Agenda

08:30	Welcome and Introduction
08:45-09:15	Keynote: **NoCs: Vision, reality, trends** *Giovanni De Micheli, EPFL Lausanne, Switzerland*
09:20-10:00	Session 1: **Advances in NoCs** (invited presentations) *Chair: Ran Ginosar, Technion, Israel* • **Standards in NoC: What can we gain?** *Axel Jantsch, KTH Stockholm, Sweden* • **Networks and Applications: Are Application-Specific Networks Worth the Trouble?** *Wayne Wolf, Princeton U, United States*
10:00-10:20	BREAK
10:20-11:00	Session 2: **Advances in On-Chip Interconnects** (invited presentations) *Chair: Pol Marchal, IMEC, Belgium* • **On-Chip Interconnect: The past, present, and future** *Eby Friedman, U of Rochester, United States* • **Potential impact of emerging System-in-Packaging technologies on system design** *Eric Beyne, IMEC, Belgium*
11:00-12:00	Introducing Research Groups *Chair: Avinoam Kolodny, Technion, Israel* *Presenters:* • Shashi Kumar, U Jönköping, Sweden • Jari Nurmi, U Tampere, Finland • Hoi-Jun Yoo, KAIST, Korea • Radu Marculescu, CMU, United States • Andre Ivanov, UBC, Canada • Christer Svensson, U Linkoping, Sweden • Marcello Coppola, ST Microelectronics, France • Shinobu Fujita, Toshiba, Japan • Kees Goossens, Philips Research, The Netherlands • Philippe Martin, Arteris, France • John Bainbridge, Silistix, United Kingdom • Ian O'Connor, LEOM, France • Graham Hellestrand, VaST Systems, United States
12:00-13:00	LUNCH
13:00-14:30	Poster Session This session will enable direct interactions among all researchers. Its goals are to get to know each other and each other's work, and to create opportunities for research collaborations.
14:30-14:45	BREAK
14:45-15:45	Panel: Looking Through the On-Chip Channels *Moderator: Luca Benini, DEIS - Bologna U, Italy*
15:45-16:30	Planning session for future NoC Symposia
16:30	CLOSE

Preface

Advances in semiconductor technology, design re-use and tools are enabling designers to put complex, massively parallel multiprocessor systems on a single chip. The idea of packet-switching Networks on a Chip (NoC) offers more flexibility, robustness and better resource utilization than what is provided by traditional bus-based architectures. NoCs are poised to be an alternative to buses, but should they rather be seen complementary than alternative in the longer term? What kind of new communication paradigms and signaling schemes are more suited for NoCs in order to battle inherent signal integrity and soft errors? Are there fundamental issues about on-chip interconnects that will limit performance and dependability of communication as Moore's Law gradually runs out of steam within the next decade? What about new technologies such as optical interconnects? These and other questions are in the centre of this workshop, which brings together researchers actively working on future interconnects and NoCs.

Since the early discussions about “routing packets not wires” only some five years ago there has been a rapid increase of research interest to this area. It has led to a large number of publications, including journal and conference papers, book chapters as well as a number of monographs. Leading international conferences on design automation, design and test, systems on a chip have had NoCs as a subject for a session or a tutorial. Our workshop is probably the first specially organized event dedicated solely to the subject of NoCs.

Although this is only a one-day event it will enjoy a variety of activities in its highly intensive programme.

The morning part will start with a keynote lecture by the leading experts in the area Giovanni De Micheli and Luca Benini, followed by four invited technical talks by Axel Jantsch, Wayne Wolf, Eby Friedman and Eric Beyne, who will address a range of issues from applications of NoCs to future underlying technologies for interconnects. Then, leaders of thirteen research groups, academic and industrial, from different parts of the world will give brief reviews of their groups’ research activities.

A central component of the workshop is its lunchtime and early afternoon poster session, involving practically all attendees. In these proceedings, you can find one page abstracts of 60 posters from 36 academic and industrial organizations from 15 countries of Europe, North America and Far East. These posters nicely cover the range of research themes, organized into 8 sections that fall into two main tracks: Track A “NoC Architectures, Design, Prototyping, CAD support, QoS and Applications” and Track B “Interconnection Technology, Signaling Schemes, Components, Infrastructures (Power, Testing, Fault-tolerance)”.

The afternoon part will feature a one hour industrial panel, moderated by Steve Furber. The panelists are John Bainbridge, Marcello Coppola, Kees Goossens, Philippe Martin

and Christian Sauer. They are well-known experts from the European companies, both large and SMEs, working in the NoC area. The panelists will be ventured to address the challenges and opportunities for on-chip interconnect paradigms and project their vision on the current developments in academic and industrial research.

The workshop will serve to launch a new series of high quality symposia, targeted at NoCs and Interconnects. The need for such a dedicated forum is increasingly recognized due to the steady growth in research activity and contributions that are currently spread over multiple conferences in diverse areas such as architecture, circuits, CAD and networking. It will also provide an excellent thematic framework for interaction between industry and academia.

It was a great pleasure and privilege for us to organize this workshop. We hope that those coming to it will find the meeting useful and enjoyable.

We acknowledge the support from the DATE conference organizers who kindly agreed to host this event. In particular we are grateful to Bashir Al-Hashimi for his encouragement and help at all stages of the organization. We thank Infineon for offering us help with additional poster boards. We are immensely thankful to the international NoC community who have been very proactive and without whose remarkable (record-breaking!) response to our call for posters this event would not attract such a fantastic turn-out.

Our special thanks go to Prof. Avinoam Kolodny from Technion, who has de facto acted as our “fourth musketeer” – he constantly provided us with advice, feedback and organizational help.

Workshop Organisers:

Ran Ginosar (Technion, Israel)
Pol Marchal (IMEC, Belgium)
Alex Yakovlev (University of Newcastle-upon-Tyne, UK)

Table of contents

Track A: NoC Architectures, Design, Prototyping, CAD Tool support, QoS, Applications

A.1 NoC General: Architectures, Platforms, Designs, Prototyping:

- **"A Hierarchical Approach to Network-on-Chip"**
  T. Ahonen and J. Nurmi
  Tampere University of Technology, Finland
- **"NoC-Based Platform Implementation on FPGA"**
  T. Ahonen, J. Kylliäinen, C. Brunelli and J. Nurmi
  University of Technology, Finland
- **"PROTEO – A Flexible Network-on-Chip Scheme"**
  T. Ahonen[1], D. Sigüenza-Tortosa[1] and J. Nurmi[1]
  1 Tampere University of Technology, Finland
  2 Complutense University of Madrid, Spain
- **"Towards a Communication-Centric Design Methodology"**
  A. Banerjee, R. Francis, J. Lee, J. May, S. W. Moore and R. D. Mullins
  Computer Laboratory, University of Cambridge
- **"Clockless On-Chip Networks and MANGO"**
  T. Bjerregaard, J. Sparsø and M. Stensgaard
  Technical University of Denmark, Denmark
- **"STNoC™: An Evolution Towards MPSoC Era"**
  M. Coppola[1], C. Pistritto[2], R. Locatelli[1] and Alberto Scndurra[2]
  [1] STM, AST Grenoble Lab, France
  [2] STM, HPC, Italy
- **"Introduction to KAIST BONE project"**
  K. Lee, S. Lee, D. Kim, K. Kim, J. Kim and H. Yoo
  Korea Advanced Institute of Science and Technology (KAIST), Korea
- **"Synchronous Latency Insensitive Design in Æthereal NoC "**
  X. Ru[1], J. Dielissen[2], C. Svensson[3] and K. Goossens[2]
  1 Philips Semiconductors, Switzerland
  2 Philips Research, The Netherlands
  3 Linköping University, Sweden
- **"A Lightweight NoC for the NOVA Packet Processing Platform"**
  C. Sauer[1], M. Gries[1], S. Dirk[1], J.-C. Niemann[2], M. Porrmann[2], U. Rückert[2]
  1 Infineon Technologies, Access Communications, Munich
  2 HNI Paderborn

- **“Interconnects and network on chip at Linköping University”**
  C. Svensson, Fellow IEEE,
  Dept. Electrical Engineering, Linköping University, Sweden
- **“FAUST, an Asynchronous Network-on-Chip based Architecture for Telecom Applications”**
  P. Vivet, F. Clermidy, D. Lattard
  CEA-LETI France
- **“Exploration of CDMA-Based Network-on-Chip”**
  X. Wang and J. Nurmi
  University of Technology, Finland

A.2 NoC Development Support: Synthesis, Simulation, Verification methods, tools:

- **“Network-on-Chip Generation and Optimization Tool”**
  T. Ahonen, H. Bin and J. Nurmi
  Tampere University of Technology, Finland
- **“CHAINworksTM – an EDA tool suite and IP libraries for self-timed interconnect design synthesis”**
  J. Bainbridge, A. Bardsley and R. McGuffin,
  Silistix
- **“µSpider NoC Road Map”**
  S. Evain[1], J. Diguet[1] and D. Houzet[2]
  1 University of South Brittany, France
  2 INSA / IETR, Rennes Cedex, France
- **“Modeling and synthesis of Asynchronous Network on Chip using SystemC”**
  C. Koch-Hofer and M. Renaudin
  TIMA, France
- **“SystemC Simulation Model of a Flexible Network-on-Chip”**
  S. Määttä and J. Nurmi
  Tampere University of Technology, Finland
- **“Methodology, benefits and tools for NoC topology exploration”**
  P. Martin and J. Lecler
  Arteris SA
- **“NoC Performance Optimization via Long-range Link Insertion”**
  U. Ogras and R. Marculescu
  Carnegie Mellon University, USA
- **“Constraint-Driven Communication Synthesis of Networks On-Chip”**
  A. Pinto, L. Carloni, and A. Sangiovanni-Vincentelli
  U.C. Berkeley, USA
  Columbia University, USA

- **“HW/SW Partitioning and Interface Synthesis in NoCs”**
  F. Regazzoni and M. Lajolo
  ALaRI - USI, Lugano, Switzerland
  NEC Laboratories America, Princeton, NJ, USA
- **“Introduction to the Princeton Polaris Project”**
  V. Soteriou, N. Eisley, H. Wang, B. Li and L. Peh
  Princeton University, USA
- **“A Flexible System Level Design Methodology Applied to NoC”**
  A. Vander Biest, A. Leroy and F. Robert
  ULB, Belgique

A.3 QoS Issues: Results on Performance Analysis, Optimisation, Simulation:

- **“QoS Oriented Configurable Network on Chip”**
  M. Al Faruque, X. Ye, G. Weiss and J. Henkel
  University of Karlsruhe, Germany
- **“QNoC: QoS Network on Chip”**

  E. Bolotin, I.Cidon, R. Ginosar and A.Kolodny
  Electrical Engineering Department, Technion, Haifa, Israel
- **“Efficient Link Capacity and QoS Design for Network-on-Chip”**
  Z. Guz, I. Walter, E. Bolotin, I. Cidon, R. Ginosar and A. Kolodny
  Electrical Engineering Department, Technion, Haifa, Israel
- **“Empirical Optimization of Multiprocessor Systems and their Heterogeneous Communication Fabrics”**
  G. Hellestrand, J. Torossian and C. Alford
  VaST Systems Technology Corporation, USA
- **“On Options for Accessing Regions in NoC”**
  R. Holsmark and S. Kumar
  Jönköping University, Sweden
- **“Spatial Division Multiplexing: A Novel Approach for Guaranteed Throughput on NoCs”**
  A. Leroy[1/2], P.Marchal[1], F. Robert[2] and F. Catthoor[1/3]
  1 IMEC Belgium
  2 VUB Belgium
  3 KUL Belgium
- **“Evaluation of On-chip Networks Using *Deflection* Routing”**
  Z. Lu, M. Zhong and A. Jantsch
  Royal Institute of Technology, Sweden
- **“Dynamic Time-Slot Allocation for Networks on Chip with TDMA QoS”**
  T. Marescaux[1], B. Brické[1/2], P. Debacker[1/2], V. Nollet[1] 1, H. Corporaal[3]
  1 IMEC V.Z.W., Belgium
  2 Katholieke Universiteit Leuven, Belgium
  3 Technical University Eindhoven (TU/e), The Netherlands

- **“Application Parallelism Exploitation using NoC with Multithreaded Processors”**
  R. Pop and S. Kumar
  Jönköping University, Sweden
- **“Curing Hotspots in Wormhole NoCs”**
  I. Walter, I. Cidon, R. Ginosar and A. Kolodny
  Electrical Engineering Department, Technion, Haifa, Israel

A.4 Routing methods:

- **“STAR: An Efficient Routing Strategy for NoC with Mixed QoS Requirements”**
  D. Andreasson and S. Kumar
  Jönköping University, Sweden
- **“Assessing Adaptive Routing Behaviour in NoCs Through Cycle-Accurate Functional Simulation**
  F. Angiolini[1], L. Lerinz[2],D. Bertozzi[2],L. Benini[1]
  1 DEIS, University of Bologna, 40136 Bologna, Italy
  2 ENDIF, University of Ferrara, 44100 Ferrara, Italy
- **“Performance Enhancement through Early Release and Buffer Optimization in Network-on-Chip Router Architectures”**
  J. Kim, D. Park, C. Nicopoulos, N. Vijaykrishnan and C. Das
  The Pennsylvania State University, USA

A.5 NoC Applications:

- **“Emerging Networks on Chip: Lessons Learned to Win the Future”**
  Iyad Al Khatib
  Royal Institute of Technology, Sweden
- **“Interoperability Protocol and Algorithms for Network-on-Chip Autonomous-System Communications for the Next Generation Biomedical Sensor-Networks”**
  Iyad Al Khatib[1], A. Jantsch[1], R. Nabiev[2] and L. Onana Alima[3]
  1 Royal Institute of Technology, Sweden
  2 Karolinska University Hospital Huddinge, Sweden
  3 Université de Mons-Hainaut, Belgium
- **“EEG Biomedical Sensors Network on Chip (EBSNOC)”**
  M. Saleh and A. Jantsch
  Royal Institute of Technology (KTH), Sweden
- **“Design of a simple clockless Network-on-Chip for an audio DSP chip”**
  M. Stensgaard[1], T. Bjerregaard[1], J. Sparsø[1] and J. Pedersen[2]
  1 Technical University of Denmark(DTU)
  2 William Demant Holding

- **“Application-Specific Networks-on-Chips”**
  J. Xu and W. Wolf
  Princeton University, USA

Track B: Interconnection Technology, NoC Signalling Schemes, Components, Infrastructures (Power, Testing, Fault-tolerance)

B.1 Interconnection Technology: NanoCMOS, Optical, PICMOS

- **"Extended Global Interconnect Architecture for Nano-CMOS Technologies",**
  J.Balachandran[1], S.Brebels[1], G.Carchon[1], M.Kuijk[2], W.De Raedt[1], B.Nauwelaers[3] and E.Beyne[1],
  1 IMEC Belgium
  2 VUB Belgium
  3 KUL Belgium
- **"Novel Performance of Three-dimensional (3D) On-chip Crossbar Bus using non-Silicon Transistors"**
  S. Fujita[1/2], K. Abe[1], K. Nomura[1] and T. Lee[3],
  1 Frontier Research Laboratory[1], Toshiba Corporation, Komukai-Toshiba-cho, Saiwai-ku, Kawasaki, Japan
  2 Toshiba America Research, USA
  3 Stanford University, USA
- **"Reconfigurable Optical Networks for On-Chip Multiprocessors"**
  W. Heirman[1], J. Dambre[1], I. O'Connor[2], J. Van Campenhout[1]
  1 ELIS, Ghent University, Belgium;
  2 LEOM, Ecole Centrale de Lyon, France
- **"Optical interconnect for on-chip data communication"**
  I. O'Connor[1], F. Tissafi-Drissi[1], D. Navarro[1], F. Mieyeville[1], F. Gaffiot[1], J. Dambre[2], M. De Wilde[2], D. Stroobandt[2] and D. Van Thourhout[3]
  1 Ecole Centrale de Lyon, France
  2 Ghent University, Belgium
  3 IMEC/Ghent University, Belgium
- **"Heterogeneous integration of light sources on an SOI waveguide platform for photonic interconnects on CMOS (PICMOS)"**
  G. Roelkens[1], J. Van Campenhout[1], D. Van Thourhout[1], R. Baets[1], P. Rojo-Romeo[2], C. Seassal[2], P. Regreny[2], P. Viktorovitch[2], I. O'Connor[2], L. Di Cioccio[3] and J. Fedeli[3]
  1 Ghent University Belgium
  2 Ecole Centrale de Lyon France
  3 CEA-DRT/LETI Franc

