

A LOSSLESS COMPRESSION SYSTEM REALIZATION UTILIZING PHIT-SERIAL NETWORK-ON-CHIP PARADIGM

P. Dziurzynski[†], T. Maki[†], G. Ułacha[†], R. Stasinski[‡]

[†]Faculty of Computer Science and Information Systems, Szczecin University of Technology
Zolnierska 49, 71-210 Szczecin, Poland
e-mail: <pdziurzynski,tmaki,gulacha>@wi.ps.pl

[‡]Institute of Electronics and Telecommunications, Poznan University of Technology
Piotrowo 3A, 60-965 Poznan, Poland
e-mail: rstasinski@et.put.poznan.pl

ABSTRACT

The routing efficiency of on-chip networks remains an open problem for system engineers dealing with multi-core processors. The presented research includes consideration of phit-serial links in a Network-on-Chip (NoC) architecture. We propose measures for communication overhead for a single wormhole router embedded in a square mesh plane. As a source of data bit-stream, we use a lossless compression system based on blending predictors. According to the experimental results, the influence of flow control unit (flit) and physical control digit (phit) lengths on total transfer time is of a non-linear character. The transfer time is proportional to the flit length and inverse proportional to the phit size.

1. INTRODUCTION

As it is foreseen in The International Technology Roadmap for Semiconductors [11], System on Chips (SoCs) are likely to grow up to 4 billion transistors running at 10 GHz by 2010. In order to maintain the performance growth in the future SoCs, it is desirable to apply a new paradigm that allows interconnect scaling performance that connects functional blocks (cores or intellectual-property (IP) blocks) synthesized from high-level descriptions. The new on-chip interconnect paradigm that attracts most attention of researchers is the packet-based Network-on-Chip (NoC) approach [3]. This communication technique in distributed systems allows to decrease the contention level in a chip with a large number of cores, where traditional on-chip buses are considered as the main obstacle in the SoC efficiency. The performance of both dedicated wiring and buses have been reported as inferior to NoCs due to a worse reusability, scalability, problems of wiring latency and noise [3].

In order to obtain high efficiency of a NoC-based algorithm implementation, the algorithm is to be decomposed into a number of stages that communicate rarely. Then, it is essential to ensure that the amount of computation performed by a single core between synchronization points are sufficient to benefit from this realization. According to our earlier research [10], one of the algorithms that is beneficial when implemented using the NoC paradigm is a lossless compression system based on blending predictors.

However, a plain implementation of the NoC approach results in large chip area and significant capacitive load caused by a large number of multi-bit links with varying throughput demands and long standby periods [1] on certain links. We decided thus to utilize high speed on-chip serial interconnects.

2. LOSSLESS COMPRESSION

Due to the spatial correlation between adjacent pixels in digital pictures, it is often possible to decrease the amount of resources needed for storing such pictures by eliminating mutual information carried by neighbor pixels. The entropy value informs about minimum average number of bits needed to encode one pixel using one of static entropy encoding techniques (eg. Huffman, Golomb, arithmetic coding).

A typical lossless image compression technique can be decomposed into two separate stages: (i) modelling, where different linear and non-linear prediction methods are usually applied, and (ii) encoding of prediction errors.

An r -th order linear predictor computes an estimated value of the n -th pixel taking into account values of the previous r neighboring pixels. Then, only the estimation errors are encoded, i.e., the difference between the actual and the predicted values, which is usually close to zero. The obtained error values of the prediction are considerably lower than the initial values of the variance, which decreases the value of unconditional entropy.

Blending predictors are an alternative technique to the selection of the temporarily best predictor. The blending predictors are a linear combination of m subpredictors whose weights depend on the error level from the closest neighborhood. The importance of a subpredictor is inversely proportional to the prediction errors obtained in the closest neighborhood.

According to [8], it is possible to achieve better effectiveness improvement when a few subpredictors are attached. These subpredictors, for example GradNorth, GradWest, Plane-2, give poor results when using separately (in comparison with, e.g., Pirsh predictor). Moreover, in the proposed algorithm named Blend-14, the selected pixels are used together with the two improved versions of the non-linear predictors from CALIC and JPEG-LS instead of the Pirsh predictor. The former is the modification named GAP₊ and the latter is MED₊ [6].

