

System Level Assessment of an Optical NoC in an MPSoC Platform

M. Brière, B. Girodias, Y. Bouchebaba, G. Nicolescu
École Polytechnique de Montréal
Montréal – Canada

F. Mieyeville, F. Gaffiot, I. O'Connor
École Centrale de Lyon
Écully – France

Abstract

In the near future, Multi-Processor Systems-on-Chip (MPSoC) will become the main thrust driving the evolution of integrated circuits. MPSoCs introduce new challenges, mainly due to growing communication through their interconnect structure. Current electrical interconnects will face hard challenges to overcome such data flows. Integrated optical interconnect is a potential technological improvement to reduce these problems. The main contributions of this paper are i) the optical network integration in a system-level MPSoC platform and ii) the quantitative evaluation of optical interconnect for MPSoC design using a multimedia application.

1. Introduction

In the near future, Multi-Processor Systems-on-Chip (MPSoC) devices will become impossible to circumvent on the integrated electronics market [1]. Processor data rates will be critical and will quickly reach a bandwidth of several tens GHz. Interconnects will play a significant role for MPSoC design in order to support these high data rates. Several electrical interconnect architectures exist or are being developed, each of them trying to overcome current limitations of bandwidth, contention and latency. Despite these efforts, the International Technology Roadmap for Semiconductors (ITRS) [2] still predicts that interconnects will become the MPSoC bottleneck.

Integrated optical interconnects are considered as an alternative to traditional interconnects [2]. Optics increase bandwidth and decrease latency. Moreover, wavelength routing introduces a totally new dimension by improving the functionality of the routing devices since it is possible to devise fully contention-free structures. The research in on-chip optical interconnects field is technology-oriented and new devices and architectures for optical NoC were proposed [3,4]. The system-level vision showing the impact of these solutions for a global MPSoC is not yet considered. This requires a multi-disciplinary cooperation between physical and system-level designers.

This paper presents the results of such cooperation by proposing a novel approach that enables the complete optical network integration in a system-level MPSoC platform. The application of this approach to a quantitative

evaluation of optical interconnect for MPSoC design is also presented.

The remainder of this paper is organized as follows: section 2 provides an overview on interconnects; section 3 presents the related work, section 4 details an innovative optical interconnect structure; section 5 describes the methodology used in this work, section 6 gives simulation results; section 7 summarizes the lessons learned from the use of an optical network on chip and section 8 concludes this paper.

2. Electrical and optical network overview

Traditional macro-interconnect topologies [5] are applied to interconnect individual MPSoC components. There are three main SoC interconnect architecture families: *i)* traditional crossbar (traditional full crossbar links together SoC components by point to point interconnect), *ii)* shared bus and *iii)* Network on Chip (NoC) [6], such as the STBus [7]. Optic is widely used in long distance communication. A well-known example is in fiber-based technologies where metallic communication networks are outperformed for interconnect lengths greater than a few meters. In the optical integrated circuit field, the development of compact optical devices is used for short integrated communication links. They are technologically compatible with electronic technology (commonly CMOS technology) [8].

3. Related work

Comparing electrical and optical communication architectures is the subject of several academic and industrial analyses. Collet *et al.* [8] have compared simple optical and electrical point to point interconnect using a Spice-like simulator. Tosik *et al.* [3] have studied more complex interconnect by comparing optical and electrical clock distribution, using accurate physical simulations, synthesis techniques and predictive transistor models (130 nm to 45 nm). Both works study power consumption and bandwidth. Intel has also studied performance improvements including technological costs between copper and optical clock distribution [9].

These previous analyses only show comparisons at the physical level without the global view of a complex

system. The work presented in this paper is different from these analyses as it adds the communication view (*e.g.* contention, data flows and processing time) between several processing units. The main contribution of this paper is the complete optical network integration in an innovative system-level MPSoC platform. This integration will allow to study MPSoC performance for electrical and optical interconnect structures.

4. Optical network on chip

The integrated optical communication system studied in this work, also called Optical Network on Chip (ONoC), is composed of three types of blocks: *i)* transmitters, *ii)* a passive integrated photonic routing structure (λ -router) and *iii)* receivers. Fig. 1 presents an overview of this ONoC. It is a heterogeneous structure, being a combination of passive and active optical devices, and mixed analog/digital integrated circuits. The next sections briefly describe the different parts of the ONoC.