B.2 Signalling and Components: Transmission Modes, Protocols, Encoding, Timing and Synchronisation Issues, Components for NoC

- **"A 20Gb/s Transceiver for Network-On-Chip"**
  Z. Asgar, J. Zou, P. Jain, R. Kamath and R. Harjani
  University of Minnesota, USA
- **"Buffer output capacitance effect on NoC line drivers performance"**
  G. Cappuccino, A.Pugliese, G. Cocorullo
  DEIS-University of Calabria, Italy
- **"Multiple-rail phase-encoding for NoC"**
  C. D'Alessandro[1], D. Shang[1], A. Bystrov[1], A. Yakovlev[1], O. Maevsky [2]
  1 University of Newcastle upon Tyne, UK
  2 Intel Labs, Moscow
- **"High-Speed Serial Interconnect for NoC "**
  Rostislav (Reuven) Dobkin, Ran Ginosar and Avinoam Kolodny
  Electrical Engineering Department, Technion, Haifa, Israel
- **"Dynamic Multi-grain Pipelined Interconnect"**
  M. Imai, T. Azuma, K. Watanabe, M. Kondo, H. Nakamura and T. Nanya
  Research Center for Advanced Science and Technology, The University of Tokyo
- **"Clock Synchronization in NoCs using Distributed FIFOs"**
  P. Pande and J. Nyathi
  Washington State University, USA
- **"A Distributed, Interleaved FIFO for SoC Interconnect"**
  S. Sood, M. Greenstreet and R. Saleh
  University of British Columbia
- **"Design of fast and reliable synchronisers for NoCs "**
  J. Zhou, D. Kinniment, G. Russell, and A. Yakovlev
  University of Newcastle upon Tyne, UK

B.3 NoC Infrastructures: NoC Testing, Fault-tolerance, Power management, Testing

- **"An improved method for delay fault testing of NoC interconnections"**
  T. Bengtsson[1], S. Kumar[1], A. Jutman[2] and R. Ubar[2]
  1 Jönköping University, Sweden
  2 Tallinn University of Technology, Estonia
- **"A dedicated communication layer for Power Management of NoC based SoCs"**
  N. Genko[1], D. Atienza[1/2], A. Aquaviva[1] and G. De Micheli[1]
  1 EPFL, Switzerland
  2 DACYA/UCM, Spain.
- **"BIST for NoC Interconnects"**
  C. Grecu[1], P. Pande[2], A. Ivanov[1] and R. Saleh[1]
  1 University of British Columbia, Canada;
  2 Washington State University, USA
- **"Fault-Diagnosis-And-Repair System for Improving the Fault-Tolerance and Manufacturability of MPSoCs"**
  H. Kariniemi and J. Nurmi
  University of Technology, Finland
- **"Versatile XGFT Network-On-Chip with Improved Fault-Tolerance for Multi-Processor Systems-on-Chip "**
  H, Kariniemi and J. Nurmi
  University of Technology, Finland
- **"Utilizing NoC Switches as BIST-structures in 2D Mesh Network-on-Chip"**
  K. Petersén and J. Öberg
  Royal Institute of Technology (KTH), Sweden

A.1 NoC General: Architectures, Platforms, Designs, Prototyping:

- **"A Hierarchical Approach to Network-on-Chip"**
  T. Ahonen and J. Nurmi
  Tampere University of Technology, Finland
- **"NoC-Based Platform Implementation on FPGA"**
  T. Ahonen, J. Kylliäinen, C. Brunelli and J. Nurmi
  University of Technology, Finland
- **"PROTEO – A Flexible Network-on-Chip Scheme"**
  T. Ahonen[1], D. Sigüenza-Tortosa[1] and J. Nurmi[1]
  1 Tampere University of Technology, Finland
  2 Complutense University of Madrid, Spain
- **"Towards a Communication-Centric Design Methodology"**
  A. Banerjee, R. Francis, J. Lee, J. May, S. W. Moore and R. D. Mullins
  Computer Laboratory, University of Cambridge
- **"Clockless On-Chip Networks and MANGO"**
  T. Bjerregaard, J. Sparsø and M. Stensgaard
  Technical University of Denmark, Denmark
- **"STNoC™: An Evolution Towards MPSoC Era"**
  M. Coppola[1], C. Pistritto[2], R. Locatelli[1] and Alberto Scndurra[2]
  [1] STM, AST Grenoble Lab, France
  [2] STM, HPC, Italy
- **"Introduction to KAIST BONE project"**
  K. Lee, S. Lee, D. Kim, K. Kim, J. Kim and H. Yoo
  Korea Advanced Institute of Science and Technology (KAIST), Korea
- **"Synchronous Latency Insensitive Design in Æthereal NoC "**
  X. Ru[1], J. Dielissen[2], C. Svensson[3] and K. Goossens[2]
  1 Philips Semiconductors, Switzerland
  2 Philips Research, The Netherlands
  3 Linköping University, Sweden
- **"A Lightweight NoC for the NOVA Packet Processing Platform"**
  C. Sauer[1], M. Gries[1], S. Dirk[1], J.-C. Niemann[2], M. Porrmann[2], U. Rückert[2]
  1 Infineon Technologies, Access Communications, Munich
  2 HNI Paderborn

- **“Interconnects and network on chip at Linköping University”**
  C. Svensson, Fellow IEEE,
  Dept. Electrical Engineering, Linköping University, Sweden
- **“FAUST, an Asynchronous Network-on-Chip based Architecture for Telecom Applications”**
  P. Vivet, F. Clermidy, D. Lattard
  CEA-LETI France
- **“Exploration of CDMA-Based Network-on-Chip”**
  X. Wang and J. Nurmi
  University of Technology, Finland

A Hierarchical Approach to Network-on-Chip

Tapani Ahonen and Jari Nurmi
Tampere University of Technology, P.O. Box 553, FIN-33101 Tampere, Finland
[tapani.ahonen, jari.nurmi]@tut.fi

Introduction

The latest branch of our on-chip interconnect research aims at creating a guaranteed service low overhead heterogeneous multiprocessor Network-on-Chip (NoC) using a hierarchical approach. We focus especially on the efficiency of the memory to memory transfer mechanism, since minimizing cache latency is essential for multiprocessor systems. The network at the global level of an FPGA prototype system being developed is depicted in figure 1 below. At this level, all blocks are equal that is they are capable of initiating transactions as well as responding to them. Figure 2 illustrates the local structure of the proposed hierarchical multiprocessor network. The leaf level of this memory-mapped split-transaction NoC has two initiators and an unlimited number of slave devices behaving as targets. The two initiators are the local master (processor or similar device) and the interface to the global level of the NoC, through which the rest of the initiators in the system make their requests. Access latency and bandwidth are controlled using a programmable time division multiple access (TDMA) scheme. The TDMA logic pipelines the transactions with no wait cycles.

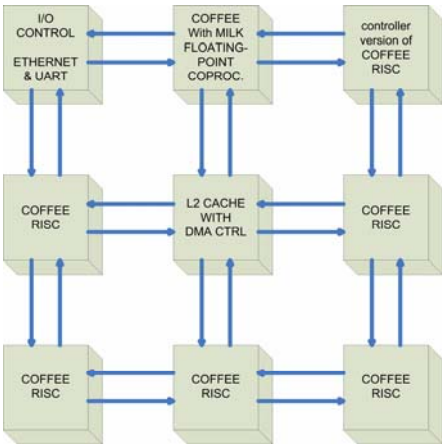

Fig. 1. Global level of an FPGA prototype.

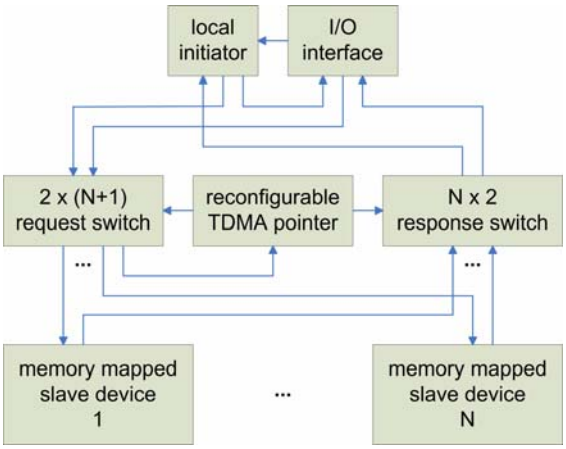

Fig. 2. Generic view of the local structure.

Discussion

Memory bandwidth forms a serious bottleneck in processor-based systems. Shared buses lengthen memory access cycles narrowing the potential bandwidth and require strong, power hungry drivers. These wasteful features are especially prominent in FPGAs where memories operate at higher frequencies than conventional processor designs and buses get spread over a large area. Our approach addresses this situation by enabling higher operating speed and simultaneous global and local access. The structure in figure 2 assumes single-channel slave devices. The switch design can be further simplified through the adoption of dual-channel devices such as dual-port memories readily available in FPGAs. Then the request and response switches would be split into two serving the initiators in parallel without the need for arbitration. The drawback with this variation is that if single channel devices exist, they can only be accessed through the initiator they are dedicated to.

NoC-Based Platform Implementation on FPGA

Tapani Ahonen, Juha Kylliäinen, Claudio Brunelli and Jari Nurmi
Tampere University of Technology, P.O. Box 553, FIN-33101 Tampere, Finland
[tapani.ahonen, juha.p.kylliainen, claudio.brunelli, jari.nurmi]@tut.fi

Introduction

We have developed a multimedia processing platform [1] using previously designed IP components. These components include the Proteo network-on-chip, the Coffee processor, the Milk floating-point coprocessor, and the transport triggered TACO for protocol processing. The main design goal was to enable efficient utilization of the communication resources through the bus-oriented standard interfaces used. Unlike shared buses, networks-on-chip support varying levels of communication parallelism depending on the topology. This design case illustrates the need to have well matching network topology, interfaces, and computation models.

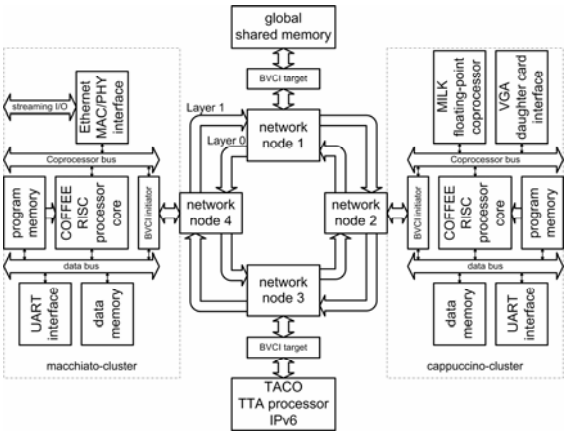

Fig. 1. FPGA prototype of the NoC-based multimedia processing platform.

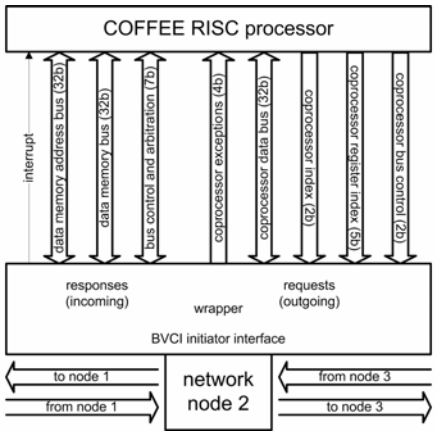

Fig. 2. NoC interface for the Coffee core.

Discussion

Figures 1 and 2 illustrate the prototype structure on FPGA. The sequential processor IPs as network hosts present limited absorption and injection capabilities. The desired I/O parallelism could be added to the TTA architecture by using additional internal buses or by applying the NoC paradigm within the TTA, which could be an interesting topic for future research. In a NoC-based system the division to master and slave devices assumed by current interface standards results in control overhead and favors a shared memory communication model between the masters. Unveiling the full potential of NoC-based systems requires not only an efficient network but also processing architectures and interfaces suitable for the topology.

References

[1] T. Ahonen, and J. Nurmi, “Integration of a NoC-based multimedia processing platform”, In Proc. *2005 International Conference on Field Programmable Logic and Applications* (FPL’05), pages 606-611, Tampere, Finland, 24-26 August 2005.

PROTEO – A Flexible Network-on-Chip Scheme

Tapani Ahonen, David Sigüenza-Tortosa * and Jari Nurmi
Tampere University of Technology, P.O. Box 553, FIN-33101 Tampere, Finland
* Complutense University of Madrid, Madrid, Spain
[tapani.ahonen, jari.nurmi]@tut.fi, *sigtor@fdi.ucm.es

Introduction

Proteo [1] is a Network-on-Chip (NoC) model that fulfills most future requirements while avoiding major overheads. The architecture compares favorably to most NoCs in terms of simplicity and economy. There are two different paradigms for NoC: one emphasizing run-time programmability and another proposing simpler NoCs closely tied to EDA tools. Proteo can be counted with the latter. We propose a development framework in which methodology and software tools help embed all the knowledge about the application and communication requirements in the network design.

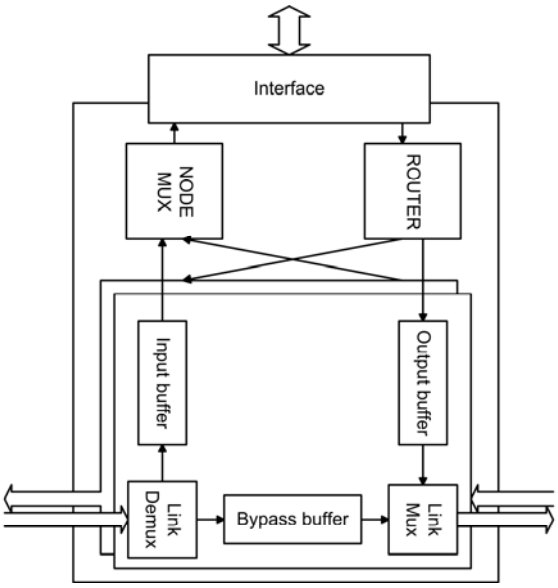

Fig. 1. Proteo node architecture with two I/O links, as used in a bi-directional ring.

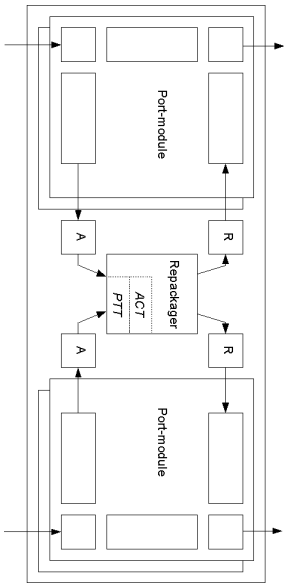

Fig. 2. Bridge structure for connecting two bi-directional rings.

Discussion

Dynamic routing and adaptability require computing power and/or extra storage. Besides, sophisticated protocols must be defined to deal with potential unstability, which add to the overall complexity. A simple NoC like Proteo can offer higher performance than more complex architectures. On the other hand, in order to keep this advantage, we have had to sacrifice the ability to give a high level of QoS support and programmability.