In Fig. 1, the comparison of the average values of entropies obtained by measuring 9 benchmark images of the size 720x576 pixels (*ballon, barb, barb2, board, boats, girl, goldhill, hotel, zelda*) is provided. In the first case (GAP₊), one predictor has been used, the remaining results are produced by three competitive subpredictor blending methods. The results obtained with the proposed technique are 1.16 per cent better than by the second-best WAVE [4] technique, which constitutes a significant improvement in the lossless




Figure 1: Comparison of unconditional entropy for various prediction methods

compression field.

More details on the applied predictor blending methods are provided, e.g., in [9].

3. NETWORKS ON CHIPS

Networks on Chip (NoC) is a System-on-Chip (SoC) paradigm aiming at improving communications within large systems implemented on a single chip [3]. In this approach, system modules (such as processor cores, memories, or other IP blocks) exchange data using routers connected to an on-chip network. An arbitrary module is capable of sending data to another module, so the data is transferred between a number of routers, which make a routing decision. As the majority of NoC implementations use a packet switching approach, a portion of data is usually split into packets. These packets are then typically transferred between modules using the wormhole switching mode [5]. Wormhole routing is a simple routing system that is implementable when all links are fixed. In this routing type, a transmitted packet is additionally split into equal-length units called flits (flow control units). The first flit of a packet contains a header, which includes necessary routing information, including a destination address. After obtaining the header flit, a virtual path is established. The remaining flits of the packet follow the header via the same route.

Links in a NoC can be designed in a way that they transfer a flit fraction, known as phit (physical control digit), in a single clock cycle.

In our work, we use a predominant NoC topology - a square mesh, where each node (except boundary ones) is connected to four neighbors. In this topology, it is possible to address the modules with its X and Y coordinates for the horizontal and vertical position, respectively. This addressing type allows to use a simple XY routing algorithm, briefly described later in this paper.

Usually in various realizations of the NoC architecture the phit length is equal to the flit length, i.e., the whole flit is transmitted into one clock cycle or one request-acknowledge pair of events. However, it has been demonstrated in [1], [2] that decreasing the phit length value can be beneficial with respect to numerous criteria. We will refer to the links where the phit length is lower than the flit length as to phit-serial links.

Phit-serial transmission type are characterized with better routability due to fewer wires. The serializer and deserializer modules have short critical path which allows us to use a

high clock frequency. It can reduce both the area and cost significantly, as well as diminish the complexity of on-chip routers. Thanks to the phit-serial links, there is a possibility of coexistence of a high-level parallelism with a high-speed processing [2].

Using the phit-serial instead of n -bit width bit-parallel links requires n times more clock cycles to obtain the same communication performance. Therefore, composing a bit-serial stream into m -bit (where $m < n$) slices can assure better practical compromise between routing and processing outlays [1].

4. SYSTEM ARCHITECTURE

The proposed approach consists of blocks that form two clusters: control and predictor ones (Fig. 2). The blocks in the control cluster are an I/O and a decision unit. The former one reads data from a local buffer and sends it to the predictor blocks grouped in the predictor cluster. These blocks compute various predictors, as mentioned earlier in this paper. Thanks to hardware implementation, the subpredictor computations are performed in parallel. The I/O block sends packages to the members of the predictor cluster in serial. When the data describing the required neighborhood pixels is transferred from the I/O block into predictor blocks, the I/O block is capable of working on the next image pixel. When all the predictor blocks finish computation of the subpredictors values, they send them together with the neighborhood weights to the blending unit, the second member of the control cluster. It returns a single prediction value by performing the subpredictors blending. Finally, the main prediction error is transmitted to the arithmetic/run-length encoder, which forms the final output stream. The functionality and implementation of this block is beyond the scope of this paper, but may be easily incorporated into the described mesh NoC architecture.

Two connected routers are synchronized with a common clock named *Communication clock*. This guarantees that one *Clock domain* is small enough to avoid problems with synchronization and clock signal distribution. Moreover, the critical path in both *Serializer* and *Deserializer* is very short. These features allow the *Communication clock* to increase significantly and thus guarantee that the frequency of the *Communication clock* is much higher than the frequency of the *Master clock*, which is a sufficient condition for benefiting from the phit-serial links.

