time scheduling of individual processes and communications during execution. That is performed by the individual tiles and network routers. The CCN performs the feasibility analysis, spatial mapping, process allocation and configuration of the tiles and the NoC before the start of an application.

It is expected that the on-chip communication networks of these future SoC will be one of the limiting factors for performance and possibly energy consumption [4]. In this paper we describe a new architectural concept for on-chip communications. Using the communication characteristics of three wireless communication standards we propose a new reconfigurable circuit-switched on-chip network. This network benefits from the common characteristics of these wireless standards.

1.2. Organization of the Paper

The paper is organized as follows. We start with related work for on-chip communication architectures. Section 3 describes three wireless standards (HiperLAN/2, UMTS and Digital Radio Mondiale) and determines their common communication characteristics. Section 4 gives the reasons for reconsidering a circuit-switched network. Using the characteristics of other NoC solutions and the requirements of the wireless applications a new NoC architecture is developed. Section 5 describes this architecture. Section 6 describes several scenarios for benchmark the power consumption of our architecture. The paper concludes with a comparison of our router with an equivalent packet-switched router.

2. Related work

Network-on-chip (NoC) architectures [5–10] have been proposed as a solution for the problem of on-chip communication in multi-tile SoC architectures. The architectures are presented as replacements of the on-chip time-division multiplex buses (e.g. the AMBA bus from ARM Inc. [11]).

All the proposed solutions are based on routers interconnected through network links. The solutions differentiate in the topology of the network and the implementation of the individual routers. The two dimensional mesh is the most common topology compared with other topologies such as hexagons, butterflies, tree or hypercube structures. The implementations for the routers vary widely using techniques of packet or circuit switching, dynamic or static scheduling, wormhole or virtual-cut through routing.

The majority of the current router implementations for network-on-chip are based on a packet-switched, synchronous networks [5–9]. Using known routing protocols the number of buffers are minimized and best-effort traffic can be served. In the circuit-switched solution [10] buffering is not necessary. To handle guaranteed throughput traffic several techniques are used, such as: contention free routing [5], static scheduling [10], virtual channels [6], virtual circuits [8] and priorities [9].

The routers are benchmarked using a local area network approach where the benchmarks use random traffic patterns. New (more specific) NoC benchmarks can be necessary, because the on-chip traffic patterns have other characteristics [12] and demands [13].

3. Application domain

To determine the requirements of the on-chip network we have investigated the common characteristics of three wireless applications: the baseband processing of HiperLAN/2, UMTS and Digital Radio Mondiale (DRM). The block diagram of DRM is similar to HiperLAN/2, but the communication requirements are a factor 1000 less compared to HiperLAN/2. The exact figures are presented in [14].

3.1. HiperLAN/2

WLAN networks use radio technologies such as IEEE 802.11a or HiperLAN/2 to provide secure, reliable, fast wireless connectivity. They operate in the unlicensed 2.4 and 5 GHz radio bands, with data rates up to 54 Mbps. The physical layer of HiperLAN/2 is described in [15]. The task of the physical layer in HiperLAN/2 is to modulate bits that originate from the data link control layer on the transmitter side and to demodulate them on the receiver side. For modulation it uses Orthogonal Frequency Division Multiplexing

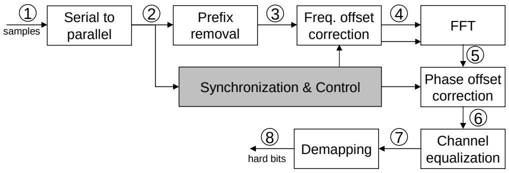

Figure 2. HiperLAN/2 Baseband Processing

Edge(s)	Stream	Bandwidth [Mbit/s]
S/P → Pre-fix removal	1-2	640
Pre-fix removal → FFT	3-4	512
FFT → Channel eq.	5-6	416
Channel eq. → De-map	7	384
Hard bits	8	12 (BPSK) up to 72 (QAM-64)

Table 1. Communication in HiperLAN/2

(OFDM). The physical layer of the HiperLAN/2 standard has been mapped on a multi-tile architecture [16]. Fig. 2 depicts the block diagram of this implementation.