Fig. 1. ONoC overview (I=Initiator, T=Target).

4.1. Transmitter and receiver blocks

In this work, each SoC component (also called core) requires a transmitter block which enables the electro-optical conversion (*cf.* Fig. 2(a)). It is mainly composed of Vertical Cavity Surface Emitting Lasers (VCSELs), drivers and a serializer (SER - 32 bits). The use of serializers is mandatory since it is hardly feasible (*e.g.* area, power consumption and floorplanning) to integrate as many lasers as the number of bits to transmit per packet.




Fig. 2. Electro-optical (a) and opto-electrical interfaces (b).

Similarly to the transmitter block, each core requires a receiver block which enables the opto-electronic conversion (*cf.* Fig. 2(b)). It is mainly composed of a PIN photodiode (conversion of flow of photons into photocurrent), a TransImpedance Amplifier (TIA), a

decision circuit (digital signal regeneration) and a deserializer (DES). The deserializer is mandatory for the same reasons as those discussed earlier for the serializer.

4.2. λ -router block

The λ -router is a passive optical network composed of 4-port optical switches (based on add-drop filters [4]) designed to route data through SoC components. Fig. 3(a) presents an example of a $N \times N$ λ -router architecture (each grey square representing an add-drop filter where a physical architecture example is shown Fig. 3(b)).

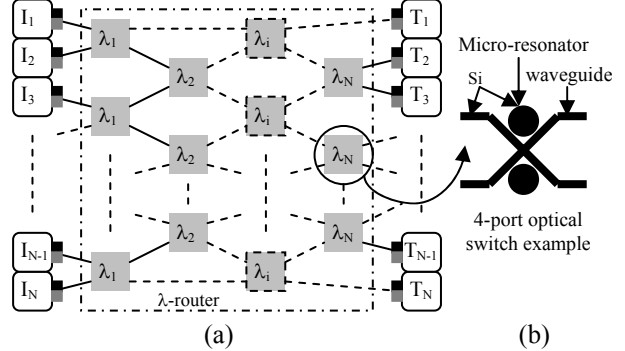


Fig. 3. $N \times N$ λ -router architecture (a), 4-port optical switch architecture example (b).

These add-drop filters (typically taking up $10 \mu\text{m}$ by $10 \mu\text{m}$) are composed of optical waveguide and optical micro-resonators and operate in a similar way to classical electronic switches from a functional point of view. From any input port, switching is obtained to one of the two opposite output ports depending on the wavelength value of the optical signal injected into the optical filter (Fig. 4).

The add-drop is bidirectional and compact devices have been demonstrated in CMOS compatible Silicon on Insulator (SOI) technology (Si/SiO₂ structures accept 1.3-1.55 μm wavelength). As illustrated in Fig. 4, there are three possible switch states depending on the input signal:

- Straight state 4(a) occurs when specific wavelengths, called resonant wavelengths (λ_i , depending on micro-resonator geometry and material) are injected in the filter and are routed through the micro-resonator.
- Diagonal state 4(b) occurs when other wavelengths (λ_j) are injected in the filter and are not routed through the micro-resonator.
- Cumulative state 4(c) occurs when signals of both resonant and non-resonant wavelengths (λ_i and λ_j) are injected into the filter using the WDM technique¹ and are either routed or not routed through the micro-resonator. Because of this property and the fact that the four add-drop ports can be used simultaneously, a contention-free network can be built.

¹ Wavelength Division Multiplexing. Several signals at different wavelengths can be injected into the same waveguide.

- Example 4(d) shows a possible exploitation of the optical switch. This example shows both unidirectional and bidirectional behaviors (several wavelengths simultaneously injected in opposite way).




Fig. 4. Functional states of 4-port optical switch.

The main advantage of this architecture is its high scalability. However, currently, up to 32 cores (16 initiators and 16 targets) can be plugged onto an ONoC, where the limit is due to the lithographical tolerance in add-drop manufacturing.

In a λ -router, only one physical path associated with one wavelength exists between I_i and T_j . The broadcast is also possible with this architecture. Truth table 5(a) represents the operation for a 4×4 network (Fig. 5(b)). For example, if I_1 communicates with T_2 , data must use the wavelength λ_3 to be sent through the λ -router (bold line in Fig. 5(b)). At the same time I_4 can communicate with T_1 using the wavelength λ_4 (dash line in Fig. 5(b)). These optical switches and λ -router have been manufactured and tested. The observed network routing corresponds to theory [10].