5. IMPLEMENTATION DETAILS

We implemented the proposed architecture in SystemC hardware description language using Synopsys CoCentric SystemStudio¹, a system-level design creation, simulation and analysis tool.

In Fig. 2, two clusters are presented utilizing the NoC environment that is comprised of nine wormhole routers connected with local cores and neighbor routers. In the proposed approach, we use seven different predictors, so seven predictor IP blocks forms the Predictor cluster. The second, *Control cluster*, is comprised of the I/O block and the *Blending unit*.

In order to assign cores functionality into mesh NoC cells, we have taken into account the number of flits ex-

¹The Software described in this document is furnished under a license from Synopsys (Northern Europe) Limited.

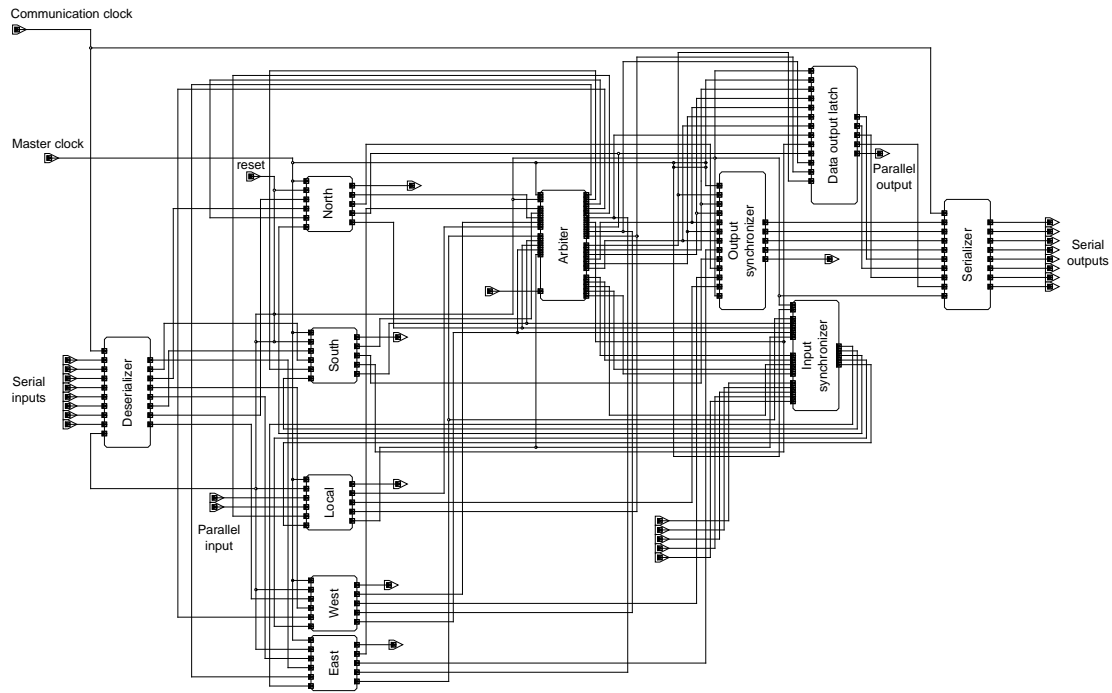


Figure 4: Blocks in a wormhole router

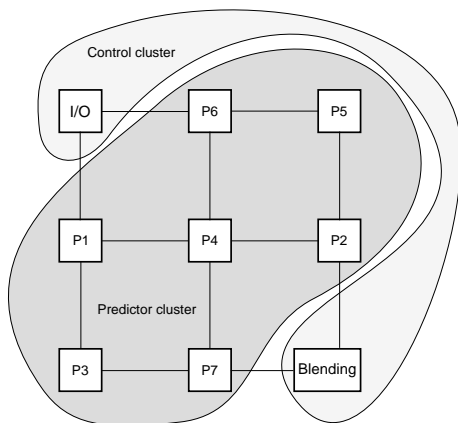


Figure 2: Two clusters in the developed system




Figure 3: Communication stages between two neighbor routers (F - flit length, P -phit length)

changed between all pairs of the cores and use a greedy mapping approach.

Two neighbor routers are connected to each other with a set of two unidirectional data buses (each in different direction) with a number of wires equal to the phit length together with *Communication clock*, as presented in Fig. 3. *Serializer* is connected to *Input buffer*, which uses the 2-stage bundled data asynchronous protocol to communicate with a wormhole router. The length of each flit is equal to 8 bits, so it may be split up to the eight phits that are transmitted in a single clock cycle.