One major property of the signaling in HiperLAN/2 is the grouping of samples in OFDM symbols. An OFDM symbol has a fixed length of 80 samples. All operations in the physical layer are performed on these OFDM symbols. The processing in the OFDM receiver is performed in block-mode on OFDM symbols. This results in a block-based communication stream between the successive processing tiles. One should guarantee that each 4 μs a new OFDM symbol can be processed. This requires a guaranteed throughput service for our NoC. Table 1 gives the required communication bandwidth, based on 16 bits quantization for the baseband processing.

3.2. UMTS

The Universal Mobile Telecommunications System (UMTS) standard [17] is an example of a Third Generation (3G) mobile communication system. UMTS is based on Wideband Code Division Multiple Access (W-CDMA). In CDMA every transmitted bit is coded with a spreading code of a higher rate, which means that every transmitted bit is multiplied with a spreading code. The spreading code consists of a sequence of so-called chips, whose rate is referred to as the chip rate. In this way the information is transmitted at the chip rate, so the spectrum is spread. As a consequence, many correlations with the spreading code have to be performed in the UMTS receiver.

Rauwerda [18] mapped the downlink of an UMTS W-CDMA receiver on a set of (reconfigurable) processors. In contrast to the DRM and HiperLAN/2 implementation the data processing and communication between the processors is streaming oriented. This streaming communication results in new requirement for our NoC. At a regular short interval a very small packet, containing 1 sample, has to be transported to the successive processor.

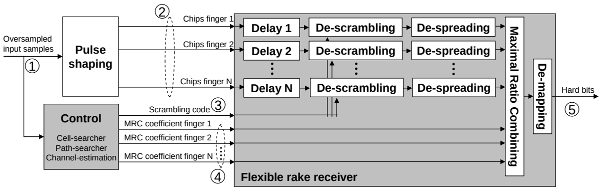

Figure 3. A UMTS Receiver W-CDMA with N RAKE Fingers

Edge	#	Bandwidth [Mbit/s]
Chips (per finger)	2	61.44
Scrambling code	3	7.68
MRC coefficient (per finger)	4	61.44/SF
Received bits	5	7.68/SF (QPSK) 15.36/SF (QAM-16)

Table 2. Communication in UMTS

Fig. 3 depicts the basic block diagram of the implemented W-CDMA receiver. The properties of the communication streams between the processes are listed in Table 2, where every chip or coefficient is represented by 8 bits. For example, the total communication bandwidth for processing 4 RAKE fingers with a spreading factor (SF) of 4 is $\sim$ 320 Mbit/s.

3.3. Common Characteristics

Analyzing the common characteristics of the three wireless applications we made the following observations:

- The applications have a fixed amount of processing per data sample (manifest loops). This is typically found in wireless baseband processing and audio and video filtering. It does not apply for audio and video compression, where the amount of processing is data dependent (non-manifest).
- Input data arrives at a fixed rate, which cause periodic data transfers between the successive processing blocks.
- We have semi-static life-time of a communication stream, which means a stream is fixed for a relatively long time.

To describe the network traffic in a system, we adopt the notation used in [19]. According to the type of services required, the following types of traffic can be distinguished in the network:

- *GT (guaranteed throughput)* this is the part of the traffic for which the network has to give real-time guarantees (i.e. guaranteed bandwidth, bounded latency).
- *BE (best effort)* this is the part of the traffic for which the network guarantees only fairness but does not give any bandwidth and timing guarantees.

We observed that in the described applications the majority of the data streams through the successive processes.

This continuous flow of data needs guaranteed throughput, because the front-end is not allowed to drop data. Depending on the standard we can use block-based communication (OFDM) or have to use streaming (UMTS) because the blocks get too large. If we compare the required bandwidth between the processes of the different applications this varies widely from several kbit/s (DRM) up to more than 0.5 Gbit/s (HiperLAN/2).

Beside the main-stream of the communication we foresee a minor part (assumed to be less then 5%) of best effort communication e.g. control, interrupts and configuration data. This communication can have more relaxed requirements for the network and hence can use the best effort services.