Fig. 5. 4×4 λ -router.

4.3. ONoC characteristics

The previous sections have shown that ONoC is a potential solution to interconnect SoC components. This solution becomes even more attractive when coupling the wavelength routing method to the WDM technique. Combining both techniques increases the ONoC functionality and the amount of bandwidth. However, ONoC performances are deteriorated by the interface circuits. The main limitation of optical interconnect is its lack of maturity and strong heterogeneity.

In prior work, the ONoC has been completely designed at the physical level and allows for an accurate estimation of the various ONoC performance metrics [11]. Table 1

summarizes the main ONoC characteristics. This maximal ONoC data rate is limited by the interface circuits, mainly due to the TIA and the decision circuit (technological node limits are reached). Concerning the latency, its value is extracted from low level simulations of each ONoC part (analog/digital circuits with Spice-like, optoelectronic devices with VHDL-AMS and λ -router with finite difference time domain algorithm).

Component	Technology	Limitation	Advantage
VCSEL	III-V	Reactivity	CMOS compatibility
PIN	III-V	Capacitance	CMOS compatibility
λ -router	SOI	Maturity	CMOS compatibility
Analog/Digital interfaces	0.13 μ m CMOS	Bandwidth	integration
Metric		Value	
Datapath width		32	
Max. data rate		3.2 Gb/s / port	
Max. frequency		100 MHz / bit	
Latency (16×16)		1 ns	

Table 1. ONoC specifications.

Currently, the ONoC is a 32-bit device involving only the communication of 32-bit data between MPSoC processing units. There is no contention in the optical interconnect because of passive behavior and use of the WDM technique. There is always an optical path to access an unoccupied recipient. This access is strongly advantageous in the MPSoC approach. However, the data size is smaller than in an electrical network (up to 32 bits).

5. Methodology

5.1. Presentation

Transaction Level Modeling (TLM) approach was applied to develop the ONoC design methodology and to integrate it into a system-level platform. TLM can link software development and SoC design at several abstraction levels higher than the Register Transfer Level (RTL), to decrease the simulation time of complex models. Thus, 3 layers are used in the TLM approach [12]: *i*) message layer (event driven-based and untimed), *ii*) transaction layer (timed and uses communication protocol) and *iii*) transfer layer (byte and clocked cycle-accurate). These three levels can be applied to SystemC-based models. This is one of the reasons why SystemC [13] has been chosen to model the ONoC.

5.2. ONoC modeling

To be able to compare electrical and optical interconnect, the ONoC must be integrated in a system-level MPSoC platform. At the system-level, the description of ONoC operation can be simplified. It can be considered as a full crossbar (point to point connection between SoC cores) with very low latency. In the high-level model, it is possible to integrate optical phenomena due to optoelectronic conversion [11]. However, due to high simulation time and since this is not the objective of a

high level platform, all optical phenomena (*e.g.* optical crosstalk or device manufacturing defects) will be neglected. The high level focuses on testing the global system functionality (MPSoC architecture and application).

At this level, SystemC is used to model the optical network, using the transaction layer of the TLM approach (*cf.* section 5.1.). SystemC facilitates its integration with the MPSoC platform since the platform itself also uses this language (*cf.* section 5.3.). A SystemC model has been created to manage the latency depending on the optical interconnect technology. These latency parameters were extracted from previous results (section 4.3.). This model is called *onoc32* (32 bits) and can be easily linked to other components in the MPSoC model.

It is essential to use an existing communication protocol which will allow the performance comparison between the ONoC and electrical interconnects. Thus, at the system-level, the ONoC can be made compatible with the SystemC Open Core Protocol (SOCP) [14] used by the MPSoC platform. SOCP follows the same high-level semantics as the Open Core Protocol (OCP) and the Virtual Component Interface (VCI) but has no notion of signals or detailed timing. Between each SoC component, SOCP read and/or write requests are simply carried out after a certain ONoC model cycle time.