References

[1] D. Sigüenza-Tortosa, T. Ahonen, and J. Nurmi, "Issues in the development of a practical NoC: the Proteo concept", *INTEGRATION, the VLSI journal*, volume 38, pages 95-105, December 2004.

Towards a Communication-Centric Design Methodology

A. Banerjee, R. Francis, J. Lee, J. May, S. W. Moore and R. D. Mullins
Computer Laboratory, University of Cambridge
Simon.Moore@cl.cam.ac.uk

The design of high performance computing devices has reached a major turning point. The next 20 years will require a step change in computer architectures and the way in which we approach many VLSI design tasks. In contrast to the past two decades, the scaling of CMOS fabrication technologies will no longer provide a simple route to improved performance. Central to many of these challenges is the ability to communicate on-chip in a low-power and robust fashion. In fact, technology scaling trends suggest that communication, not computation, will dominate delay, area and power budgets. Communication is also central to supporting an increasingly distributed design style where many heterogeneous or homogeneous IP blocks are integrated on a single chip to create a complete system. In many cases it will be the scheduling and management of these compute resources and the provision of the necessary communication resources that will be the major design and implementation challenge.

Our work to date has focused on the design and implementation of high-performance on-chip networks [3, 4]. Our latest test chip, Lochside, implements a 4x4 mesh network of low-latency virtual-channel routers. Our novel speculative router architecture permits each network hop (router and link) to be traversed in a single clock cycle.

The characteristics of these new *communication-centric* VLSI architectures also forces the choice of system-timing regime to be carefully reevaluated. The interconnection of many different, physically distributed, IP blocks operating at different or even adaptive clock frequencies poses significant challenges for existing synchronous design techniques. Our work has suggested new techniques for generating and distributing clocks while minimising power and skew [1]. Work has also explored the use of local clocks and event-driven synchronous systems [2]. Efficient interconnects may also be constructed using purely asynchronous design techniques.

The shift to communication-centric architectures involves much more than simply the replacement of on-chip buses with a scalable packet-switched interconnect. Future work is set to examine how communication-centric techniques can serve as a basis for reducing system complexity and cost, managaging local power and thermal budgets, improving reliability and manufacturability and boosting performance.

[1] S. Fairbanks and S. Moore. High precision timing signals using asynchronous control rings,. In *10th International Symposium on Asynchronous Circuits*, 2004.

[2] R. D. Mullins. Asynchronous versus synchronous design techniques for NoCs. Tutorial at the International Symposium on System-on-Chip, 2005.

[3] R. D. Mullins, A. F. West, and S. W. Moore. Low-Latency Virtual-Channel Routers for On-Chip Networks. In *Proceedings of the 31st Annual International Symposium on Computer Architecture (ISCA)*, 2004.

[4] R. D. Mullins, A. F. West, and S. W. Moore. The design and implementation of a low-latency on-chip network. In *Proceedings of the 11th Asia and South Pacific Design Automation Conference (ASP-DAC)*, 2006.

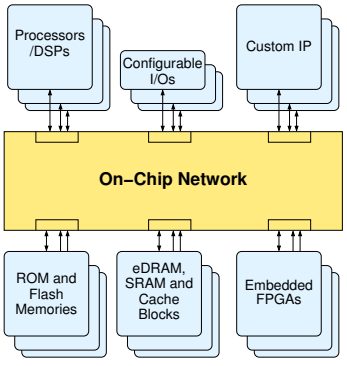

Figure 1. A Flexible SoC platform

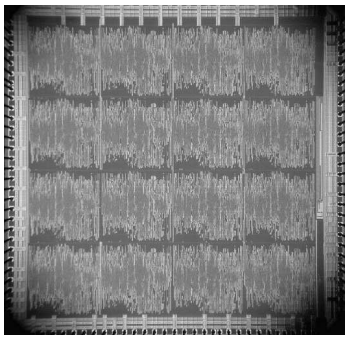

Figure 2. Lochside Die Micrograph

Clockless On-Chip Networks and MANGO

Tobias Bjerregaard, Jens Sparsø and Mikkel Bystrup Stensgaard
Informatics and Mathematical Modelling, Technical University of Denmark (DTU)
{tob, jsp, mibs}@imm.dtu.dk

Abstract

The demand for IP reuse and system level scalability in System-on-Chip (SoC) designs is growing. Networks-on-chip (NoC) constitute a viable solution space to leveraging these demands. In this poster we describe advantages of using clockless circuit techniques for the implementation of NoC architectures. We also explain key features of the MANGO clockless NoC that we have developed.

Advantages of using clockless circuit techniques for NoC:

(i) Micro-level flow control. The flow of data in clockless systems is handled by local handshaking mechanisms. Flow control is hence an integral part of the design methodology.

(ii) Zero dynamic idle power consumption. The local handshaking causes latches and flip flops to be acticated only when there is data to be moved. This corresponds to ultra fine-grained clock gating.

(iii) Low forward latency. Pipeline stages are activated as soon as input data becomes available, not needing to await a discretely timed clock signal. Hence forward latency is lower than in a clocked system. Also, in multi clock domain systems, only a single synchronization event is required per end-to-end path: crossing from the clockless domain into a clocked one.

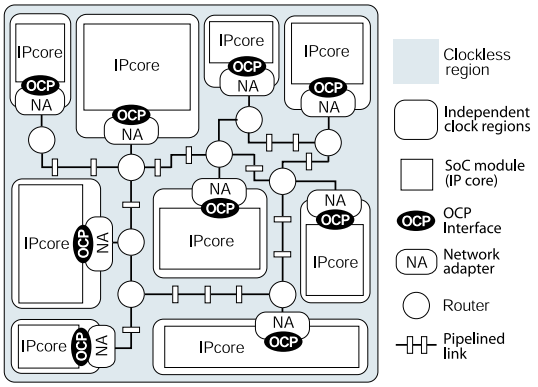

Figure 1. A MANGO-based SoC is a heterogeneous system of independently clocked IP blocks.

(iv) Always operates at at maximum speed. Since no clock is used, clockless circuits can go from idle to maximum speed instantly. This is especially advantageous during periods of peak communication.

(v) Globally asynchronous operation. Clockless systems facilitate an inherent timing-wise decoupling of IP blocks. Delay insensitive implementations are possible, increasing timing robustness further. Link level synchronizers are avoided, as synchronization is needed only at the network borders.

MANGO (Message-passing Asynchronous Network-on-Chip providing Guaranteed services over OCP interfaces), is a packet-switched on-chip network being developed at DTU [1, 2, 3]. Key features of MANGO, which help leverage a modular and scalable SoC design flow, are:

(i) Clockless implementation. Links and routers are constructed using asynchronous (clockless) circuits. This enables a Globally Asynchronous Locally Synchronous (GALS) system.

(ii) Guaranteed communication services. Connections in the network provide hard bounds on bandwidth and latency. This facilitates real-time systems, and makes it feasible to verify systems analytically rather than through simulation.

(iii) Standard OCP socket access points. Access points in the MANGO adhere to the Open Core Protocol (OCP).

References

[1] T. Bjerregaard, S. Mahadevan, R. G. Olsen, and J. Sparsø. An OCP compliant network adapter for GALS-based SoC design using the MANGO network-on-chip. In *Proceedings of International Symposium on System-on-Chip 2005*. IEEE, 2005.

[2] T. Bjerregaard and J. Sparsø. A router architecture for connection-oriented service guarantees in the MANGO clockless network-on-chip. In *Proceedings of Design, Automation and Testing in Europe Conference 2005 (DATE05)*. IEEE, 2005.

[3] T. Bjerregaard and J. Sparsø. A scheduling discipline for latency and bandwidth guarantees in asynchronous network-on-chip. In *Proceedings of the 11th IEEE International Symposium on Advanced Research in Asynchronous Circuits and Systems*. IEEE, 2005.

STNoC™: An Evolution Towards MPSoC Era

Marcello COPPOLA[1], Carlo PISTRITTO[2], Riccardo LOCATELLI[1] and Alberto SCANDURRA[2]
[1]STM, AST Grenoble Lab, 12 Jules Horowitz 38019 Grenoble, France
{marcello.coppola, riccardo.locatelli}@st.com
[2] STM, HPC, Stradale Primo Sole 50, 95121 Catania, Italy
{carlo.pistritto, alberto.scandurra}@st.com

For many years, System-on-Chip (SoC) technology has been instrumental in delivering the ever-increasing price/performance ratios that enabled the creation and development of advanced products. At the heart of SoC technology is the CMOS integration increasing capability, expressed by Moore's law, which provides exciting opportunities to develop new applications and market scenarios. However, it also presents severe challenges for the design community, which must exploit this increasing complexity with no impact on design times. One of the key issues in the evolution of MPSoC is the nature of the on-chip communications architecture. In today's SoC devices, all the IPs are connected by global on-chip buses which are increasingly dominating the delay, power and area of the device. There is a strong industry consensus that a new on-chip communications architecture is needed in the future.

The STNoC™ is a state-of-the-art low-cost on-chip interconnect that will play a vital role in enabling multiprocessor system-on-chip by providing structure, performance, and modularity. The STNoC topology (called Spidergon) is based on three basic components: a Network Interface providing modularity through uniform network-on-chip access from any IP subsystem (e.g. processing or storage element or FPGA), a high performance wormhole Router with reduced buffering and a physical communication link. The Spidergon NoC topology is a regular chordal ring with vertex-transitivity. Thus, all nodes have global knowledge of the network, providing for simple, local shortest-path routing and scheduling based on virtual circuits. In addition, for (SoC) network sizes, Spidergon topology is a low-cost tradeoff compared to Mesh or Torus, providing competitive performance metrics with respect to number of links, diameter, average distance, size granularity (just 2), and embedding properties for mapping MPSoC application traffic.

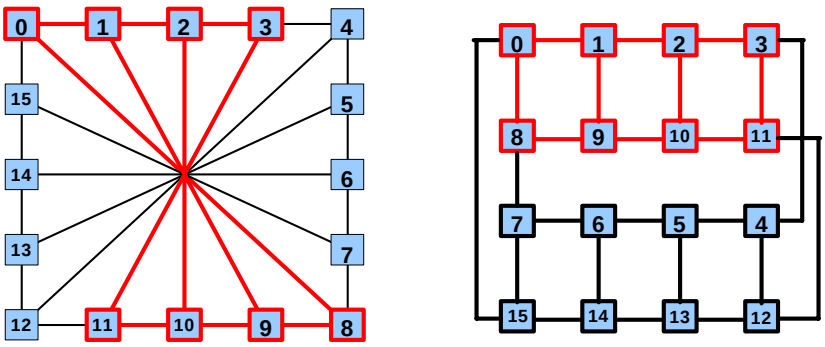

Figure 1: Spidergon's conceptual elegance translates to low-cost silicon implementation

On top of that, STNoC™ has got low implementation and maintenance costs thanks to a simple architecture. It is important to note that increasing transistor density, higher operating frequencies, short time-to-market and reduced product life cycles will characterize the future semiconductor industry scenario. Thanks to this cost/performance scalability, STNoC™ provides a full roadmap from current to future needs, providing an available solution to today's SoC integration issues and enabling the next generations of complex architecture integration.

Introduction to KAIST BONE project

Kangmin Lee, Se-Joong Lee, Donghyun Kim, Kwanho Kim, Joo-Young Kim, and Hoi-Jun Yoo
Korea Advanced Institute of Science and Technology (KAIST), Daejeon, Republic of Korea
(Presentation Author: kangmin@eeinfo.kaist.ac.kr)

Abstract

For the unique purpose of realizing the new NoC technology through implementation, the BONE (Basic On-chip Network) project was launched in 2002 at KAIST. As the results of the project, new NoC techniques and the implementations have been published and demonstrated in every year as summarized in Fig. 1. The first prototype, BONE-1, proposed plesiochronous communications and on-chip serial wires on a star topology. BONE-2 demonstrated various low-power techniques on a physical layer such as low-swing signaling, crossbar partial activation, serial encoding and frequency scaling on a multiprocessor SoC. BONE-3 devised high-performance techniques such as serial wave-front-train signaling, adaptive bandwidth control, programmable synchronizer and adaptive circuit/packet switching technique. Recently, BONE-4 presented on-chip traffic monitoring system on a FPGA NoC platform. In this year, we will focus on a reconfigurable and programmable NoC as a BONE-5 project.

This project covers circuit level design, architectural researches and system integration on a NoC platform.

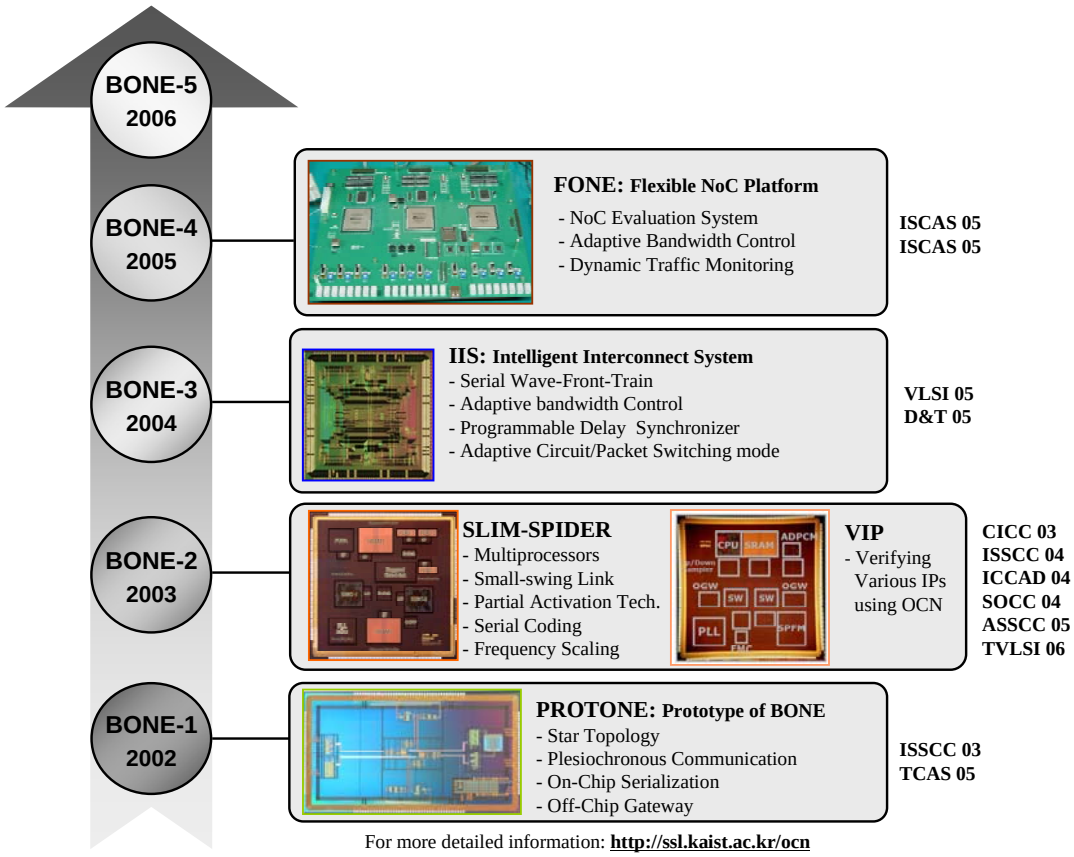

Figure 1 BONE series roadmap since 2002

Synchronous Latency Insensitive Design in Æthereal NoC

Xiao Ru[^], John Dielissen[#], Christer Svensson[*] *Fellow IEEE*, Kees Goossens[#]

[*] Presenting Author, Linköping University, Department of Electrical Engineering, SE-58183 Linköping, Sweden (*christer@isy.liu.se*)
[#] Philips Research, Prof. Holstlaan 4, 5656 AA Eindhoven, The Netherlands *({John.Dielissen, Kees.Goossens}@philips.com)*
^Philips Semiconductors, Binzstrasse 44, CH - 8045 Zürich, Switzerland (*Xiao.Ru@philips.com*)

Abstract

In this paper the Synchronous Latency Insensitive Design [1] (SLID) method to mitigate timing problems due to global wire delays in Æthereal NoC is proposed. This method follows closely a fully synchronous design flow and utilizes only true digital library elements, which does not change the original design flow and libraries of Æthereal. A few clock cycles latency is inserted to every connection between different functional units. This latency is used to automatically absorb unknown global interconnect delays, unknown global clock skews and other timing uncertainties. The introduced SLID method may substantially cut down the timing closure effort in Æthereal NoC system design, and other large scale, high frequency digital designs carried out in Deep Sub-Micron (DSM) technologies.

