MPEG-4 Performance Analysis for a CDMA Network-on-Chip

Manho Kim, Daewook Kim and Gerald E. Sobelman
Department of Electrical and Computer Engineering
University of Minnesota, Minneapolis, MN 55455, USA
Email: {mhkim,daewook,sobelman}@ece.umn.edu

***Abstract*— Realistic traffic patterns for a multi-processor MPEG-4 architecture are used to evaluate the performance of network-on-chip (NoC) implementations. In particular, we study the characteristics for a design that is based on CDMA switching techniques and a star-network topology. The results are compared to those for a more conventional mesh-topology NoC. We evaluate metrics for bandwidth requirements, latency and area overhead and show that the CDMA star design is a good candidate for the implementation of these systems.**

I. Introduction

Multimedia applications are widespread and will become even more important in the future. Video telephony, digital television, video games, virtual reality simulators, etc. are growth areas of the future and the MPEG-4 standard has emerged as a key ingredient in many of these systems [1], [2], [3]. Therefore, efficient hardware platforms to perform the set of algorithms within the standard are of great interest. Since these computations are varied, it is necessary to include a range of hardware resources in the system such as a DSP processor, RISC CPU, graphics engine, etc. [10]. As a result, there is a need for high-throughput communications links between these blocks, and this can become a performance bottleneck. A bus-based interconnect scheme is a shared medium which does not scale well to large systems requiring very high aggregate bandwidth.

Networks-on-chip have been proposed as a way to overcome this limitation and provide a scalable interconnect environment [4]. Several types of network switches and topologies have been proposed, but most of the performance analysis is done using random traffic models where the computational blocks are simply modeled as random number generators without respect to any particular application. These types of analyses are only of limited utility, since they do not address the actual traffic requirements that one would find in an actual application. Some recent papers have used more realistic traffic models. For example, Varatkar et al proposed on-chip traffic analysis using self-similar processes [7]. Ching et al discussed integrated modeling and traffic generation for a reconfigurable NoC [8]. Murali and De Micheli [5] proposed automatically mapping cores onto a network architecture using a mesh/torus topology by making use of the Xpipes library [6]. Wiklund et al use both random and 'data plus control' traffic models in their analyses [9].

In this paper, we study the performance and area overhead of NoCs for MPEG-4 system implementations. In particular, we focus on the properties of a recently proposed NoC that is based on using CDMA switching techniques to concurrently route multiple data streams between computational resources [11]. We use the average data rates between the blocks in a realistic MPEG-4 implementation that uses a CDMA star network topology. In addition, we compare the performance and area overhead of the CDMA NoC to that of a more conventional mesh topology NoC for this application.

The remainder of this paper is organized as follows. In Section II, we describe the mapping methodology that is used to create the MPEG system implementation and traffic model. In Section III, we briefly describe the properties of our CDMA-based star network topology. Then, in Section IV, we give the specific mapping that results for the MPEG-4 application onto our NoC, as well as on a baseline mesh topology NoC. Section V presents our results for performance and area overhead and our conclusions are given in Section VI.

II. Mapping Methodology

Our overall procedure for characterizing the performance of the NoC for a particular application such as MPEG-4 is illustrated in Figure 1.

We start with the given communication characteristic parameters such as bandwidth requirements, payload size, buffer size, and operating clock period, etc. The traffic generator, which is actually a resource (IP) model written in a hardware description language, generates traffic with a given probability of the packet being sent out to a given destination resource. The generated input traffic data is used to simulate the application on the NoC platform. Using the specified communication characteristics, a mapper groups the resources which communicate frequently onto the same switch in order to reduce latency. The transmitted and received packet traffic is traced in a log file. In a post-processing phase, we use the log file to analyze the performance. In particular, the latency of packet transmission and the FIFO buffer full signal are monitored for their performance. These steps can be iterated until all the requirements are met. After finding the best case structure and mapping for the NoC platform, we can then synthesize it and thereby obtain the estimated frequency and area overhead of the system.