5.3. MPSoC platform

The StepNP platform was selected [14] to integrate the ONoC. The StepNP platform includes models of processors (standard or configurable), networks-on-chip, configurable H/W processing elements, as well as networking-oriented I/O's. Apart from these domain-specific I/O's, this is a general-purpose platform. The StepNP platform was developed specifically for MPSoC systems. Several programming models are supported in StepNP. These models are inspired by leading-edge approaches for large distributed systems development, adapted to the SoC domain. The StepNP platform traditionally uses electrical interconnects (*e.g.* STBus and crossbar) to facilitate the communication between SoC components. In this work, the ONoC behavior has been integrated into the StepNP platform interconnect library.

5.4. Platform configuration for global validation of ONoC implementation in MPSoC

Fig. 6 shows a high level simulation approach which consists of: *i)* a library of basic elements (*e.g.* interconnects, initiators and slaves), *ii)* applications and *iii)* parameters for the platform configuration. Using configuration scripts, all these elements are linked to create a model and to extract simulation results from the platform.

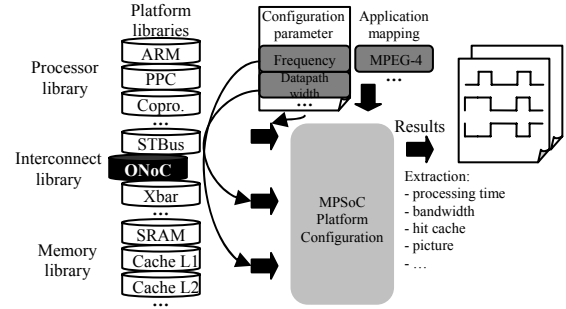


Fig. 6. Methodology.

The steps of the methodology used in this work are:

1. To choose the parameters (architecture and number of processors, number of memories, type and number of interconnects, system frequency, application).
2. To configure the platform with previous parameters.
3. To map the application to the platform.
4. To extract the results: mainly application processing time (estimation of the number of clock cycles required to complete the application).

Table 2 presents the platform configuration. Monothread ARM processors (with configurable cache memory), SRAM external memory and interconnect (ONoC/crossbar/STBus) are the principal components used in the platform. At the same technological node, the ONoC working frequency will be up to 100 MHz compared to the electrical interconnect operational frequency up to 233 MHz.

Parameter	Value
Electrical interco. clock frequency	up to 233 MHz
Optical interco. clock frequency	up to 100 MHz
Initiator type	ARM processor or specific coprocessor
Target type	SRAM Memory or internal slave
Interconnect	ONoC, full crossbar, STBus
Protocol	SOCP
Datapath width	32 bits
Application	MPEG-4
Extracted parameter	Processing time

Table 2. Platform configuration.

In this work, one of the electrical interconnects used in the platform is a full crossbar which has no contention. Only the latency and jitter management are taken into account between an initiator and a target. Regarding the STBus model used in this paper, for all simulation results presented in the next sections, the “transfer layer” STBus model is used and configured to have 4 request/response lines and the possibility to manage up to 14 initiators and 12 targets.

6. Case study: tests and results

This section presents the quantitative evaluation of interconnect architectures (ONoC, STBus and crossbar) by extracting and comparing meaningful performance metrics (latency and processing time).

6.1. MPEG-4 application

The tested application is the MPEG-4 audio/video coding standard, introduced in late 1998 by the ISO/IEC Moving Picture Experts Group (MPEG). MPEG-4 is a massive communication application, and is useful to test interconnect behaviors. Moreover, the MPEG-4 algorithm is easily parallelizable for the MPSoC approach. MPEG-4 introduces a complex algorithm with many mathematical transformations.

In this paper, a 640×480 Audio Video Interleave (AVI) movie was used as source in the MPEG-4 encoder. The simulation time for this application can become strongly prohibitive. Moreover, the first frame is not very significant in terms of processing time since the global algorithm is not executed (*e.g.* motion evaluation code) due to the absence of the previous frame. For both reasons, to study the system performance, the mean of the processing time of the five first frames will be used.

Table 3 presents the average processor statistics (statistic are expended to estimate 30 seconds of MPEG-4 coding). The number of instructions is important. Thus, it implies a large quantity of memory access which may cause some contention in the interconnect.

Number of instructions	5 189 149 313
Number of reads	8 388 753 375
Number of writes	942 367 500
Instruction cache data	
Number of reads	7 050 238 313
Number of writes	517 288 875
Number of misses	1 034 577 375

Table 3. Average processor statistics.