The overall important characteristic is the life-time of a communication stream. We aim to develop a SoC for a multimedia terminal where we can assume that the data streams are semi-static and have periodic behavior. This means that for a long period of time subsequent data items of a stream follow the same route. This will last for seconds and more, because a user will listen to its radio or has a phone conversation for a considerable time. However, the control system might change some settings of processes due to changing environmental conditions.

4. Circuit-Switched Revisited

The application analysis of section 3 shows that the amount of expected guaranteed throughput traffic is much larger compared to the amount of best effort traffic. Fundamentally, sharing resources and giving guarantees are conflicting, and efficiently combining guaranteed traffic with best-effort traffic is hard [20]. By using dedicated techniques for both types of traffic we aim for reducing the total area and power consumption. In this paper we concentrate on the architecture for communication of guaranteed throughput traffic.

For guaranteed throughput we use reconfigurable circuit switching that create dedicated connections between two processing tiles. The reasons for reconsidering circuit switching instead of using packet switching are:

- The flexibility of packet switching is not needed, because data streams are fixed for a relatively long time. Therefore, a connection between two tiles is required for a long period (e.g. seconds or longer). This connection can be configured by the CCN.
- Large amount of the traffic between tiles will need a guaranteed throughput, which can be easier guaranteed in a circuit-switched connection.
- Current SoC have a large amount of wiring resources that give enough flexibility for streams with different bandwidth demands.
- Circuit switching eases the implementation of asynchronous communication techniques, because data and control can be separated. A control free pipelined asynchronous data stream does not require much design effort.
- The circuit switching has a minimal amount of control in the data pad (e.g. no arbitration). This increases the energy- efficiency per transported bit and the maximum throughput.

Further, we see some benefits when guaranteed throughput traffic has to be scheduled:

- Scheduling communication streams over non-time multiplexed channels is easier, because by definition a stream will not have collisions with other communication streams. The Æthereal [5] and SoCBUS [10] routers have large interaction between data streams (both have to guarantee contention free paths). Determining the static time slots table requires considerable effort.
- Because data-streams are physically separated, collisions in the crossbar do not occur. Therefore, we do not need buffering and arbitration in the individual router. An established physical channel can always be used.

5. Architecture

As described in the introduction, we have a hardware architecture containing a heterogeneous set of processing tiles interconnected by a Network-on-Chip (NoC). For the moment the tiles and NoC are synchronized by the same clock. The NoC will use dedicated techniques for BE and GT traffic. For the GT traffic we developed a reconfigurable circuit-switched network, that is configured by the CCN. For BE traffic we aim for a packet-switched solution.

5.1. Proposed Circuit-Switched Network

The task of a NoC is to transport data from one tile to the other. The NoC consists of a set of routers interconnected by links. Using the benefits of the circuit switching and required application demands we defined a reconfigurable circuit-switched router (see Fig. 4).

The reconfigurable circuit-switched router consists of three major parts: the data-converter, crossbar and the crossbar configuration. For the moment the router has five bidirectional 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.

Wiklund [10] concluded that their circuit-switched NoC had a high latency for setting up a new circuit. This was caused by the blocking of the routers, because a reserved physical link could not be used for other connections. They

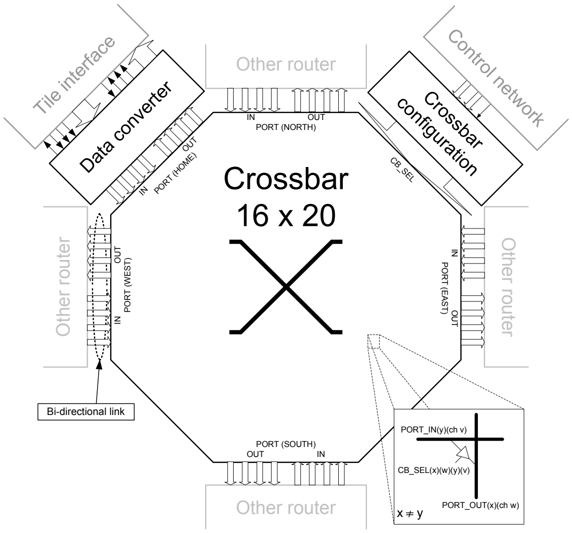

Figure 4. Block Diagram of the Router

solved this problem by suggesting a static timing schedule for the different occurring data-streams.