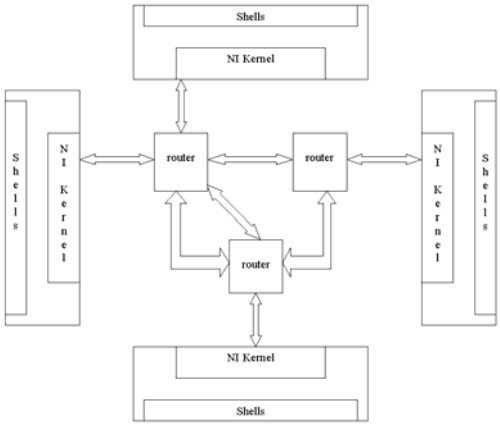

Figure 1: Æthereal NoC Example

Æthereal [2] Network on Chip (NoC) (See Figure 1) is an advanced intercommunication solution to large scale System on Chip (SoC). It plays a central role in integrating IPs with diverse communication requirements. Æthereal and other NoCs usually have long interconnects in their design. A general disadvantage of long interconnects is that they introduce uncertainties like transmission delay over buses or global clock skews. It has been recognized that integrated circuits in DSM technologies exhibit increased timing problems due to increased clock frequencies, increased complexity and increased wire delays. These problems manifest themselves as a severe increase in verification cost (timing closure), increased problems to scale up the clock frequency and increased effort for clock distribution.

This paper proposes the SLID scheme to ease the timing problems in Æthereal NoC. The main goal is to keep its fully synchronous design flow even for large high-speed designs in DSM processes. The idea is to manage the inevitable wire delays already at architecture level, and then guarantee that the functional description at this level is valid all the way to layout.

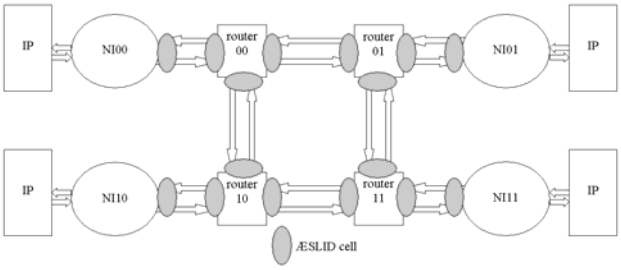

Figure 2: SLID scheme implemented 2*2 Æthereal NoC

The implementation of the SLID scheme (See Figure 2) in Æthereal is based on the partition of a large design into isochronous blocks (e.g. routers, NIs), still keeping global synchronism. Each router/NI should be small enough not to exhibit severe wire delays. Between the routers and NIs an extra delay (pipelining) is inserted. This extra delay is later applied to automatically mitigate unknown wire delays and clock skews. Thus the proposed SLID scheme simplifies timing closure and relaxes clock distribution constraints of Æthereal NoC.

[1] A. Edman and C. Svensson, "Timing Closure through a Globally Synchronous, Timing Partitioned Design Methodology", Proc 41st Design Automation Conference, pp. 71-74, June 2004
[2] John Dielissen, Andrei Rădulescu, Kees Goossens and Edwin Rijpkema, “Concepts and Implementation of the Philips Network-on-Chip”. IP-Based SOC Design, Grenoble, November 2003

A Lightweight NoC for the NOVA Packet Processing Platform

HEINZ NIXDORF INSTITUTE
University of Paderborn
System and Circuit Technology
Prof. Dr.-Ing. Ulrich Rückert

C. Sauer[1], M. Gries[1], S. Dirk[1], J.-C. Niemann[2], M. Porrmann[2], U. Rückert[2]

[1] Infineon Technologies, Access Communications, Munich [2] Heinz Nixdorf Institute, Paderborn

Motivation & Goal

- Modular platform construction kit to investigate design criteria, such as flexibility, programmability, area, and performance
- Reuse of of-the-shelf components (e.g., programmable embedded cores) and deployment tools where possible
- Refinements (instruction set, co-processors, etc.) where needed as determined by profiling of reference applications
- Synergy: Hardware platform directly supports message passing semantics of application; NoC is a natural choice

Application Domain

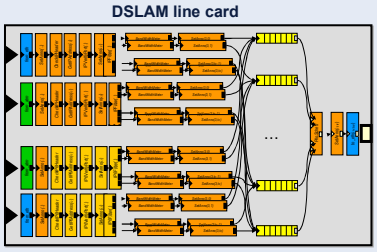

- Packet processing in access networks
- Data-flow driven processing

- Message passing semantics between computational kernels
- Subject to **tight constraints on costs and performance**
- Flexibility required to support broad variety of protocols and customer requirements

Network-Optimized Versatile Architecture (NOVA) Platform: Concepts

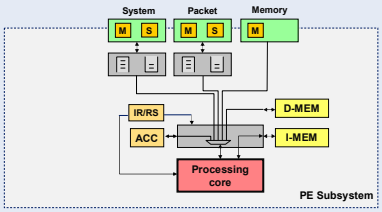

3

Systematic development of platform building blocks

- Profiling of embedded general-purpose cores (and their compilers) and specialized packet processing engines
- Packet processing engines provide high performance; general purpose cores are flexible and have mature compilers

⇒ Provides a first estimate of required parallelism

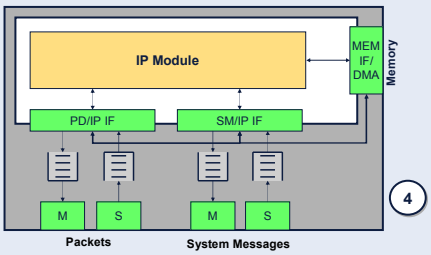

4

NOVA Socket concept

- Homogeneous interface
- Separation of IP-specific interfaces from interconnect and memory interfaces

1

Modular platform

- Number and type of processing cores
- Number and type of co-processors
- Number and type of I/O interfaces
- Heterogeneous memory hierarchy
- Communication architecture

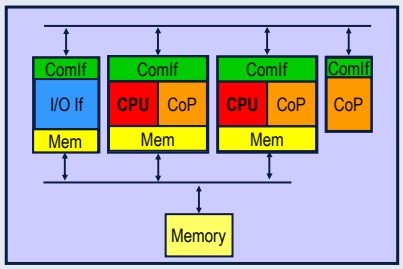

Efficient implementation

- Two NoC interfaces, buffering 16 64B messages each, need less than 1/3 of the CPU subsystem area

5

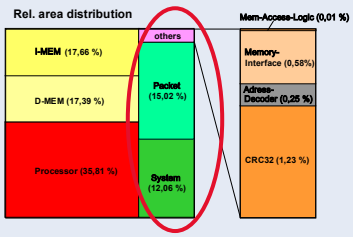

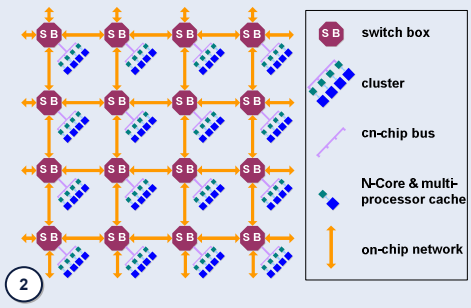

2

GigaNoC – hierarchical Network-on-chip

- Bus-based network for small clusters (< 10 components)
- Switch boxes for connecting clusters

NoC principles

- Message passing interface with priorities: computation can continue while message is transmitted to destination
- Incorporation of flow control by backpressure signaling to avoid overload and loss
- Lightweight one-word header, which contains destination, message type, and context ID for fast flow-through processing

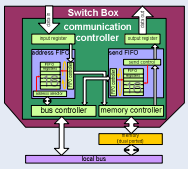

SoC main components	Area [mm²] 130nm	Area [mm²] 90nm	Frequency [MHz] 130nm	Frequency [MHz] 90nm
32 Cores [N-Core]	32 x 0.16	32 x 0.12	205	285
8 switch-boxes [with 5 ports]	8 x 1.129	8 x 0.53	560	650
32 local RAMs, (32 KB) + 8 local packet buffers (2 x 16 KB)	32 x 0.875 + 8 x 2 x 0.466	32 x 0.875 + 8 x 2 x 0.466	400	450
8 local on-chip busses	8 x 0.05	8 x 0.02	211	290
Total	50.01	43.7	205	285

Status/Next Steps

- Verification of approach and concepts
- Currently implementing initial NOVA platform prototype comprising approx. 20 message passing clients
- Application-driven analysis using system-level benchmark
- Mapping IP-DSLAM reference application to the platform
- Enables detailed quantitative exploration of design trade-offs
- Demonstrator at CeBIT 2006 trade fair
- Evaluation of GigaNoC for larger systems [2]

Partners/Funding

Infineon Technologies
- Communications, Access System Engineering (COM AC SE NP)
- 81726 Munich, Germany
- E-mail: Christian.Sauer@infineon.com

Heinz Nixdorf Institute
- University of Paderborn, System and Circuit Technology
- Fürstenallee 11, 33102 Paderborn, Germany
- E-mail: niemann@hni.upb.de

This project is supported, in parts, by the research grants GigaNetIC and PlaNetS of the German government (BMBF).

References:
[1] C. Sauer, M. Gries, S. Sonntag, "Modular Reference Implementation of an IP-DSLAM," 10th IEEE Symposium on Computers and Communications (ISCC 2005), Cartagena, Spain, Jun 2005
[2] J.-C. Niemann, M. Porrmann, C. Sauer, U. Rückert, "An Evaluation of the Scalable GigaNetIC Architecture for Access Networks," 2nd Advanced Networking and Communications Hardware Workshop (ANCHOR) along with ISCA 2005, Madison, Wisconsin, USA
[3] R. Eickhoff, J.-C. Niemann, M. Porrmann, U. Rückert, "Adaptable Switch boxes as on-chip routing nodes for networks-on-chip, " In: Rettberg, Zanella, Mauro, Rammig, "From Specification to Embedded Systems Application," IESS 2005. Manaus, Brazil

Interconnects and network on chip at Linköping University

Christer Svensson, Fellow IEEE, Dept. Electrical Engineering, Linköping University, Sweden (chs@isy.liu.se).

Abstract: Interconnects and network on chips has been studied extensively at Linköping University recent years. Four aspects of this theme, all related to emerging deep submicron problems related to system on chips, are briefly described in this poster.

We have studied limits to latency and data-rate of on-chip wires in deep-submicron technologies, and methods to reach these limits. We have shown that latency can be limited to velocity-of-light delay by utilizing upper-level metal for long distance data transport. We have further demonstrated that data-rate is limited by wire dispersion (due to skin-effect) and crosstalk (due to variable delays), and shown how dispersion can be partly compensated for by overdriving the wire. Experiments have demonstrated 3 Gb/s data transport over 5mm on-chip wires with nearly velocity-of-light delays.

We have proposed a new method for managing long wire delays in synchronous designs, Synchronous Latency Independent Design (SLID). The method utilizes FIFO synchronizers, which latencies are preset to a value, large enough to absorb all unexpected wire delays and clock skews. The technique has been experimentally demonstrated at a data rate of 3Gb/s (Fig. 1).

We have proposed a new scheme for Networks on chip, SoCBUS, using a hybrid packet/circuit switching technique, Packet Connected Circuit (PCC). We have further investigated its applicability to various demanding applications, for example in 3G WCDMA/FDD basestations.

Finally, we study low power clock distribution, based on an on-chip adiabatically generated resonant clock. We have experimentally demonstrated 2.3x clock-power saving at 1.55GHz clock frequency for a core with 1792 flip-flops.

References

[1] P. Caputa and C. Svensson, "Well-behaved Global On-Chip Interconnect", IEEE Trans. On Circuits and Systems, vol. 52, p. 318, February 2005.

[2] P. Caputa and C. Svensson, "A 3Gb/s/wire Global On-Chip Bus with Near Velocity-of-Light Latency", VLSI Design 2006 Conference, Hyderabad, January 2006.

[3] R. Källsten, P. Caputa and C. Svensson, "Capacitive Crosstalk Effects on On-Chip Interconnect Latencies and Data-Rates", Proc. Of the Norchip Conference, p. 281, November 2005.

[4] A. Edman and C. Svensson, "Timing Closure through a Globally Synchronous, Timing Partitioned Design Methodology", Proc. Of the 2004 Design Automation Conference, p. 71, 2004.

[5] P. Caputa and C. Svensson, "An On-Chip Delay- and Skew-Insensitive Multi-Cycle Communication Scheme", International Solid-State Circuits Conference, San Francisco, February 2006.

[6] D. Wiklund and D. Liu, "SoCBUS: Switchced Network on Chip for Hard Real Time Systems", Proc. Of the International Parallel and Distributed Processing Symposium, Nice, April 2003.

[7] D. Wiklund and D. Liu, "Design, Mapping and Simulation of a 3G WCDMA/FDD Basestation Using Network on Chip", Proc. Of the International Workshop on SoC for Real-Time Applications, Banff, July 2005.

[8] M. Hansson, B. Mesgarzadeh and A. Alvandpour, "!.55GHz On-Chip Resonant Clocking with 2.3x Clock Power-Saving in 130nm CMOS", unpublished.

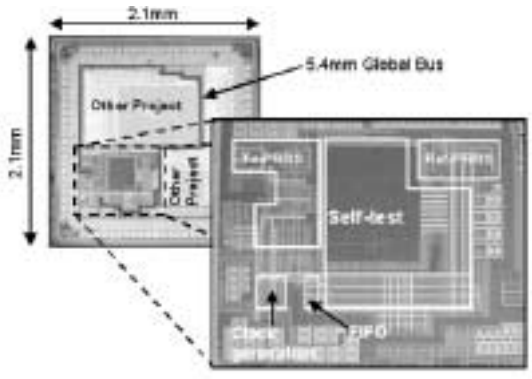

Fig. 1. Chip in 0.18 µm CMOS demonstraring latency insensitive transport of 3Gb/s per wire over 5.4mm

FAUST, an Asynchronous Network-on-Chip based Architecture for Telecom Applications

P. Vivet, F. Clermidy, D. Lattard
CEA-LETI, 38054, Grenoble, France
{ pascal.vivet ; fabien.clermidy ; didier.lattard }@cea.fr

With many tens of million transistors available on a single chip, the era of System-on-Chip has become a reality. Design and IP reuse is mandatory : chips integrate processor cores, DSPs, on-chip memories, IP-blocks, etc… A globally shared bus cannot meet the increasing demands of System-on-Chip interconnects, because the long-wire loads and resistances result in slow signal propagation. Alternative synchronous solutions, such as bus hierarchy, crossbars, still face the same issues : difficulties in timing validation and in connecting blocks running at different speeds, limited throughput versus power efficiency.

On the other hand, the advantages of "Network-on-Chip" (NoC) are numerous : high scalability and versatility, high throughput with good power efficiency. By separating communications from computations, NoC based architecture and packet switching provide dynamic communication possibilities. For very large scale integration, NoC are fully scalable because the number of nodes can be increased to provide higher global throughput, without degrading local network link performances, while the dissipated power is reduced to only active NoC links. The NoC distributed communication architecture is perfectly adapted to the Globally Asynchronous Locally Synchronous (GALS) paradigm where the NoC nodes and links are implemented using asynchronous logic while the NoC functional units are implemented with standard synchronous design methodologies.