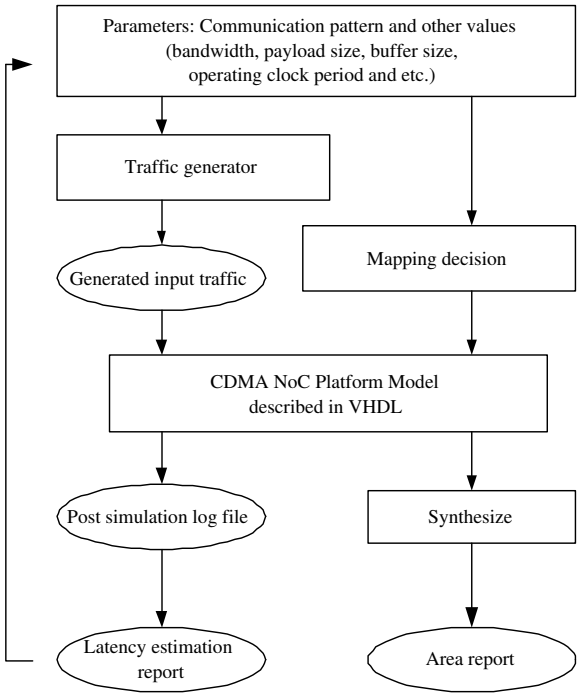

Fig. 1. Mapping Methodology

We consider a particular MPEG-4 video processing application that was presented in [12]. We have used the IP resources and the inter-resource bandwidth requirements specified in that paper as our starting point. There are a total of 12 IP blocks in this design. They are listed as follows, where the number in parentheses is used as an index to identify each of the blocks: audio output processor (1), audio DSP processor (2), media CPU (3), video output processor (4), 3D graphics processor (5), SDRAM (6), SRAM1 (8), quantization unit (9), SRAM2 (10), RISC CPU (11), scaling unit (12) and upsampling unit (13). The communication requirements between these blocks are specified in the data structure of Figure 2. In this table, the entry (i, j) specifies the average communications bandwidth from block i to block j. (The bandwidth data specified in Ref. [12] are the total bidirectional bus traffic between pairs of blocks. For simplicity, we have split the total value equally amongst the two directions of data transfer between each pair of IP blocks.)

	1	2	3	4	5	6	8	9	10	11	12	13
1						0.25						
2						0.25						
3						25	20					
4						95						
5						300		20				
6	0.25	0.25	30	95	300						16	455
8			20									
9					20				250			
10								250		125	87	335
11									125			
12						16			87			
13						455			335			

Fig. 2. Bandwidth requirements for the MPEG-4 system.

The entries in this table can be sorted to indicate the most frequently communicating blocks in the design. These blocks can then be given the highest priority in terms of scheduling. Also, once this table has been constructed, the relative rate of transmitting a packet can be determined for each of the IP blocks and a corresponding packet traffic file can be generated for use in simulations. The transmission probabilities are calculated as follows: The highest bandwidth requirement for this application is the 455 MBytes/second between the SDRAM and the upsampling unit. All other bandwidth values in the table are then normalized relative to this value, so that this path will have a transmission rate of 100%, while all of the others will have a rate less than this. In other words, the path between those two IP blocks will always have traffic, whereas the other paths will pass traffic less often based on their transmission percentages. The normalized values for the various paths are given in Figure 3.

	1	2	3	4	5	6	8	9	10	11	12	13
1						0.05						
2						0.05						
3						5.5	4.4					
4						21						
5						66		4.4				
6	0.05	0.05	6.6	21	66						3.5	100
8			4.4									
9					4.4				55			
10								55		27	19	74
11									27			
12						3.5			19			
13						100			74			

Fig. 3. Normalized bandwidth requirements for the MPEG-4 system.

III. CDMA Star NoC Architecture

In a wired CDMA communication network, each data bit is represented as either an L-bit Walsh codeword or its one's complement depending on whether the bit is a 0 or a 1, respectively [11]. We refer to this process as modulation. While this leads to an increase in the number of bits to be transmitted by each resource, it is offset by the fact that up to $L-1$ resources transmit concurrently through a switch. Each packet is composed of group identification, source address, destination address and payload fields. The transmitter module selects a codeword to use depending on the destination field of the packet and sends out a corresponding modulated codeword. The modulated codewords from the different sources attached to a switch are then summed together using a code adder block. At the receiving side, the demodulation module recovers the original transmitted data using the same codeword that was used for transmission.

Each transmitter module includes a FIFO buffer. This buffer is used for storing packets when other transmitters also wish to send a packet to the same destination at the same time. In this case, the scheduler controls which packet to send depending on a predefined scheduling algorithm. If the FIFO is full, the packet is not dropped. Rather, the transmitter sends a "buffer full" signal to the corresponding resources. Those resources will stop sending packets until that signal is deasserted.

Figure 4 shows our core CDMA NoC switch block diagram.

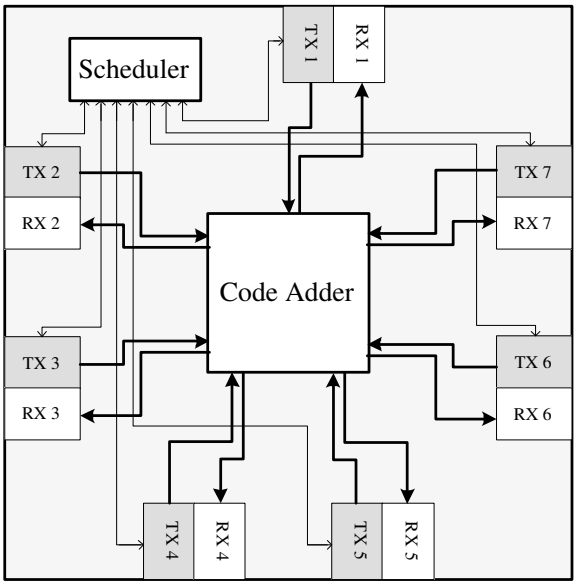

Fig. 4. CDMA NoC Switch Block Diagram

IV. Mapping of MPEG-4

In this section we consider the mapping of the MPEG-4 application onto two types of network-on-chip structures, namely our proposed CDMA star network and a crossbar mesh network.