This problem is solved in our proposal by creating small channels (e.g. four bits) called lanes. The bi-directional link between two routers consists of a concatenation of uni-directional lanes (e.g. two times four lanes) as depicted in the right-side of Fig. 5. This increases the flexibility as is available in time division multiplexed systems. The small lanes are connected to a tile interface via the data-converter. Fig. 5 depicts a data-converter that converts the 16 bit data to the width of the lanes and visa-versa. The 16 bit tile interface is compatible with the packet-switched alternative of Kavaldjiev [6]. We define the division in lanes as lane division multiplexing. Each lane can be used for a different physical connection between processors.

The width and number of lanes are adjustable parameters in the design. They can be adjusted at design-time of the SoC to meet the flexibility and bandwidth requirements of the aimed applications. For example, if more streams are needed for the north and south port their number of lanes can be increased. The tables of section 3 can be used to determine the width and number of lanes. Four lanes of four bits and a tile interface of 16 bits have been chosen to make a fair comparison with the four virtual channel configuration of the packet-switched alternative.

In the router the four lanes of one port have to be connected with all the four lanes of all the other four ports. This results in a router with 20 input and 20 output lanes. They are connected via a 16x20 fully connected crossbar (20x20 is not necessary, because data does not have to flow back). The 20 output lanes of the crossbar are registered. The speed of the total network will therefore only depend on the maximum delay in a single router plus the maximum wire delay of the link between two routers.

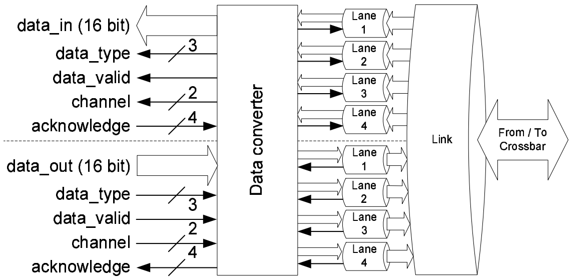

Figure 5. Data Converter Between Tile Interface and Crossbar

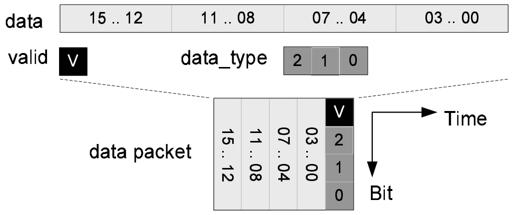

Figure 6. Organization of the Header and Data

The configuration of the crossbar (which input lane is connected with which output lane) is stored in a local configuration memory. Per output lane the input lane is stored plus an activation bit. The configuration memory size is 5x20 = 100 bits.

To minimize energy consumption the circuit switching has fully separated data and control paths. Because a data-packet cannot include routing information, we cannot serve best effort traffic. We configure the configuration memory via a small additional interface. Configuration of 1 lane requires 10 bits that are generated by the CCN based on the optimal applications mapping [3]. The configuration interface is connected to the separate BE network. We aim to transport the reconfiguration data in less than 1 ms over the BE network, because the configuration of the crossbar will not change frequently due to the long-life data streams between tiles. One single router can than be fully reconfigured within 20 ms.

5.2. Data and Flow Control

The circuit-switched network can handle synchronization of information in the data-packets. To enable this we included a small four bits header with every data-word. The header is combined with a 16-bit data-word of the tile. The result is a packet of 5x4 bits, which can be transported over a lane. The organization of this 20 bit packet is given in Fig. 6.