6.2. Platform parameters

Table 2 summarizes the platform configuration. The number of processors varies between 2, 4, 8 and 10. There are 3 coprocessors for specific calculations in the MPEG-4 algorithm and also 3 other initiators for I/O management (not covered in this paper). However, the STBus model used in this paper may connect up to 8 processors. The latency value can be set to the ONoC latency value or to the electrical latency value (*e.g.* the electrical crossbar latency values may go up to 5 clock cycles).

6.3. Simulation results

Fig. 7 illustrates the processing time depending on the number of processors, for the platform configuration given in table 2. Three crossbar clock cycle latencies are tested to study interconnect impact: ideal, 2 clock cycle latency and 5 clock cycle latency, all of them with 200 MHz and 100 MHz working frequencies. The simulation time, for each graphic node, is about 0.4-1 hour/frame using the ONoC.

Fig. 7 shows that, in terms of processing time, the ONoC gives better performance than any traditional electrical interconnect (STBus or crossbar), even with an

operational frequency twice as low. For an 8-processor MPSoC architecture, a speedup factor of ~ 1.5 is obtained against the 2 clock cycle latency crossbar, a speedup factor of ~ 3.2 against the 5 clock cycle latency crossbar and a speedup factor of ~ 2.5 against the STBus interconnect.

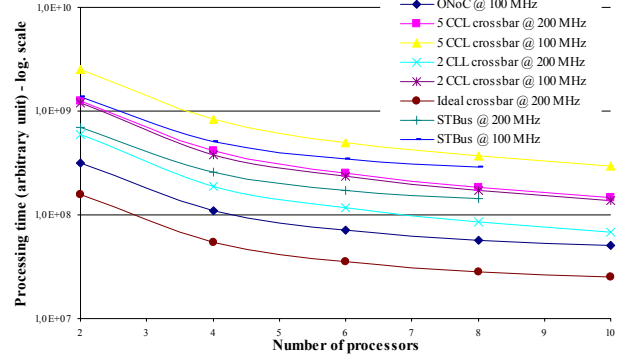


Fig. 7. Simulation results for MPEG-4 application.

Table 4 summarizes the ratio between electrical (STBus, 5 clock cycle latency, 2 clock cycle latency and ideal crossbar at 200 MHz) and optical (ONoC at 100 MHz) processing. These results confirm that the ONoC demonstrates significant advantages over electrical interconnects (speedup factor of 1.35 to 3.96) for any number of processors. Obviously, the ideal crossbar (0 clock cycle latency) gives better results than the ONoC (speedup factor of 2), since it has zero latency and is a contention free interconnect with an operational frequency twice as high. However, the ONoC is closest to an ideal crossbar than the tested electrical interconnects.

	STBus	5 CCL Xbar	2 CCL Xbar	Ideal Xbar
s				
2 processors	2.18	3.96	1.87	0.49
4 processors	2.34	3.8	1.74	0.49
6 processors	2.45	3.49	1.63	0.49
8 processors	2.5	3.21	1.49	0.49
10 processors	x	2.91	1.35	0.49

Table 4. Performance summary: speedup factor between electrical (200 MHz) and optical (100 MHz) interconnect.

Concerning electrical interconnect at a 100 MHz operational frequency, with an 8-processor architecture, the ONoC has a speedup factor of 5.01 compared to the STBus and 6.4 compared to a 5 clock cycle latency crossbar.

6.4. Discussion

Optical interconnect demonstrates a significant speed up gain compared to the electrical interconnect. MPEG-4 is a well suited application for important data flow in the interconnect. About 40 % of data is communicated over the global interconnect structure between the shared memory and the cache. Traditional electrical interconnects are quickly overloaded due to their high latency and

contention which strongly reduce the system's performance. One may also note two different behaviors concerning the ratio between both electrical interconnects and the ONoC when the number of processors increases: *i)* the ratio between the crossbar and the ONoC decreases and *ii)* the ratio between the STBus and the ONoC increases. This phenomenon was not really foreseeable, and is mainly due to the fact that crossbar only has latency and is contention free, whereas the STBus has latency and contention due to its limited number of request/response lines.

For this type of application having high data flow transfer, the optical interconnect may be an interesting future solution.