A. Mapping onto Star NoC with CDMA Switch

Fig. 5 is a proposed mapping of the MPEG-4 decoder system onto our CDMA NoC architecture. In the figure, solid lines represent actual links between a switch and a processing element (PE). A dotted line represents the required communication bandwidth in MBytes/sec.

From the given communication characteristic parameters such as bandwidth requirements, payload size, buffer size, operating clock period, etc, the traffic generator creates packets with the specified probability. The generated input traffic data is used to simulate the MPEG-4 computations mapped onto the NoC platform.

B. Mapping onto Mesh NoC with Crossbar Switch

To compare our CDMA NoC platform with another implementation, we considered mapping onto a crossbar-based mesh topology NoC, as shown in Fig. 6. The crossbar switch has the same input buffer scheme and a buffer size of 8. Because there are a total of 12 IP blocks, we used a 4-by-3 mesh topology to accommodate all of these resources. The shortest path routing scheme is used for this mesh topology.

V. Simulation Results and Performance Analysis

In order to generate the results of interest, we logged all input and output packet transactions into a file. In a post-processing phase, we use a script to analyze the average time to deliver packets and the buffer utilization that was obtained.

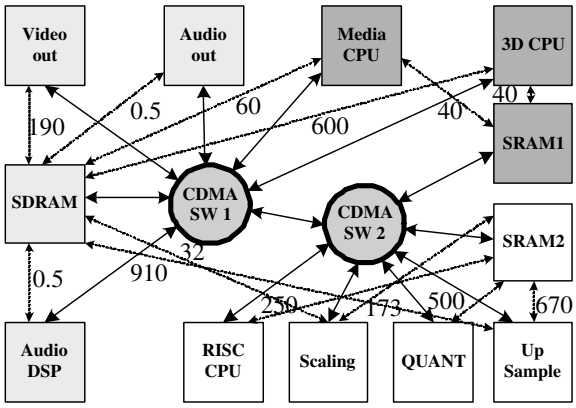

Fig. 5. Mapping onto Star topology NoC with CDMA Switch

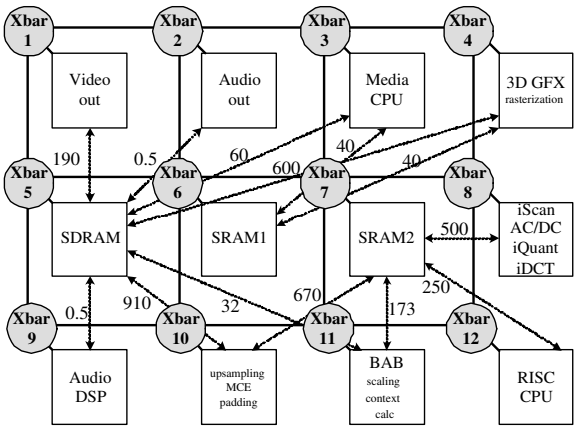

Fig. 6. Mapping onto Mesh topology NoC with Crossbar Switch

We also compared the hop count for both platforms. For the area comparison, we synthesized both the CDMA star and the crossbar mesh switches using the Synplify ASIC tool with the Chip Express CX4001 0.25 μm structured library.

A. Hop Count

TABLE I
Hop count comparison

	CDMA Star NoC	Crossbar Mesh NoC
Average	1.451	3.063
Min	1	2
Max	2	6
Standard deviation	0.65	1.39

The hop count for a packet is defined as the number of routers it has been forwarded through. Table I shows the average number of hops for both platforms. The table indicates that the CDMA star topology has favorable (i.e., lower) hop count values compared to the crossbar mesh topology.

B. Area Overhead

The synthesized area includes the total cell area for either network. In other words, it compares the total area for the 2 required CDMA star switches vs. the total area for the 12

required crossbar mesh switches. We used a buffer size of 8 in both platforms. The estimated maximum frequency is 76 MHz. The Table II shows that our CDMA NoC platform is about two times larger than the crossbar mesh topology platform. The reason is that our prototype CDMA switch uses a more complex algorithm in the TX and RX modules compared to the simple crossbar input and output buffers.