For NoC protocols based on packet switching, low latency cannot easily be guaranteed. Therefore, the main issue in such NoC architectures is the notion of Quality-of-Service (QoS). We have proposed and developed ANOC, a complete Asynchronous NoC architecture adapted to GALS systems, using virtual-channels to provide low latency and QoS, and which is implemented in Quasi-Delay-Insensitive (QDI) asynchronous logic. Moreover, when addressing SoC of million gates and designing a new communication protocol such as ANOC, it is mandatory to offer to designers and architects a high level modelling strategy and all associated simulation/debug facilities. SystemC and Transaction-Level-Modeling (TLM) methodology are good candidates for such challenges, and have been used to develop a TLM model and environment of the proposed NoC protocol.

The proposed NoC architecture and design methodology has been successfully applied to the design of a 8 Million gate prototype chip in 130um STMicroelectronics CMOS technology. The FAUST chip (Flexible Architecture of Unified System for Telecom, see Figure 1) integrates a ARM946 core, embedded memories, smart DMA engines, numerous highly programmable HW blocks and re-configurable data-paths engines. A complete prototyping platform (Figure 2) has been developed , it integrates 2 FAUST chips and 2 FPGAs connected with the same unified NoC protocol to address Software defined Radio (SDR) applications. The current targeted application is a Telecom MC-CDMA MIMO application (htttp://ist-4more.org).

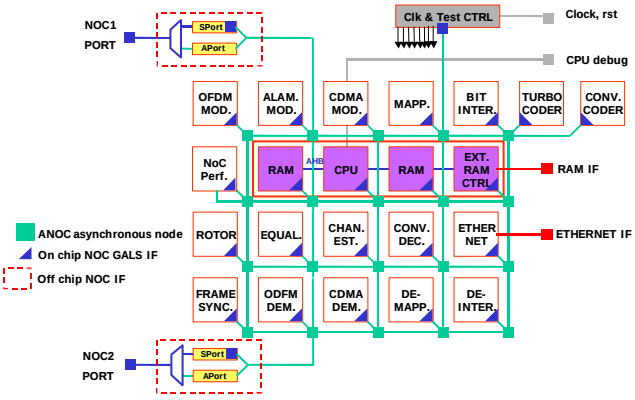

Figure 1 : FAUST chip architecture

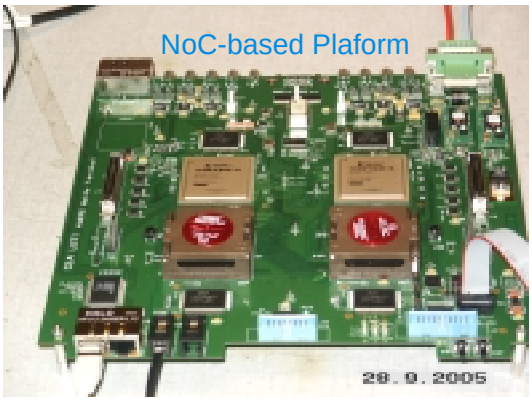

Figure 2 : FAUST prototyping platform

Exploration of CDMA-Based Network-on-Chip

Xin Wang and Jari Nurmi
Tampere University of Technology, P.O. Box 553, FIN-33101 Tampere, Finland
jari.nurmi@tut.fi

Abstract:
Code-Division Multiple Access (CDMA) technique applies a set of orthogonal codes to encode the data from different users before transmission in a shared communication media. Hence, it permits multiple users to share the communication media simultaneously by separating the data streams from different users in code domain. This concurrent transmission feature of CDMA technique is useful for handling the concurrent communications in an on-chip network. Therefore, in order to apply the CDMA technique into on-chip communication, a CDMA packet-switched on-chip network has been developed. The proposed CDMA Network-on-Chip (NoC) [1] supports Globally-Asynchronous Locally-Synchronous (GALS) communication scheme by using both synchronous and asynchronous circuits, and it has been implemented in Register-Transfer Level (RTL) using VHDL. In a packet-switched NoC which applies point-to-point connection scheme, data transfer latency varies largely if the packets are transferred to different destinations or to the same destination through different routes in the network. The proposed CDMA NoC can fix the data transfer latency into a constant value by sharing the communication media for data transfers in code domain instead of time domain. Thus, the data transfer latency can be guaranteed in an on-chip system where the proposed CDMA NoC is used to perform the communications among different components in the system.

1. Apply CDMA Technique in NoC

(1) Digital Encoding and Decoding Scheme

- Fully digital CDMA encoding and decoding schemes for CDMA NoC are proposed in Fig.1 and Fig.2.
- Transfer the binary equivalent of the summation value of the encoded data
- Accumulate the received summation values into two separate parts, positive part and negative part, according to the value of spreading code chip
- By comparing the accumulation value between positive and negative part, we can get the original data

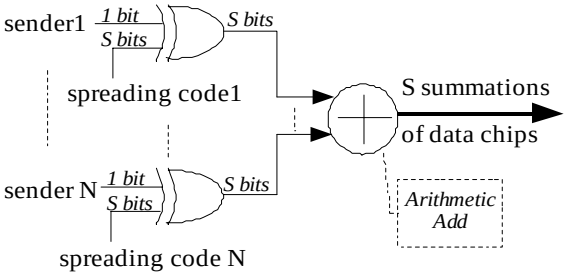

Figure 1. Digital CDMA Encoding Scheme

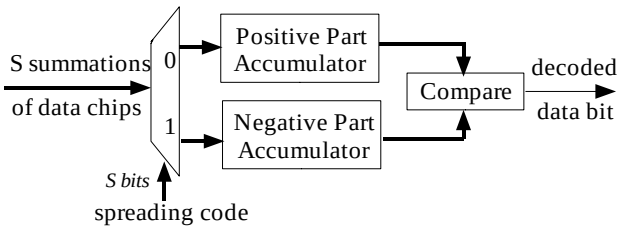

Figure 2. Digital CDMA Decoding Scheme

(2) Spreading Code Selection

Criteria:

- Orthogonal: normalized auto-correlation is 1, while the cross-correlation is 0.
- Balance: The number of bit '0' = The number of bit '1'

Walsh code meets the criteria.

(3) Spreading Code Protocol

Arbiter-Transmitter-Based (A-T) Protocol is proposed for the CDMA NoC:

- A unique spreading code is allocated to each user for data transfer
- Before starting data transfer, the sender will tell the arbiter who is the expected receiver for the data.
- Arbiter will inform the requested receiver to prepare the decoding code according to the sender

2. The Proposed CDMA NoC Structure

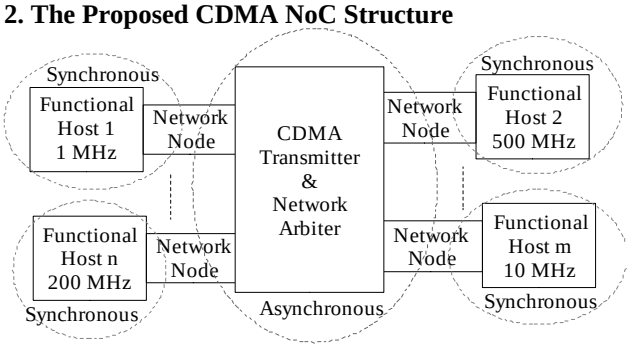

Figure 3. The Proposed CDMA NoC Structure

All the functional hosts which work in different clock domains in the system are connected to the CDMA NoC through individual 'Network Node'. The main part of CDMA NoC (CDMA Transmitter and Network Arbiter) applies asynchronous design, while 'Network Node' applies both synchronous and asynchronous designs to connect functional hosts and CDMA network together.

References

[1] X. Wang, and J. Nurmi; "An On-Chip CDMA Communication Network"; *Proceedings of 2005 International Symposium on System-on-Chip*; Finland, Nov. 2005.

A.2 NoC Development Support: Synthesis, Simulation, Verification methods, tools:

- **"Network-on-Chip Generation and Optimization Tool"**
  T. Ahonen, H. Bin and J. Nurmi
  Tampere University of Technology, Finland
- **"CHAINworksTM – an EDA tool suite and IP libraries for self-timed interconnect design synthesis"**
  J. Bainbridge, A. Bardsley and R. McGuffin,
  Silistix
- **"µSpider NoC Road Map"**
  S. Evain[1], J. Diguet[1] and D. Houzet[2]
  1 University of South Brittany, France
  2 INSA / IETR, Rennes Cedex, France
- **"Modeling and synthesis of Asynchronous Network on Chip using SystemC"**
  C. Koch-Hofer and M. Renaudin
  TIMA, France
- **"SystemC Simulation Model of a Flexible Network-on-Chip"**
  S. Määttä and J. Nurmi
  Tampere University of Technology, Finland
- **"Methodology, benefits and tools for NoC topology exploration"**
  P. Martin and J. Lecler
  Arteris SA
- **"NoC Performance Optimization via Long-range Link Insertion"**
  U. Ogras and R. Marculescu
  Carnegie Mellon University, USA
- **"Constraint-Driven Communication Synthesis of Networks On-Chip"**
  A. Pinto, L. Carloni, and A. Sangiovanni-Vincentelli
  U.C. Berkeley, USA
  Columbia University, USA

- **“HW/SW Partitioning and Interface Synthesis in NoCs”**
  F. Regazzoni and M. Lajolo
  ALaRI - USI, Lugano, Switzerland
  NEC Laboratories America, Princeton, NJ, USA
- **“Introduction to the Princeton Polaris Project”**
  V. Soteriou, N. Eisley, H. Wang, B. Li and L. Peh
  Princeton University, USA
- **“A Flexible System Level Design Methodology Applied to NoC”**
  A. Vander Biest, A. Leroy and F. Robert
  ULB, Belgique

Network-on-Chip Generation and Optimization Tool

Tapani Ahonen, Hong Bin and Jari Nurmi
Tampere University of Technology, P.O. Box 553, FIN-33101 Tampere, Finland
[tapani.ahonen, bin.hong, jari.nurmi]@tut.fi

Introduction

EDA tools hold the key to the success of NoC. We presented in [1] our first prototype tool, called OIDIPUS advocating the use of a cost evaluation criterion that contemplates both power consumption and performance estimates, linked to heuristics (simulated annealing) to search the design space. The tool is able to generate network instances that meet a given set of requirements. An algorithmic addition for the tool to measure the channel interaction, possibly resulting in congestion in some links, was presented in [2].

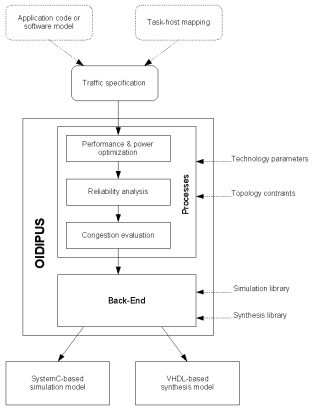

Fig. 1. Overview of an automated design flow for NoC-based systems.

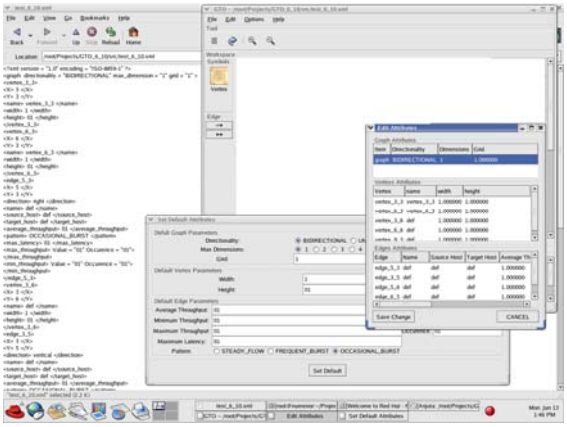

Fig. 2. Screenshot of the graphical user interface to OIDIPUS.

Discussion

We are currently working on refining the optimization processes to take into account more attributes of both the application and the underlying technology, and better analysis tools that are able to detect problematic conditions in the network in a wider range of configurations. Another open problem is the integration of OIDIPUS in more general design methodologies, especially obtaining an accurate traffic specification early in the design process. We have developed a convenient graphical user interface, illustrated in figure 2 above to speed up the process of system specification through close coupling to the analysis features of the toolset.

References

[1] T. Ahonen, D. A. Sigüenza-Tortosa, B. Hong, and J. Nurmi, “Topology optimization for application-specific networks-on-chip”, In Proc. *SLIP’04 2004 International Workshop on System Level Interconnect Prediction*, pages 53-60, Paris, France, 14-15 February 2004.
[2] D. Sigüenza-Tortosa and J. Nurmi, “Topology Design for Global Link Optimization in Application Specific Networks-on-Chip”, In Proc. *SOC’04 2004 International Symposium on System-on-Chip*, pages 135-138, Tampere, Finland, 16-18 November 2004.

CHAINworks™ – an EDA tool suite and IP libraries for self-timed interconnect design synthesis

John Bainbridge, Andrew Bardsley, Roy McGuffin, Silistix
wjb@silistix.com

The Silistix CHAINworks tool suite takes a description of the initiator and target ports of an SoC design and synthesizes a structural netlist for the required interconnect. The suite comprises three tools: CHAINdesigner™, CHAINcompiler™ and CHAINlibrary™.

CHAINdesigner captures a description of the connectivity requirements and ports of a design. It allow the user to guide the synthesis process to achieve the best trade-off of resource, performance, power consumption and timing closure cost. Analysis of the networks generated is facilitated through System-C models.

CHAINcompiler takes the output from CHAINdesigner and generates structural Verilog netlists for use with the back-end design flow. CHAINcompiler also creates test patterns for automatic test equipment (ATE), timing data for STA, and constraints and guidance for automated layout .

CHAINlibrary™ is an interconnect component library containing self-timed primitives and macro modules, together with templates for protocol converters for legacy synchronous bus and socket protocols.

CHAINworks will support the most-used bus protocols. The first product release supports ARM's AMBA protocols using a delay-insensitive, best-effort network fabric to facilitate globally asynchronous locally - synchronous systems. This approach allows an incremental approach to moving from bus-based architectures, whilst exploiting the principle benefits of self-timed technology. These benefits include timing closure, power management and congestion reduction. Each network client operates as its own independent timing domain greatly simplifying clock distribution and timing validation, and reducing switching activity when bursty traffic is encountered.

About Silistix
Silistix is a venture funded start-up, led by former members of the AMULET research group at the University of Manchester, UK. The team at Silistix have many years of self-timed design, synchronous design and SoC design experience which the company is applying to the creation of an EDA tool suite for the synthesis of networks-on-chip for use in large, multiple time domain SoCs.

µSpider NoC Road Map

Samuel Evain[1], Jean-Philippe Diguet[1], Dominique Houzet[2]
[1]University of South Brittany, LESTER, FRE 2734 CNRS, Rue de Saint-Maudé, 56325 Lorient Cedex – France
[2]INSA / IETR, UMR 6164 CNRS, 20, avenue des Buttes de Coësmes, 35043 Rennes Cedex – France
E-mail: evain@univ-ubs.fr

Abstract

µSpider is a tool suite providing designers with an ad-hoc NoC, with respect to today and future design constraints (figure 1). This is an open framework where evolving NoC concerns have been and will be implemented. The first idea implemented was the NoC code (VHDL) generation automation based on designer architectural choices [1]. The main parameters available are the NoC topology, the port buffer depth, the virtual channels features and for each of them: routing, arbitration and flow control policies. Then the network interface has been added including TDMA slot allocation, end-to-end flow control and different wrappers such as Amba and Avalon. The initial work was the implementation of the automation of design decision such as topology, IP mapping and path allocation. We are currently testing some new concepts. The first one is motivated by the question of clock synchronization over large SoCs. Our solution (synchronizer and time router) consists in solving this issue with an extension of TDMA techniques for real-time adapted to the context of Sub-NoCs communicating as GALS islands. The second one is the use of specific techniques for security management including sender and path authentications [2]. In addition, to get inter-operability with interconnect standards and to get full advantage of NoC parallel potential, specific wrappers and APIs (Application Programming Interface) are necessary and are currently under test.