6.5. Summary

These results confirm that the interconnect latency is a key parameter for MPSoC system performance. Optical interconnect will undoubtedly offer significant advantages once integrated photonic manufacturing reaches full production-quality maturity (area and power consumption). This production-maturity should be reached in a decade [9,15]. However it should be noted that future application innovation and optimization may drastically change the influence of the interconnect structure. Moreover, with current technologies, the optical network working frequency is twice as small as the electrical network frequency, which explains the less significant results between optical network and electrical network (in processing time point of view). Concerning the ONoC power consumption, some estimation has been done in previous work [11,16]. Comparing to STBus [7], the ONoC is not power friendly (mainly due to lasers' driver and SERDES blocks).

7. Lessons learned

Experience on system-level ONoC integration has shown new considerations for MPSoC design integrating optical technology:

- ONoC guarantees low latency and a contention free interconnect,
- ONoC impact on MPSoC performance is limited by its interface circuits (*i.e.* its insufficient technological maturity); at the same technology node, optical network working frequency is twice smaller than electrical network,
- ONoC can decrease processing time compared to traditional electrical interconnects,
- ONoC significantly improves performance for intensive communication applications,
- System-level models enable early and rapid exploration for new technologies impact in MPSoC design.

8. Conclusion and future work

A system-level model of an optical network on chip has been defined and integrated into an MPSoC platform. SystemC has been used to model the ONoC. This system-level model is defined using parameter extractions from physical level. This integration will allow one to estimate performance of future MPSoC using optical or electrical interconnect. Experiments have shown that ONoC may improve performances in MPSoC running intensive communication application.

Future work will study the impact of new applications to obtain different data flows in the interconnect.

References

- [1] P. Paulin, The inexorable progression of parallelism in systems-on-a chip, *System Design Frontier*, vol. 3, no. 2, Feb. 2006.
- [2] ITRS. (2005) International Technology Roadmap for Semiconductors. [Online]. Available: <http://public.itrs.net/>
- [3] G. Tosik, *et al.*, Power dissipation in optical and metallic clock distribution networks in new VLSI technologies, *IEEE Electronics Letters*, vol. 40, no. 3, pp. 198-200, Feb. 2004.
- [4] C. Manolatos and H. Haus, *Passive Components for Dense Optical Integration*. Springer, 2002.
- [5] J. Duato, S. Yalamanchili, and L. Ni, *Interconnection Networks: an Engineering Approach*. Morgan Kaufmann, 2003.
- [6] L. Benini and G. De Micheli, Networks on Chips: A New SoC Paradigm, *IEEE Computer*, vol. 35, no. 1, pp. 70-78, Jan. 2002.
- [7] A. Bona, V. Zaccaria, and R. Zafalon, System level power modeling and simulation of high-end industrial network-on-chip, in *Proc. of DATE*, Feb. 2004.
- [8] J. H. Collet, *et al.*, Performance Constraints for On chip Optical Interconnects, *IEEE J. Sel. Topics Quantum Electron.*, vol. 9, no. 2, pp. 425-432, Mar. 2003.
- [9] M. Kobrinsky, *et al.*, On-chip optical interconnects, *Intel Technology Journal*, vol. 8, no. 2, pp. 129-142, May 2004.
- [10] A. Kazmierczak, *et al.*, Design, simulation, and characterization of a passive optical add-drop filter in silicon-on-insulator technology, *IEEE Photon. Technol. Lett.*, vol. 17, no. 7, pp. 1447-1449, July 2005.
- [11] M. Brière, *et al.*, Heterogeneous modelling of an Optical Network-on-Chip with SystemC, in *Proc. of RSP*, June 2005.
- [12] F. Ghenassia, *Transaction Level Modeling With SystemC: TLM Concepts And Applications for Embedded Systems*. Kluwer Academic Publishers, 2005.
- [13] OSCI. (1999) SystemC. [Online]. Available: <http://www.systemc.org/>
- [14] P. Paulin, *et al.*, *Embedded Systems Handbook*. CRC Press, Nov. 2004, ch. A Multi-Processor SoC Platform and Tools for Communications Applications.
- [15] European Commission funded program. (2004) PICMOS. [Online]. Available: <http://picmos.intec.ugent.be/>
- [16] I. O'Connor, *et al.*, Integrated optical interconnect for on-chip data transport, in *Proc. of NEWCAS*, June 2006.