TABLE II
AREA COMPARISON

payload size	Area [μm^2]	
	CDMA star NoC	Crossbar Mesh NoC
8 bits	109,090.0	42,383.5
16 bits	165,003.6	69,535.8
32 bits	273,014.0	128,946.0
64 bits	498,256.4	232,813.8

C. Latency

After analyzing the traffic log file, we can obtain the travel times of the packets during the transmission. Note that the latency values include the effects due to contention between packets destined for the same address at the same time. Latency represents one of the important performance parameters of the NoC platform and is computed as follows:

$$average\ latency = \sum_{i=1}^{n} \frac{T(i)_{received} - T(i)_{transmitted}}{N}$$

where N is the total number of received packets.

Table III shows that our CDMA star topology is around 9 times faster than the general crossbar mesh topology.

TABLE III
LATENCY COMPARISON

	CDMA Star NoC	Crossbar Mesh NoC
Average	28 clock cycles	269 clock cycles

D. Bandwidth Constraints

The highest bandwidth requirement in our system, as given in Figure 2, is 455 MBytes/sec. We would like to determine if our CDMA NoC can meet that constraint. The largest possible bandwidth is the maximum clock frequency, which was found to be 76 MHz, multiplied by the payload size in bytes. Of the cases considered, only the 64-bit payload size is sufficient to meet this constraint: 76 MHz $\times$ 8 bytes = 608 Mbytes/sec. This is a best-case value that does not take into account possible effects due to contention. However, the number is sufficiently high to strongly suggest that the network is fast enough to meet the MPEG throughput requirements.

VI. CONCLUSIONS

We have obtained performance and area overhead results for a high-performance a CDMA-based network-on-chip implementation of an MPEG-4 processor. Realistic traffic rates between the IP resources in the design were used to determine the average latency for packet transmission. In addition, we compared our proposed NoC implementation to one based on a traditional mesh topology and crossbar switches. It was determined that the latency for the CDMA design is about one-ninth of that for the crossbar mesh network. However, synthesis results show that the area overhead is about 2 times as much for the CDMA network. We also illustrated the basic mapping and traffic generation techniques that can be applied to any multimedia application.

VII. ACKNOWLEDGMENTS

We thank Sang Woo Rhim, Bumhak Lee and Euiseok Kim of Samsung Advanced Institute of Technology (SAIT) for their help with this manuscript. This research work is supported by a grant from SAIT.

REFERENCES

[1] T. Ebrahimi and F. Pereira, *The MPEG-4 Book*, Prentice Hall PTR, 2002.
[2] I. Richardson, *H.264 and MPEG-4 Video Compression: Video Coding for Next Generation Multimedia*, John Wiley & Sons, 2003.
[3] P. Kuhn, *Algorithms, Complexity Analysis and VLSI Architectures for MPEG-4 Motion Estimation*, Kluwer Academic Publishers, 1999.
[4] H. Tenhunen and A. Jantsch, *Networks on Chip*, Kluwer Academic Publishers, 2003.
[5] S. Murali and G. De Micheli, "Bandwidth-Constrained Mapping of Cores onto NoC Architectures," *Design, Automation and Test in Europe Conference and Exhibition*, Vol. 2, pp. 896-901, 2004.
[6] D. Bertozzi and L. Benini, "Xpipes: A Network-on-Chip Architecture for Gigascale Systems-On-Chip," *IEEE Circuits and Systems Magazine*, Vol. 4, No. 2, pp. 18-31, 2004.
[7] G. V. Varatkar and R. Marculescu, "On-chip Traffic Modeling and Synthesis for MPEG-2 Video Applications," *IEEE Transactions on Very Large Scale Integrated Systems*, Vol. 12, No. 1, pp. 108-118, 2004.
[8] D. Ching, P. Schaumont and I. Verbauwhede, "Integrated Modeling and Generation of a Reconfigurable Network-on-Chip," *18th International Parallel and Distributed Processing Symposium*, pp. 139-145, 2004.
[9] D. Wiklund, S. Sathe and D. Liu, "Network on Chip Simulations for Benchmarking," *4th IEEE International Workshop on System-on-Chip for Real-Time Applications*, pp. 269-274, 2004.
[10] M. Pastrnak, P. Poplavko, P. N. H. de With and D. S. Farin, "Data-flow Timing Models of Dynamic Multimedia Applications for Multiprocessor Systems," *4th IEEE International Workshop on System-on-Chip for Real-Time Applications*, pp. 206-209, 2004.
[11] D. Kim, M. Kim and G. E. Sobelman, "CDMA-Based Network-on-Chip Architecture," *IEEE Asia Pacific Conference on Circuits and Systems*, pp. 137-140, 2004.
[12] E. B. van der Tol and E. G. Jaspers, "Mapping of MPEG-4 Decoding on a Flexible Architecture Platform," *Proceedings of SPIE - Media Processors*, Vol. 4674, pp. 1-13, 2002.