The first flit of a package includes the (X,Y) address of a destination IP block and is necessary to establish a route between the source and the destination IP block. The second flit includes a payload data length count equal to a number of remaining flits.

Having obtained the first flit, a wormhole router determines the next router based on the target IP block address. Then, it transmits the first and the remaining flits using the selected output port. It initializes a flit counter with the number given in the second flit and decreases it after receiving the following flits as long as it is equal to zero, i.e., the whole package has been transmitted by this router.

A single router is capable of transferring two packets at the same time provided that they use different ports. Besides, every port is equipped with a small buffer, which accumulates flits when the output port is busy to increase its throughput.

In Fig. 4, a block of a single wormhole router placed in the center of a 3x3 mesh is presented. This block consists of four input buffers (described as *North*, *South*, *East*, and *West*), each connected with one of the four neighbors, and the *Local* IP block. All these input buffers are connected with an *Arbiter*. The role of the *Arbiter* is to decide which input

buffer is capable of sending data through a router output. In order to select the next router in a route to the destination IP core, the *Arbiter* uses the well-known XY routing algorithm, where a flit is firstly routed to the appropriate horizontal position and then in the vertical direction. It has been proved that XY routing produces minimal paths without redundancy and is deadlock-free [7]. The *Output synchronizer* takes the information from the *Arbiter* where to send the data from a particular input buffer, and then transfers the appropriate control signal to the *Serializer*. Having obtained this signal, the *Serializer* sends the flit obtained from *Data output latch* to an appropriate neighbor router. The flit is decomposed into a number of phits, which are transmitted synchronously in serial to the next-hop router. The *Deserializer* in the next hop router gathers all the received phits and sends it into an appropriate input buffer.

The connections with the local IP core are realized without serializing/deserializing. So, these connections are labelled as *Parallel input* and *Parallel output* in Fig. 4.

6. PERFORMANCE ANALYSIS

In the proposed approach, the flit length, F , is to be higher than $2 \cdot \lceil \log_2 N \rceil$, where N is a size of the mesh (in our case $N = 3$). This is caused by the fact that it is necessary to write the whole (X,Y) address into one flit. The length of a phit, P , belongs to range $\langle 1, F \rangle$. *Master clock* and *Communication clock* have frequencies f_m and f_c , respectively. In order to benefit from the phit-serial communication, we have to guarantee that $f_c \gg f_m$.

To transfer a packet with l payload bits we have to split it into flits and add two extra flits with a header, so the real number of transferred bits is equal to

$$L = \left\lceil \frac{l}{F} \right\rceil F + 2F. \quad (1)$$

The value of $\Gamma = \frac{L}{F}$ denotes the number of flits in a packet.

In case of a lack of contention and $P = F$, a constant number of clock cycles is required to transmit the first flit (t_f) and the remaining packages (t_r), $t_r < t_f$. In case of a phit-serial realization, however, the corresponding values t_{fp} and t_{rp} are equal to

$$t_{fp} = \frac{t_f}{f_m} + \frac{F}{P \cdot f_c}, \quad (2)$$

and

$$t_{rp} = \frac{t_r}{f_m} + \frac{F}{P \cdot f_c}, \quad (3)$$

respectively. (Transmission of the first flit requires more time as the wormhole router has to determine the next router, the remaining flits are sent to the same router.)

The time of transfer the whole package t_p for a single router with no contention is equal to

$$t_p = t_{fp} + t_{rp} \cdot \left(\left\lceil \frac{l}{F} \right\rceil + 1 \right). \quad (4)$$

In a regular NoC $N \times N$ mesh architecture there are $2(N-1) \cdot N$ links each utilizing $2(P+1)$ wires.