We are now starting a new direction of research currently based on SystemC simulations, we address the hot topic of autonomous NoCs able to decide their own configuration based on environment and traffic evolutions.

To conclude, a full suite of tools is needed to achieve the ad-hoc NOC design to reach the future challenges of SoC design. µSpider is a user-friendly CAD tool for design space exploration and synthesis that implements currently required architectural options and remains open to upcoming NOC evolutions such as monitoring and self-configuration.

REFERENCES

[1] S. Evain, J. P. Diguet, D. Houzet, "A Generic CAD Tool for Efficient NoC Design", in IEEE ISPACS 2004, International Symposium on Intelligent Signal Processing and Communication Systems, Seoul, Korea, November 18-19, 2004.

[2] S.Evain, J-Ph.Diguet, "From NoC Security Analysis To Design Solutions", IEEE Work. On Signal Processing Systems, SIPS'05, Nov. 2005, Athens

Figure 1: µSpider Biotope

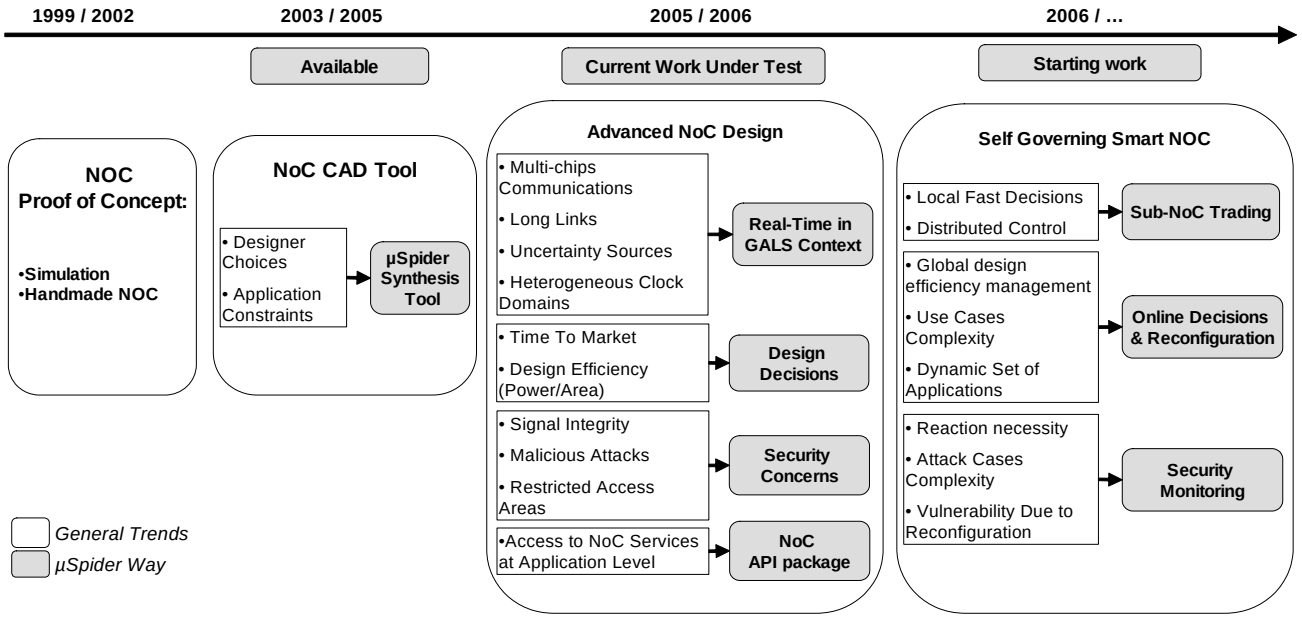

Modeling and synthesis of Asynchronous Network on Chip using SystemC

Cedric Koch-Hofer and Marc Renaudin

TIMA, France

cedric.koch-hofer@imag.fr

Abstract

This poster presents the ASC SystemC library which is intended to be used for modeling asynchronous circuits and more specifically asynchronous Network on Chip (NoC) and Globally Asynchronous Locally Synchronous (GALS) systems. This library is also intended to be used for synthesis as an input of the TAST framework.

Like common CSP based languages for asynchronous modeling (CHP, Tangram, Balsa) the communication between modules is supported by channels and ports offering the following primitives of communication: send, receive and probe. This library also defines some methods and operators for synchronizing processes.

For modeling the non deterministic behavior of asynchronous arbiters, this library defines two new statements as_choice and as_guard. These two statements allow us to model and simulate non-deterministic choice over a set of guards.

One of the fundamental properties of the asynchronous circuits is that they are not sensitive to delays. For checking this fundamental property, the selection of a process to execute among the set of runnable processes must be non-deterministic. One of our on-going researches on ASC is to enable a non-deterministic scheduling of the processes of the SystemC OSCI simulator. Our last current investigation on ASC is to formally define the semantic of SystemC and ASC for being able to properly synthesize asynchronous circuits.

* Denotes Contact Author

SystemC Simulation Model of a Flexible Network-on-Chip

Sanna Määttä and Jari Nurmi
jari.nurmi@tut.fi
Tampere University of Technology, P.O. Box 553, FIN-33101 Tampere, Finland

Abstract

A high level SystemC based simulation model of a flexible Network-on-Chip (NoC), Proteo [1], is presented. Proteo architecture is a packet switched communication network concept. Proteo aims at creating structural and flexible communication systems based on a small library of pre-designed, parameterized, and reusable communication blocks. The SystemC simulation model is implemented for Proteo network architecture and communication protocol exploration. Then the simulation times of a simple network topology, ring, are measured for different amount of network nodes. The simulation results are compared with the similar VHDL simulation model [2]. Figure1 illustrates the simulation times when having either 2, 4, 8 or 16 network nodes connected to a ring topology. The X-axis illustrates the number of network nodes and the Y-axis tells the simulation time (CPU-time) in seconds. The blue curve with square shape symbols presents the VHDL model simulation times and the red curve with diamond shape symbols the SystemC model simulation times. The SystemC model is simulated in its own light weight simulation kernel, whereas the VHDL model needs an application, ModelSim v.5.7. ModelSim is used through command prompt to make the simulation results more comparable. With 2, 4, 8, and 16 nodes, the SystemC model is 54, 53, 57, and 65 percent faster than the VHDL model. Even though ModelSim probably cause extra overhead in the VHDL model's simulation times, the results shows, that the SystemC model is faster than VHDL model, which allows faster architectural and communication protocol exploration and therefore faster design flow.

References:

[1] Sigüenza-Tortosa D., Nurmi J.: Proteo: A New Approach to Network-on-Chip, Proceedings of *IASTED Communication Systems and Networks (CSN)*, Spain, September 2002.

[2] Sigüenza-Tortosa D., Nurmi J.: VHDL-Based Simulation Environment for Proteo NoC, Proceedings of *IEEE High Level Design, Validation and Testing (HLDVT)*, France, October 2002.

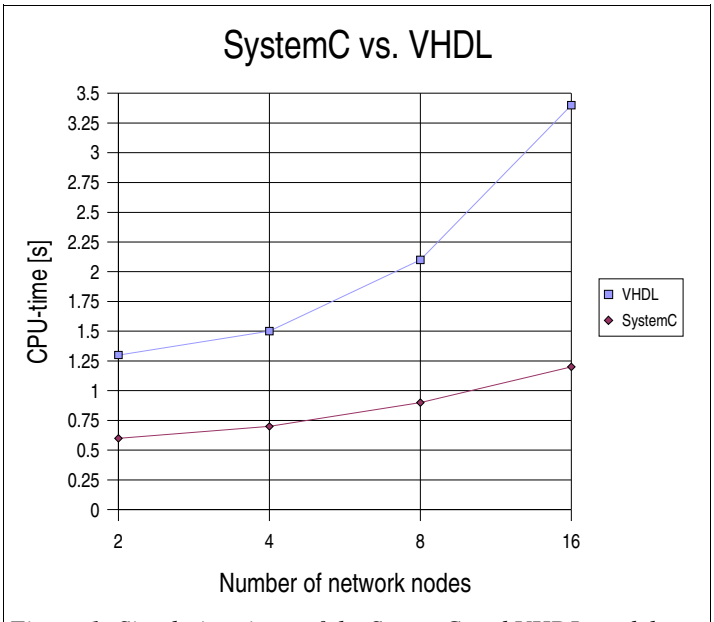

Figure 1: Simulation times of the SystemC and VHDL models

Methodology, benefits and tools for NoC topology exploration

Philippe Martin – Product Marketing Director, Arteris SA
Jean-Jacques Lecler – Senior Architect, Arteris SA

Network-on-Chips are a necessary evolution for on-chip communications, and their advent is driven by major evolutions in SoC designs:

a) SoC architectures are increasingly complex, integrate more and more IP blocks of higher performance. A SoC manages many data flows, of varied bandwidth, latency, and more generally Quality of Service requirements, which may vary over the software application spectrum of the SoC

b) Managing global wires and associated topics such as large fanout and clock skew become dominant issues in SoC designs. NoCs clearly solve these issues by providing single-fanout, unidirectional communication links, and supporting Globally Asynchronous, Locally Synchronous design styles.

Many NoC research studies have focused on specific aspects of NoC (performance predictability, self-timed logic, fault tolerance …) and used regular NoC topologies (such as meshes, hypercubes, etc..) without any application-specific input, while applications are usually very heterogeneous and asymmetrical. Regular topologies usually lead to overdesigned communication systems, for example 2D mesh links may be oversized for many flows, undersized for a few, and matching such a topology to floorplan constraints may lead to performance issues.

We believe that none of the standard topologies address the extremely variable floorplanning, silicon cost, complexity and performance issues of actual production SOCs, and that SoC architects must be empowered to make application-specific NoC topology tradeoffs early in the design cycle, leading to decisions that are critical for the actual SoC performance, its shelf lifetime, its design cost and its unit cost.

We provide an architecture exploration methodology, based on a NoC modeling tool called NoCexplorer™. IP characteristics (IP socket type, frequency), application communication requirements (burst and temporal characteristics of data flows, QoS requirements), and topology of the NoC under elaboration, (including individual link characteristics, clock domains, packetization overhead, arbitration schemes, etc…) are captured in a simple fashion, throughput-accurate simulations over many microseconds of real-time executed in a few seconds of runtime. Simulation results provide the SoC architect with feedback on many crucial questions such as:

- does my memory hierarchy fulfill performance expectations ?
- does my arbitration scheme fulfill SoC performance constraints ?
- do traffic burst lengths provide a reasonable trade-off between NoC, DRAM efficiency and latency on one side, and latency, buffering area on the other side ?
- have I put enough buffering in the NoC to adapt throughput rates of data producers and consumers ?
- can I further reduce the count of long wires between my top-level IP clusters ?

The posters will describe the methodology and tool, some examples and results of actual SoC architectures using NoCexplorer.

NoC Performance Optimization via Long-range Link Insertion

Umit Y. Ogras, Radu Marculescu
Electrical and Computer Engineering Department,
Carnegie Mellon University
{uogras, radum}@ece.cmu.edu

Abstract

Networks-on-Chip (NoCs) represent a promising solution to complex on-chip communication problems. The NoC communication architectures considered so far are based on either completely regular or fully customized topologies.
We propose a methodology to automatically synthesize an architecture that consists of a superposition of a few long-range links inserted on top of a standard mesh network. The application-specific long-range links are inserted to optimize the performance of the network by increasing the critical traffic workload at which the network transitions from a free to a congested state. This way, we exploit the benefits offered by both complete regularity and partial topology customization. Indeed, extensive simulations and an FPGA prototype demonstrate that significant reduction in the average packet latency and a major improvement in the achievable network throughput are obtained.

Constraint-Driven Communication Synthesis of Networks-On-Chip

Alessandro Pinto, Luca P. Carloni, and Alberto L. Sangiovanni-Vincentelli
U.C. Berkeley and Columbia University
{apinto,alberto}@eecs.berkeley.edu luca@cs.columbia.edu

The focus of the electronic design automation (EDA) research community is moving from computation-centric design methodologies to communication-centric ones. This ongoing shift can be observed in all areas where EDA tools and design flows can provide major benefits: multi-core systems-on-chip (SOC), wireless sensor networks, distributed embedded software synthesis. In particular, the design of SOCs increasingly entails the capability to reuse and assembly hundreds of intellectual properties (IP) cores that have been either built for previous design generations within the same company or acquired off-the-shelf from specialized vendors. The consequent challenge is the ability of addressing the communication and synchronization issues that naturally arise while assembling pre-designed components [3]. Networks-on-chip (NOC), have been proposed as a robust and efficient solution to some of these issues. The idea, borrowed from macro-level networks like the Internet, is to build robust channels, switches, and network interfaces in order to route reliably data packets among the chip components in a standardized manner. Besides the similarities, however, the reality of the semiconductor industry and the peculiar aspects of on-chip communication call for the development of *ad hoc* design methodologies for NOCs.

In the framework of the platform-based design paradigm [4, 7], we have proposed the idea of constraint-driven communication synthesis [5] and we have applied it to the particular case of NOC synthesis [6]. In our approach a NOC is derived in a *meeting-in-the-middle* fashion: we developed a common semantic domain to match a set of constraints that capture the required system-level communication performance and a synthetic library that abstracts the physical-level characteristics of the NOC components. In particular, the NOC design specification is given as a set of end-to-end communication constraints, each represented by a pair (b, d) denoting the bandwidth and distance of the corresponding channel. The emphasis on a *bandwidth-oriented approach* is motivated by the nature of SOC design, where the average-case performance is the main goal, and the advantages of managing latency-throughput trade-offs in the context of latency-insensitive design [1, 2]. In particular, latency-insensitive protocols provide a robust solution to implement flexible and scalable flow control mechanisms. The library of elementary NOC components contains the abstraction of the silicon characteristics of wires and switching circuits, e.g. the bandwidth that a wire built in a given metal layer can sustain over a certain distance.

In [5, 6] we discussed the foundations of this methodology, but we presented algorithms that limit the design exploration only to bus-based and point-to-point network topologies. More recently, we have extended the proposed methodology to all stages of the design process with a particular emphasis on minimizing energy consumption. At the specification level we have developed an algebraic framework that is compositional and that can take into account different classes of design constraints. At the implementation level, we have refined our characterizations of the physical components in order to model dynamic and static power consumption of wires as well as the energy trade-offs between channel communication and switching/storage capabilities. Finally, as the common semantics domain still represents all possible network topologies that can be built with the given library of elementary components, we have developed a new set of algorithms that are able to explore the entire solution space and find the optimal topology, which minimizes power consumption while satisfying all communication constraints.

References

[1] L. P. Carloni, K. L. McMillan, A. Saldanha, and A. L. Sangiovanni-Vincentelli. A methodology for "correct-by-construction" latency insensitive design. In Andreas Kuehlmann, editor, *The Best of ICCAD - 20 Years of Excellence in Computer-Aided Design*, chapter 12, pages 143–158. Kluwer Academic Publishers, 2003.

[2] L. P. Carloni, K. L. McMillan, and A. L. Sangiovanni-Vincentelli. Theory of latency-insensitive design. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 20(9):1059–1076, September 2001.

[3] L. P. Carloni and A. L. Sangiovanni-Vincentelli. Coping with latency in SOC design. *IEEE Micro*, 22(5):24–35, Sep-Oct 2002.