With only a four bit forward lane from source to destination and no feedback, we have to assume the destination can consume the data. In this case we do not support end-

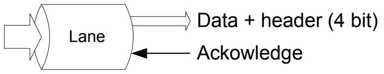

Figure 7. Details of a Single Lane

to-end flow control. To overcome this problem an acknowledgement signal is added in the reverse direction. The new lane consists of 4 data signals and 1 acknowledge signal in the reverse direction (see Fig. 7). Depending on the application one or more lanes and zero, one or more acknowledgements signals can be used.

The acknowledgement signal is used in combination with a window counter mechanism. This mechanism will prevent a buffer overflow at the destination of a connection. Every source has a local window counter of size WC. This local window counter indicates how many data-packet the source is allowed to send to the destination. The destination will send an acknowledgement signal when it has read X data-packets, where $X \leq WC$. When the source receives an acknowledge signal it increases its local window counter (WC) by X. By configuring the use of the acknowledgement signal and size of X and WC we can support both blocking and non-blocking communication.

6. Traffic Patterns

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

1. The average load of every data stream. This varies between 0% and 100% of the available bandwidth of a single lane.
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).

6.1. Used Traffic Patterns

To test the parameter sensitivity of our router we defined a test set 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 four scenarios. The scenarios are a combination of concurrent data-streams that are listed in Table 3. All three data streams have a load of 100%.

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. Scenario two, three and four are depicted in Fig. 8. Scenario II simulates the communication between the tile interface and a link. Scenario III extends Scenario II with communication between a link and the tile interface. Scenario IV also simulates a data stream that passes the router. This scenario also gives an indication of the difference between time and lane multiplexing, because two streams will be routed to the output port East.

Stream	Input port	Output port
1	Tile	Router (East)
2	Router (North)	Tile
3	Router (West)	Router (East)

Table 3. Stream Definitions

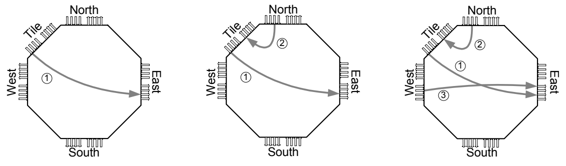

(a) Scenario II (b) Scenario III (c) Scenario IV

Figure 8. Examples of Three Test Scenario's

7. Results

The proposed circuit-switched router of this paper has been compared with a packet-switched equivalent of Kavaldjiev [6]. Both routers have bi-directional links of 16 bits. At the same frequency they have the same maximum bandwidth and bounded latency for guaranteed throughput traffic.

7.1. Synthesis Results

Both routers are synthesized in the same 0.13 μm technology. We used a TSMC low voltage, nominal VT (TCB013LVHP) standard cell library with dielectric constant (Low-k) insulators. Table 4 gives the synthesis results. Furthermore, the last column of Table 4 includes the synthesized and layouted results of the Æthereal router [5].

7.2. Power Estimation

We estimated the power consumption of our circuit-switched router and the packet-switched router of Kavaldjiev using the Synopsys Power Compiler [21]. Power Compiler distinguishes between two types of power consumption: static and dynamic. The static power consumption is the power dissipated by a gate when it is not switching. The dynamic power is the power dissipated when the circuit is active. The dynamic power is composed of two kinds

Router	Circuit switched	Packet switched	Æthereal [5]
Ports Width of data	5 16 bit	5 16 bit	6 32 bit
Crossbar	0.0258 mm^2	0.0706 mm^2	n.a.
Buffering	-	0.1034 mm^2	n.a.
Arbitration	-	0.0022 mm^2	n.a.
Configuration	0.0090 mm^2	-	-
Data converter	0.0158 mm^2	-	-
Misc	-	0.0038 mm^2	n.a.
Total	0.0506 mm^2	0.1800 mm^2	0.1750 mm^2
Max freq. Bandwidth/link	1075 MHz 17.2 Gb/s	507 MHz 8.1 Gb/s	500 MHz 16 Gb/s