Figure 5: Transfer time of an average-length package for a single router

7. EXPERIMENTAL RESULTS

We simulated the described system at bus cycle-accurate abstraction level using CoCentric SystemStudio simulator. According to this simulation, a wormhole router requires eight cycles to route the first flit of the package, i.e., $t_{fp} = 8$, and four cycles for the remaining packages, i.e., $t_{rp} = 4$. We also measured, that the average payload size in a package, l , is equal to 34.46 bits for the described above compression technique. Using these parameters, we determined the time of transfer the whole package t_p for a single router and presented it in Fig. 5 for four different phit lengths. In the described system, the average number of wormhole routers crossed by a single package is equal to 2.08, which is slightly better than a random mapping technique returning 2.29 on average. From the graph it follows that t_p increases gradually with an increase of a flit length up to the point where it falls steeply due to the non-linear changes in the length of transferred bits (1). An increase of the phit length results in decreasing t_p , but with the cost of the growth in the number of wires, as described below.

Our first SystemC realization was reported in [10]. In that implementation, we used dedicated wiring for connecting modules. The total number of wires was equal to 428, which may be too large to fit on a reprogrammable device. In the proposed phit-serial realizations, we set the flit length $F = 8$ and the phit length $P \in \{1, 2, 4, 8\}$ and for these values we got 96, 120, 168, and 264 wires, respectively, which is much less in comparison with dedicated wiring approach. Moreover, the wires used in the NoC approaches are very short - they connect only neighbor cells, whereas wires in the dedicated approach were longer, which may decrease the transfer time in physical realizations.

8. CONCLUSION

A routing-efficient realization of a lossless compression based on subpredictors blending utilizing the phit-serial Network on Chip paradigm was described.

The experiments on the influence of flit and phit lengths on a package transfer time confirmed its non-linear character. An increase of a flit size lengthens the routing time, whereas decreasing of the phit size results in diminishing the system

performance. However, in this case the synthesized chip area and routability expenditures are more suitable for hardware realizations.

The proposed phit-serial implementation of the proposed algorithm required, depending on the phit size, only from 22.43% to 61.68% of the wires used in an earlier hardware implementation utilizing dedicated-wiring technique.

The obtained solution meets all the enumerated requirements imposed on data storage and manifests its suitability for hardware designs of lossless compression systems.

Final note

Synopsys and the Synopsys product names described herein are trademarks of Synopsys, Inc.

REFERENCES

- [1] R. Dobkin, R. Ginosar and A. Kolodny, Fast Asynchronous Shift Register for Bit-Serial Communication, *12th IEEE International Symposium on Asynchronous Circuits and Systems - ASYNC'06*, France, 2006.
- [2] Y. Wang, D. Li, T. Isshiki, H. Kunieda, A Novel Fingerprint SoC with Bit Serial FPGA Engine, *IPSJ Journal*, Vol. 46, No. 6, 2005.
- [3] L. Benini, G. de Micheli, Networks on Chips: A New SoC Paradigm, *IEEE Computer*, vol. 35, no. 1, January 2002, pp. 70-78.
- [4] G. Deng, H. Ye, A general framework for the second-level adaptive prediction, *Proceedings IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP'03)*, vol. 3, April 2003, pp. III-237-240.
- [5] A. Ivanov, G. de Micheli, Guest Editors' Introduction: The Network-on-Chip Paradigm in Practice and Research, *IEEE Design & Test of Computers*, vol. 22, no. 5, September-October 2005, pp. 399-403.
- [6] J. Jiang, C. Grecos, Towards an improvement on prediction accuracy in JPEG-LS, *Optical Engineering*, SPIE, vol. 41, no. 2, 2002, pp. 335-341.
- [7] M. Li, Q.A. Zeng, W.B. Jone, DyXY: a proximity congestion-aware deadlock-free dynamic routing method for network on chip, *43rd ACM IEEE Design Automation Conference*, San Francisco, CA, USA, 2006, pp. 849-852.
- [8] T. Seemann, P. Tischer, Generalized locally adaptive DPCM, Department of Computer Science Technical Report CS97/301, Monash University, Australia, 1997.
- [9] G. Ulacha, R. Stasiski, Parameter choice for predictor blending and application in lossless image coding, *Measurement, automation, control*, no. 7b, 2006, pp. 95-97.
- [10] G. Ulacha, R. Stasinski, P. Dziurzanski, R. Olejnik, Lossless compression system architecture dedicated to Networks on Chips, *Int. Conf. on Signals and Electronic Systems (ICSES'06)*, Lodz, Poland, 2006, pp. 235-238
- [11] The International Technology Roadmap for Semiconductors. 2005 Edition. Semiconductor Industry Association, 2005.