[4] K. Keutzer, S. Malik, A. R. Newton, J. Rabaey, and A. Sangiovanni-Vincentelli. System level design: Orthogonalization of concerns and platform-based design. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 19(12):1523–1543, December 2000.

[5] A. Pinto, L. P. Carloni, and A. L. Sangiovanni-Vincentelli. Constraint-driven communication synthesis. In *Proc. of the Design Automation Conf.*, pages 783–788, New Orleans, LO, June 2002. IEEE.

[6] A. Pinto, L. P. Carloni, and A. L. Sangiovanni-Vincentelli. Efficient synthesis of networks on chip. In *Proc. Intl. Conf. on Computer Design*, pages 146–151, San Jose, CA, October 2003. IEEE.

[7] A. L. Sangiovanni-Vincentelli. Defining platform-based design. In *EEDesign*, feb 2002.

HW/SW Partitioning and Interface Synthesis in NoCs

Francesco Regazzoni
ALaRI - USI, Lugano, Switzerland
regazzoni@alari.ch

Marcello Lajolo
NEC Laboratories America, Princeton, NJ, USA
lajolo@nec-labs.com

Abstract

With deep sub-micron technology, chip designers are expected to create System on Chip (SoC) solutions by connecting different Intellectual Property (IP) blocks using efficient and reliable interconnection schemes.

Networks on Chip (NoCs) are quite compelling because, by applying networking techniques to on-chip communication, they allow to implement a fully distributed communication pattern with little or no global coordination. This avoids the problems associated with distribution of global clock over the entire chip, clock power consumption, clock skew and electromagnetic interference.

On the other hand, in order to benefit from the NoC communication paradigm, designers should perform a careful functional mapping for taking advantage of spatial locality, by placing the blocks that communicate more frequently to be closer together. This reduces the use of long global paths and the corresponding energy dissipation.

In this work we propose the wrapper organization for tile based NoC architectures shown in Figure 1. The architecture consists of multiple tiles organized in a Manhattan-like structure resulting in a 2-D mesh interconnection topology. Tiles are wrapped in a tile wrapper that provides access and isolation to each tile.

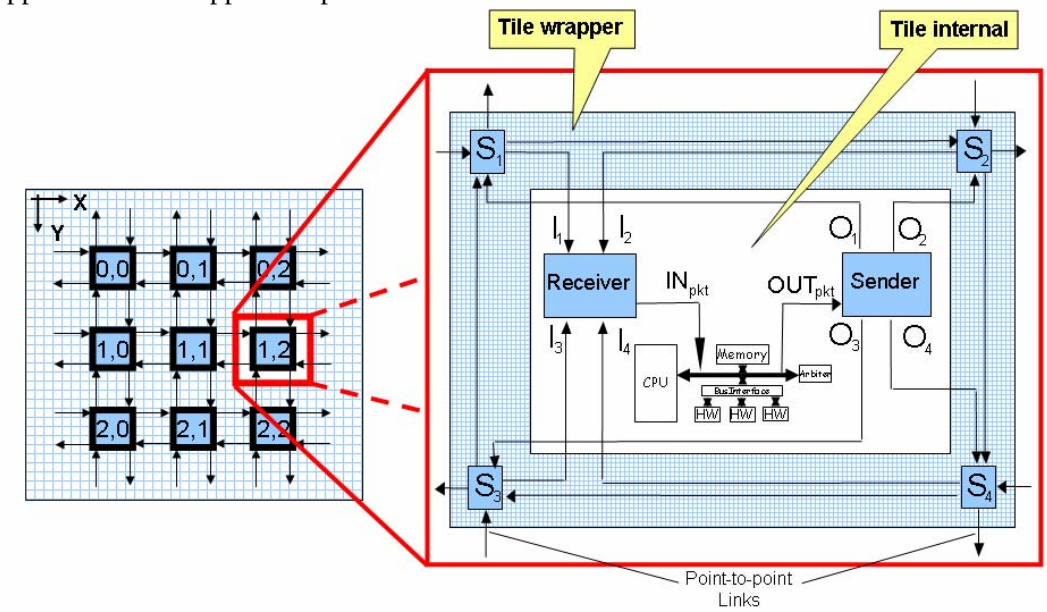

Figure 1: The proposed 2D Mesh NoC Architecture

The internal organization of our tile is shown on the right side of Figure 1. We have four switches (*S1* to *S4*), one at every corner of the tile and dedicated *Receiver* and *Sender* units that act as adaptation layers between the tile wrapper (that is the same in every tile) and the internal tile (that instead can be different in every tile). The internal tile might contain multiple buses and processors and the only assumption is that one single bus within the internal tile will be allowed to be connected to the tile wrapper.

In a tile based NoC, inter-tile communication is provided by the on-chip network once packets are generated, but interface synthesis is still needed in order to provide intra-tile communication for tasks mapped onto the same tile. Interface synthesis is also needed for implementing inter-tile communications, where it is necessary to configure the sender units in order to route the communication.

In this work we show, by means of an example, how to approach the hardware/software partitioning and interface synthesis phases in order to take advantage of the new NoC paradigm.

We discuss also suitable routing algorithms and flow control techniques necessary to provide predictable on-chip communications latency and we propose a programming model for NoCs.

Introduction to the Princeton Polaris Project

Vassos Soteriou, Noel Eisley, Hangsheng Wang, Bin Li and Li-Shiuan Peh
Department of Electrical Engineering
Princeton University, Princeton, NJ 08544
{soteriou,eisley,hangshen,binl,peh}@princeton.edu

Presenting Author: Noel Eisley

Abstract

Technology trends are driving parallel on-chip architectures both in the form of Multi-Processor-Systems-on-a-Chip (MPSoCs) and Chip Multi-Processors (CMPs). The increasing on-chip communication demand among the processors necessitates the use of scalable, high-bandwidth network on-chip (NoC) fabrics. However, myriad applications coupled with vast amounts of on-chip resources give rise to a wide range of feasible network architectures; for cost-effectiveness, in the early stages of design-space exploration a system designer critically needs to identify the architecture(s) that best balance(s) cost/performance for a set of applications, before investing time and resources in an actual detailed NoC design process. This prompted us to develop Polaris, a system-level roadmap (SLR) for on-chip networks which scans over a wide range of NoC configurations at various current state-of-the art and future process technologies guiding designers towards the most suitable on-chip network design(s) tailored to their performance needs and power/silicon area constraints, given a range of application classes that are characterized by their spatio-temporal parameters, applications that the NoC is to run. Unlike synthesis toolchains that generate designs tailored specifically to a known traffic workload, architecture and process, a roadmapping toolchain such as Polaris needs to explore designs based on ***projected*** *traffic, architecture and process characteristics. While the Polaris roadmapping toolchain is extensible so new traffic, network designs and processes can be added, the current first version of the roadmap already incorporates 7,872 points in the design space of NoCs, all iterated within a tractable run time of 125 hours on a typical desktop machine, while maintaining high relative and absolute accuracies when validated against detailed synthesis results. This first set of results, which will be made available on the web, provides insights on current* $90nm$ *and future* $50nm$ *technologies; guided by Polaris' flexibility and expandability, possible extensions to the project include, but are not limited to, the accommodation of additional NoC micro-architectures, topologies, routing and flow-control protocols, support for heterogeneous systems, incorporation of circuit reliability models and output metrics for further NoC exploration and roadmapping. Interesting avenues include questions such as "what will the communication needs of future applications look like" and "what will the architectures that will run them be like"; with Polaris providing the first steps towards identifying the most suitable on-chip communication fabric, these questions may soon be answered with confidence. In synopsis, we see Polaris effectively filling the absence of a NoC SLR, serving as a guide for NoC design projections.*

A Flexible System Level Design Methodology Applied to NoC

Alexis Vander Biest, Anthony Leroy, Frédéric Robert
Free University of Brussels, Belgium
anthony.leroy@imec.be

Nowadays, designing custom Network-on-chip has become very difficult. On one hand, many architectural choices have to be made: topology, routing algorithm, mapping of the application on the nodes,… . On the other hand, the technology shift from one generation to another is becoming more and more painful due to the increasing influence of Deep Sun Micron effects. In a highly computation intensive VLSI design, the overall performance highly relies on the network. It has to be deterministic, to offer high bandwidth and low latency communication while preferably introducing a low overhead in terms of energy consumption, design cost and silicon area [1].
After fixing these design requirements, classical design flows proceed from top to down by successively refining the NoC design along the different steps in a hierarchical way [2] using iterative optimisation methods : this kind of process is very complex and extremely time consuming [3].
In fact nothing guarantees the flow convergence : without any deep design space exploration, targeted performance values may simply be not be reachable within the chosen constraints. System level design tries to cope with that issue by providing the designer with performance prediction tools faster than the design tools themselves so that dead-end solutions could be eliminated very early. A huge area of NoC research is devoted to the study of predictive models for different parameters (mainly the silicon area, the latency and the energy consumption) but few efforts have been made to integrate all these models into a generalized framework [4].
In this poster, we present a general formalism that we have developed to provide the support for the specification of flexible system-level design tool and try to apply it to the particular case of NoC. The formalism features the following :

- Support for hierarchical design flow modelling to allow the user to explicitly choose between modelling accuracy and estimation speed.
- Possible representation of many parameters linked to design quality (like area, energy consumption) and design activity (like design time and cost)
- Flexibility in the use of models to ease their comparison
- Support for model error representation and its impact on the results to include the uncertainties in the design choices

[1] 'Key Research Problems in NoC Design: A Holistic Perspective', U. Ogras, J. Hu, R. Marculescu, Proc. CODES+ISSS, Jersey City, NJ, Sept.2005, pp. 69-74

[2] "SAGA : Synthesis Technique for Guaranteed Throughput NoC Architectures", Krishnan Srinivasan and Karam S. Chatha, Proceedings of Asia South-Pacific Design Automation Conference (ASPDAC), Shanghai, China, January 2005

[3] L. Benini and G. De Micheli, “Networks on Chip: A New SoC paradigm,” IEEE Computer, January 2002, pp. 70-78

[4] Krishnan Srinivasan, Karam S. Chatha: ISIS: A Genetic Algorithm Based Technique for Custom On-Chip Interconnection Network Synthesis. VLSI Design 2005: 623-628

A.3 QoS Issues: Results on Performance Analysis, Optimisation, Simulation:

- **"QoS Oriented Configurable Network on Chip"**
  M. Al Faruque, X. Ye, G. Weiss and J. Henkel
  University of Karlsruhe, Germany
- **"QNoC: QoS Network on Chip"**

  E. Bolotin, I.Cidon, R. Ginosar and A.Kolodny
  Electrical Engineering Department, Technion, Haifa, Israel
- **"Efficient Link Capacity and QoS Design for Network-on-Chip"**
  Z. Guz, I. Walter, E. Bolotin, I. Cidon, R. Ginosar and A. Kolodny
  Electrical Engineering Department, Technion, Haifa, Israel
- **"Empirical Optimization of Multiprocessor Systems and their Heterogeneous Communication Fabrics"**
  G. Hellestrand, J. Torossian and C. Alford
  VaST Systems Technology Corporation, USA
- **"On Options for Accessing Regions in NoC"**
  R. Holsmark and S. Kumar
  Jönköping University, Sweden
- **"Spatial Division Multiplexing: A Novel Approach for Guaranteed Throughput on NoCs"**
  A. Leroy[1/2], P.Marchal[1], F. Robert[2] and F. Catthoor[1/3]
  1 IMEC Belgium
  2 VUB Belgium
  3 KUL Belgium
- **"Evaluation of On-chip Networks Using *Deflection* Routing"**
  Z. Lu, M. Zhong and A. Jantsch
  Royal Institute of Technology, Sweden
- **"Dynamic Time-Slot Allocation for Networks on Chip with TDMA QoS"**
  T. Marescaux[1], B. Brické[1/2], P. Debacker[1/2], V. Nollet[1] 1, H. Corporaal[3]
  1 IMEC V.Z.W., Belgium
  2 Katholieke Universiteit Leuven, Belgium
  3 Technical University Eindhoven (TU/e), The Netherlands
- **"Application Parallelism Exploitation using NoC with Multithreaded Processors"**
  R. Pop and S. Kumar
  Jönköping University, Sweden
- **"Curing Hotspots in Wormhole NoCs"**
  I. Walter, I. Cidon, R. Ginosar and A. Kolodny
  Electrical Engineering Department, Technion, Haifa, Israel

QoS-Oriented Configurable Network on Chip

M. A. Al Faruque, X. Ye, G. Weiss, J. Henkel
CES - Chair for Embedded Systems,
University of Karlsruhe
Contact E-mail: alfaruque@ira.uka.de

Networks on Chip gain increasing interest since Moore's Law allows for integrating virtually hundreds of processors on a single die. Efficient on-chip communication architectures will drive the future of SoC (System-on-Chip) architectures. One of the major challenges in NoC design is the large design space spanned by an extensive set of (partly) interdependent parameters. Only a set of carefully adapted parameters will allow to unveil the potential benefits of a NoC.

The goals of our NoC project are twofold: (1) a rapid prototyping environment for NoCs that allows specification and evaluation of a customized NoC within hours through a proprietary NoC IP library. As a result, cycle-accurate transaction-level data is obtained through execution on an FPGA platform, allowing for fast design space exploration; (2) A concept for guaranteeing QoS (Quality-of-Service) of a NoC: lossless communication, in-order transmission and priority-based connection. The advantage of virtual channels, priorities, dynamic channel preempting/establishment is used to provide guaranteed connection for some highly prioritized connections that are part of, for example, a security task etc.

Previous research has not completely addressed these concerns: In [1, 3], Xpipes is discussed. It allows for customization at data flow and Network Interface (NI) level. Our approach goes further as it also considers protocol layer customization since we consider this mandatory for reliable QoS guarantee. Another significant NoC framework is Aethereal [2], which also addresses the importance of QoS. The connection establishment is rather costly.

Our approach can generate a synthesizable NoC specification through a VHDL and SystemC based proprietary NoC IP library. On the upper abstraction level it allows for fast HW/SW co-simulation, on the lower abstraction level it provides rapid prototyping on a scalable FPGA platform. Fig. 1 shows the FPGA-based prototyping environment using a ProDesign plateform. The synthesizable NoC is build from our NoC component IP library. All of these components are freely parameterizable.

To keep the NoC predictable in terms of cycle counts, the router-to-router connections and also router-to-IP connections are typically cycle-accurate. To provide QoS: lossless communication, CRC check/parity bits can be used in the data link layer. In-order transmission is guaranteed in flow control and in NI. The analytical model shows, after task mapping to the NoC IPs the number of guaranteed connections can often be kept small in typical application scenarios and thus keep QoS related overhead low. As for guaranteed throughput: it deploys virtual channels for highly prioritized connections. Packet-based communication with optimal pipelined and cycle-accurate flow control has been considered in our design concept also.

In Table 1 the *man*hour* time to build a configurable NoC from a library of components is shown. It clearly shows the suitability of our rapid prototyping NoC environment.