Table 4. Synthesis Results of Three Routers

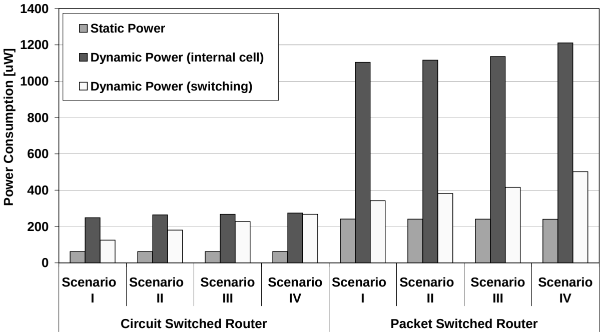

Figure 9. Dynamic and Static Power Bars for Different Scenarios (random data, 100% load)

of contributions: switching and internal cell. The switching power of a driving cell is the power dissipated by the charging and discharging of the load capacitance at the output of the cell. The internal cell power is any power dissipated within the boundary of a cell.

Fig. 9 depicts the contribution of the three power types for different traffic scenarios as described in Section 6. The clock frequency of both routers is fixed at 25 MHz. This gives a data-bandwidth of 80 Mbit/s per stream. These streams meet the required bandwidth of an edge 2 of figure 3, which is 61.44 Mbit/s. It uses random data (50% bit-flips) for every data stream as input. This 50% bit-flips is also noticed in for example edge 2 of the UMTS receiver. The simulation time is 200 μs in which 2 kB of data is transported per stream.

Fig. 10 depicts the data dependency of the dynamic power consumption. We vary the amount of bit-flips in the offered data stream. We use three cases: best-case (0%), worst-case (100%) and typical case (50%).

7.3. Discussion

Analyzing the results of the synthesis and power estimation of the circuit and packet-switched routers we made several observations:

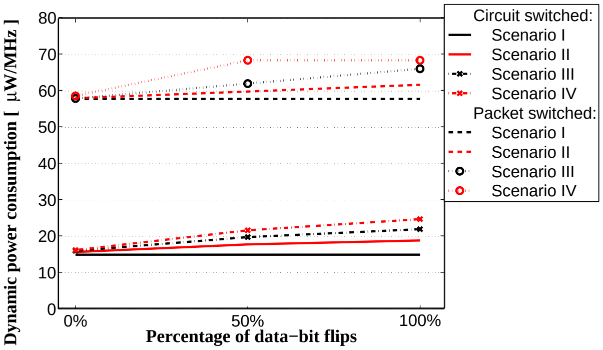

Figure 10. Data Dependency of the Dynamic Power Consumption (100% load)

- The area and power consumption of the circuit-switched router is 3.5 times less compared to the packet-switched router. The main reasons for this difference are the necessary buffers and extra control in the crossbar of the packet-switched router.
- Table 4 shows that the maximum bandwidth of both routers can meet the required bandwidth of the wireless applications as has been described in section 3.
- In contradiction to our expectations the number of bit-flips has only a minor influence on the dynamic power consumption. We can also expect that a large set of data streams will have similar toggle behavior as the used random data set. A more relevant parameter is the number of data streams that have to be routed concurrently from the input lanes to the output lanes.
- The dynamic power consumption of scenario II up to IV does not increase considerably compared with Scenario I (no transport of data). This is caused by the relative high offset in the dynamic power consumption. For the circuit-switched router we expect to reduce this power consumption with clock gating. For clock gating we can use the configuration information of the router and switch off the unused lanes. If clock gating is used, we expect that this offset will decrease. The lower offset will cause more variations in the power consumption due to variations in the traffic patterns.
- The last observation is the non-straight line in Scenario III for the packet-switched router in Fig. 10. This is caused by the collision of stream 1 and 3 at the output port East. Apparently, the time-multiplexing of the link will cause extra switching behavior in the control of the crossbar and therefore increase the power consumption.

8. Conclusion and Future Work

In this paper we proposed a new reconfigurable circuit-switched router for a Network-on-Chip. The circuit-switched Network-on-Chip was expected to benefit from the common characteristics of three wireless applications. We showed that the circuit-switched NoC satisfies the requirements for the three different wireless applications.

Compared with a packet-switched equivalent for the guaranteed throughput traffic the circuit-switched router indeed has: Lower power consumption, a smaller chip area and higher maximum throughput per direction. But it has less flexibility and no support for best effort traffic.

Further analysis of the results showed that the number of data streams is more important for the dynamic power consumption than the number of bit-flips in the data stream. To reduce large offset in the dynamic power consumption the implementation of clock gating is useful.

For future implementations of the circuit-switched router we plan to include clock gating schemes for reducing the high offset in the dynamic power consumption. Furthermore, we want to define a generic tile interface so the router can be embedded in a multi-tile SoC. This interface will support several types of communication that can be used by the application designers.

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] L. T. Smit, "Energy-efficient wireless communication," PhD thesis, University of Twente, December 2003.

[3] 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.

[4] L. Benini and G. de Micheli, "Networks on chips: A new soc paradigm," *IEEE Computer*, vol. 35, no. 1, pp. 70–78, January 2002.

[5] J. Dielissen, *et al.*, "Concepts and implementation of the Philips network-on-chip," in *IP-Based SOC Design*, Nov. 2003.

[6] N. Kavaldjiev, 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.

[7] M. Taylor, *et al.*, "The raw microprocessor: A computational fabric for software circuits and general purpose programs," *IEEE Micro*, vol. 22, no. 2, pp. 25–35, 2002.

[8] M. Millberg, *et al.*, "Guaranteed bandwidth using looped containers in temporally disjoint networks within the nostrum network on chip," in *Proceedings of the Design Automation and Test Europe Conference (DATE)*, February 2004.

[9] E. Bolotin, *et al.*, "Qnoc: Qos architecture and design process for network on chip," *The Journal of Systems Architecture*, December 2003, special issue on Networks on Chip.

[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] *AMBA bus description*, ARM, Inc., http://www.arm.com.

[12] J. Henkel, W. Wolf, and S. Chakradhar, "On-chip networks: A scalable communication-centric embedded system design paradigm," in *Proceedings of the 17th International Conference on VLSI Design*. IEEE, 2004.

[13] K. Goossens, *et al.*, "Interconnect and memory organization in SOCs for advanced set-top boxes and TV — evolution, analysis, and trends," in *Interconnect-Centric Design for Advanced SoC and NoC*, J. Nurmi, *et al.*, Eds. Kluwer, Apr. 2004, ch. 15, pp. 399–423.

[14] P. T. Wolkotte, G. J. M. Smit, and L. T. Smit, "Partitioning of a drm receiver," in *Proceedings of the 9th International OFDM-Workshop*, Dresden, Germany, September 2004, pp. 299–304.

[15] *Broadband Radio Access Networks (BRAN); Hiperlan type 2*, ETSI TS 101 475 v1.2.2 (2001-02) ed., European Telecommunication Standard Institute (ETSI), February 2001.

[16] P. M. Heysters, G. K. Rauwerda, and G. J. M. Smit, "Implementation of a HiperLAN/2 receiver on the reconfigurable montium architecture," in *Proceedings of the 11th Reconfigurable Architectures Workshop (RAW 2004)*, Santa Fé, New Mexico, USA, April 26-27 2004, iSBN 0-7695-2132-0.

[17] H. Holma and A. Toskala, *WCDMA for UMTS: Radio Access for Third Generation Mobile Communications*. John Wiley & Sons, 2001.

[18] G. K. Rauwerda and G. J. M. Smit, "Implementation of a flexible rake receiver in heterogeneous reconfigurable hardware," in *Proceedings of the 2004 International Conference on Field Programmable Technology (FPT)*, Brisbane, Australia, December 6-8 2004.

[19] E. Rijpkema, *et al.*, "Trade offs in the design of a router with both guaranteed and best-effort services for networks on chip," in *Proceedings of Design, Automation and Test in Europe Conference*, March 2003, pp. 350–355.

[20] J. Rexford and K. G. Shin, "Support for multiple classes of traffic in multicomputer routers," in *Proceedings of the First International Workshop on Parallel Computer Routing and Communication*. Springer-Verlag, 1994, pp. 116–130.

[21] "http://www.synopsys.com."