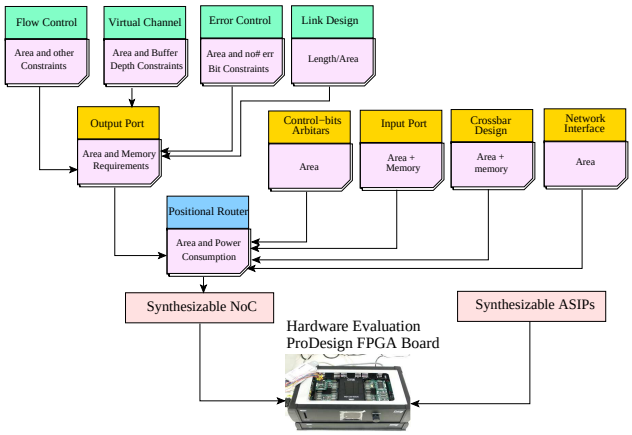

Figure 1: FPGA based NoC simulation

Configuration	Implementation Time	Test Time
1 Router (5 In / 9 Out)	15 h from Configurable parts	8 h
2*2 NoC	3 h for each positional router	2 h
	4 h connecting all parts	2 h
3*3 NoC	4 h for each positional router	2 h
	8 h connecting all parts	4 h
4*4 NoC (9 types of positional router)	10 h connection time from previously designed router	8 h

Table 1: Design and Testing time for different components and the whole NoC

Typically, resource hungry connection-oriented NoC designs perform bad for variable-length transaction and waste available bandwidth during the absense of data in a specified slot. Our prioritised connection, configurable, router-to-router cycle accurate and above all QoS supported NoC can maximize the resource utilization and bandwidth utility maximization and can improvise its usability compared to state-of-the-art methods in the Multi-Processor System on Chip(MPSoC) designs.

REFERENCES

[1] M. Dall'Osso and et. al.:*Xpipes: a latency insensitive parameterized network-on-chip architecture for multi-processor SoCs.*, & ICCD, pages 536-539, 2003.

[2] E.Rijpkema and et al.:*Trade-offs in the design of a router with both guaranteed and best-effort services for networks on chip.*, & DATE 2003, pages 350-355, March 2003.

[3] S. Stergiou and et al.:*Xpipes lite: A synthesis oriented design library for networks on chips.*, & DATE 2005, 2005.

QNoC: QoS Network on Chip

E. Bolotin, I.Cidon, R. Ginosar and A. Kolodny
EE Department, Technion - Israel Institute of Technology
bolotin@tx.technion.ac.il, {cidon;ran;kolodny}@ee.technion.ac.il

Abstract:

We define Quality of Service (QoS) and cost model for communications in Systems on Chip (SoC), and derive related Network on Chip (NoC) architecture and design process. SoC inter-module communication traffic is classified into four classes of service: signaling (for inter-module control signals); real-time (representing delay-constrained bit streams); RD/WR (modeling short data access) and block-transfer (handling large data bursts). Communication traffic of the target SoC is analyzed (by means of analytic calculations and simulations), and QoS requirements (delay and throughput) for each service class are derived. A customized Quality-of-Service NoC (QNoC) architecture is derived by modifying a generic network architecture. The customization process minimizes the network cost (in area and power) while maintaining the required QoS.

The generic network is based on a two-dimensional irregular mesh and fixed shortest path multi-class wormhole routing. Once communication requirements of the target SoC are identified, the network is customized as follows: The SoC modules are placed so as to minimize spatial traffic density, unnecessary mesh links and switching nodes are removed, and bandwidth is allocated to the remaining links and switches. The result is a low cost customized QNoC for the target SoC which guarantees that QoS requirements are met. We show several design examples of our NoC for typical SoC and compare to alternative solutions.

We analyze the generic cost in area and power of Networks on Chip and alternative interconnect architectures. Analytical calculation quantifies and clearly shows the intuitive NoC scalability advantages. We also explore possible cost tradeoffs between the number of buffers in network routers and the inter-router links capacity.

Efficient Link Capacity and QoS Design for Network-on-Chip

Zvika Guz[1], Isask'har Walter[1], Evgeny Bolotin[1], Israel Cidon[2], Ran Ginosar[2], Avinoam Kolodny[2]
Electrical Engineering Department, Technion, Haifa, Israel
[1]{zguz, zigi, bolotin}@tx.technion.ac.il [2]{cidon, ran, kolodny}@ee.technion.ac.il

Presenting author: Zvika Guz

Abstract

The allocation of link capacities is an important phase in the automated design process of a network on chip based system. The goal of this optimization phase is to minimize the network cost (in term of area and power) while maintaining the required QoS for the specific system communication demands; insufficient allocation will not meet performance requirements while lavish allocation will result in excessive power and area consumptions.

Since most NoC architecture use wormhole switching, this customization process heavily rely on simulations as no existing analysis accounts for the combination of heterogeneous traffic patterns and virtual channels, both fundamental characteristics of wormhole based NoC interconnect. The use of simulations makes the task of searching for efficient capacity allocation computationally extensive and does not scale well with the size of the problem. On the other hand, a detailed exact analytical solution of a complex NoC is intractable and simplistic approximations may lead to inaccurate results

In this work, we propose a hybrid methodology to the network design problem that combines the best of both worlds. First, we propose a novel and simple analytical delay model for a wormhole based NoC, which approximates the network behavior in a wide range of loads. Given any system (in terms of topology, routing and link capacities) and its communication demands (in terms of packets length and generation rate), the model estimates the delay experienced by every source destination pair. To the best of our knowledge, this is the first modeling and analysis of a wormhole network with non-uniform link capacities. Next, we propose an algorithm that applies the delay modeling and analysis to efficiently allocate capacities to network links without need for repetitive simulations in the inner-loop of optimization. Simulation runs are only used for final verification and fine tuning of the system. Using a design example we demonstrate that our algorithm considerably decreases the total NoC cost and significantly improves the quality and speed of the customization process.

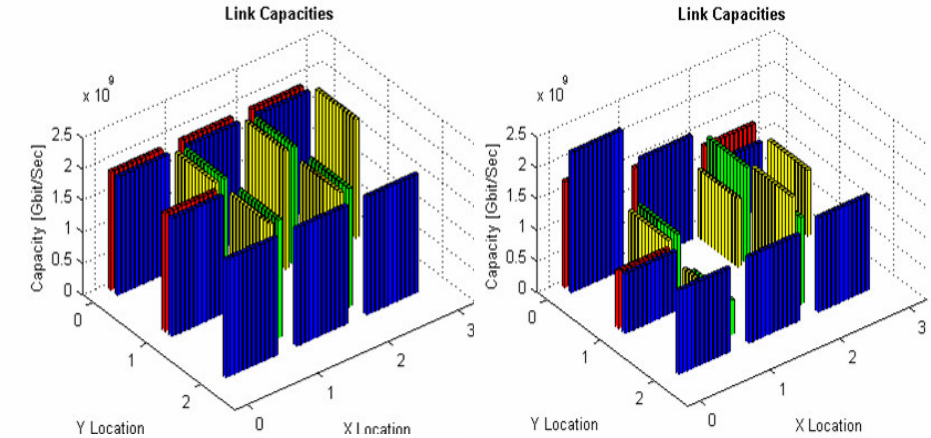

(a) **(b)**

Figure 1: NoC cost saving demonstrated on a heterogeneous SoC: (a) Wormhole classic uniform link capacity assignment; (b) Link capacities as assigned by the proposed algorithm. While both systems meet the same QoS requirements, the proposed algorithm achieves a reduction of 30% in resources.

Empirical Optimization of Multiprocessor Systems and their Heterogeneous Communication Fabrics

Graham R. Hellestrand, James R. Torossian, Casey Alford

VaST Systems Technology Corporation
1250 Oakmead Pkwy, Sunnyvale, CA USA

Abstract

The overall optimization of throughput in multiple processor-centric systems on a chip is a study of matching communication structures with peak computational demands for information. A more complex optimization, say of throughput and power consumption, likely yields very different computational and communications structures both within a computational unit (called a plex) and between plexes.

This study considers the specification of systems, their mapping to computational plexes, the communication required within a plex and between plexes, the formulation of the optimizing function, and the measures required to quantitatively drive the optimization of such systems. Identifying high bandwidth information transfers between functions (or tasks / processes) in a specification and clustering computational plexes according to required throughput rates identifies hierarchies of communications constraints that are useful in directing the mapping of the specifications onto realizable physical subsystems. Figure 1 shows a system that might have resulted from such a process. This is an analagous problem to that of optimizing placement and routing of hierarchical subsystems using floorplans derived from an HDL specification.

The mapping effort is typically undertaken in 2 or 3 steps prior to realization and the complexity of the process is beyond our powers to intuitively deduce an optimal system, or indeed often, even a system that will satisfy the constraints in the marketing requirement. The use of high performance, timing accurate parameterized models of computational engines (eg. processors, hardware devices) and communication fabrics (eg. point-to-point interconnects, buses, combinational interconnects), together with ability to measure characteristics that can be reduced to factors pertinent to the optimization function, form the basis of an empirical approach to solving the problem using simulation. The potential solution space is very large. The use of *design of experiments* strategies to reduce the potentially millions of experiments to a few hundred, and multi-variate statistics to reduce the number of factors used to drive the optimization process are necessary to help make this problem tractable.

The presentation will address these issues and show experiments and techniques being advocated to produce *optimal* systems.

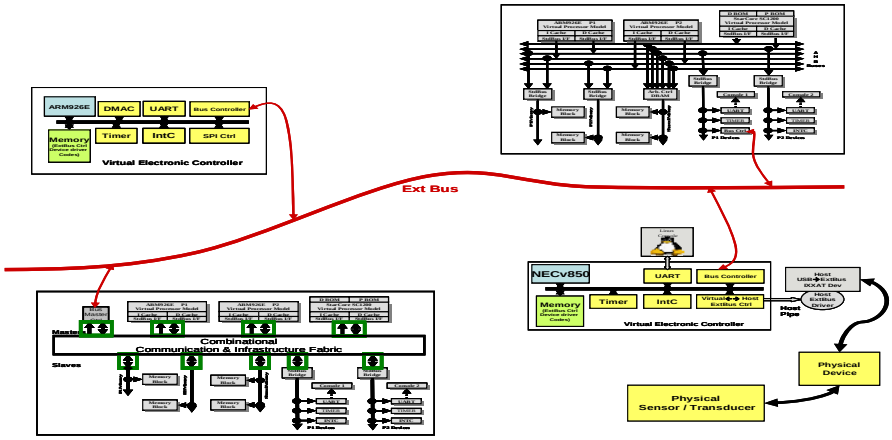

Figure 1: Clusters of High Performance, Closely-Coupled Processing Systems with Complex Infrastructure Fabrics Communicating via a Low Bandwidth Bus Fabric on a Single Chip

On Options for Accessing Regions in NoC

Rickard Holsmark and Shashi Kumar
School of Engineering, Jönköping University, Sweden
{rickard.holsmark, shashi.kumar}@ing.hj.se

A popular NoC topology is the mesh, where nodes are connected in a two dimensional array. Each node constitutes of a resource connected to a router, which is responsible for message routing. To use full advantages of this architecture, the communication facilities should be homogeneous and resources should fit in the tiles/slots. This implies that the tile size should be equal to the size of the largest core in the system. This will lead to enormous inefficiency, since there is a large variance in the sizes of cores in a SoC. For example, a large multi-port shared memory core will be much larger as compared to a single function hardware block.

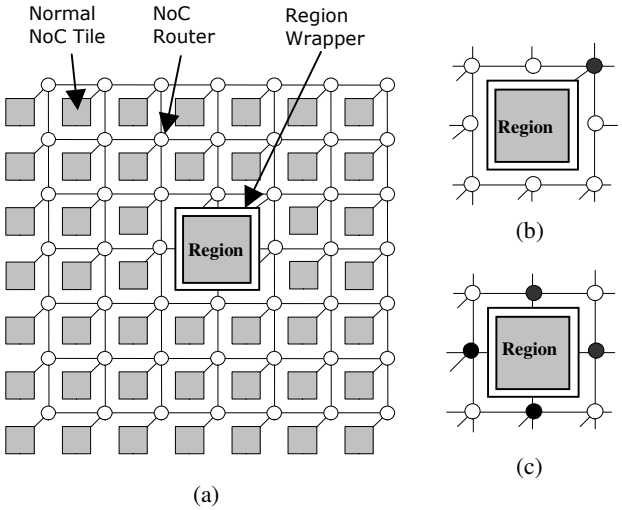

Figure 1. (a) Region within a NoC. (b) Single Access Point for Region. (c) Multiple Access Points for Region

To overcome this problem the concept of a specialized region within the NoC was presented by Kumar et al. [1]. Fig. 1a shows the definition of a region inside a mesh topology NoC. Wrappers are responsible for connection between the region and the NoC. Whereas normal sized resources are connected only to one router, regions can be connected to several *access points* due to their large size. The region acts like an obstacle when introduced in a mesh NoC. As a consequence, special routing algorithms have to be used in order to preserve connectivity between nodes, while ensuring freedom from deadlocks.

Bolotin et al. [2] presented a similar possibility to allow regions and proposed X-Y routing extended with hard-coded paths for region-affected traffic, to ensure deadlock free communication. As another solution, routing algorithms from the area of fault-tolerant routing in mesh networks have been evaluated for routing in mesh topology NoC with regions [3]. The paper presented new design issues and applications related to regions, as well as simulation results, which showed that position, size and orientation of regions strongly affect network performance. In this paper, we use the same routing algorithm as in [3] for a study on accessing regions. We compare two different configurations of access points for a 7X7 NoC with region of size 4 slots as shown in Fig. 1 (b) where the upper right corner router serves as a single access point, and Fig. 1(c) where 4 routers are used as access points (one router on each side). Fig. 2 shows the simulation results for these two cases plus, for reference, latency for a case without regions. We see that using one access point on each side results in about 5-15% lower latency, for load values stretching from 5-20%. This configuration also reaches the saturation point later. Because the region obstructs traffic passing through it, none of the region cases can match the case without region. The reason for the better performance using several access points is that the average distance is shorter between the sources and destinations in the network. This is because packets emitted from, or destined to the region do not have to circumvent it to reach the access point. Considering larger regions it is likely that the advantage of using multiple access points will be even greater.

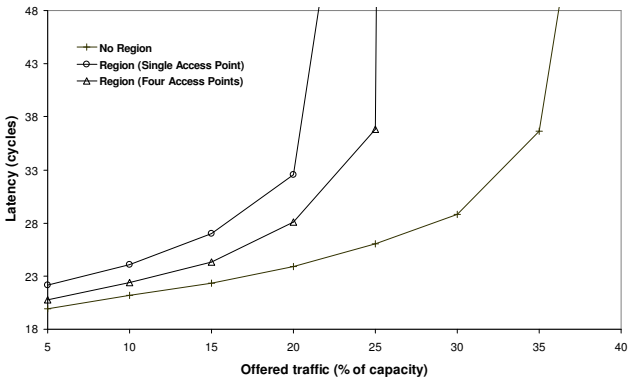

Figure 2. Access Points Effect on Latency

This enhancement in performance using multiple access points come at some cost. Routing protocol for regions with multiple access points will require extra bits to decide whether a packet is to be delivered to the access point of a region or a regular resource. This overhead depends on the number of access points. Multiple access points will also result in extra-hardware for multiplexing and buffering data received through various physical access ports and resolving conflicts. Extra hardware is also required for de-multiplexing messages leaving the regions. In conclusion, we can say that region concept is very useful for mesh topology NoC architectures. Providing multiple access points to large cores can help general network communication performance, however, there is some extra hardware and protocol cost.

REFERENCES

[1] S. Kumar, A. Jantsch, J-P. Soininen, M. Forsell, M. Millberg, J. Oberg, K. Tiensyrja and A. Hemani, "A Network on Chip architecture and Design Methodology", ISVLSI 2002

[2] E. Bolotin, A. Morgenshtein, I. Cidon, R. Ginosar and A. Kolodny, "Automatic Hardware-Efficient SoC Integration by QoS Network on Chip," ICECS 2004

[3] R. Holsmark and S. Kumar, "Design Issues and Performance Evaluation of Mesh NoC with Regions", Norchip 2005

Spatial Division Multiplexing: a Novel Approach for Guaranteed Throughput on NoCs

Anthony Leroy (anthony.leroy@imec.be), Paul Marchal, Francky Catthoor, Diederik Verkest, Frédéric Robert.

IMEC VZW, Heverlee, Belgium
Universitée libre Bruxelles, Bruxelles, Belgium
Katholieke Universiteit Leuven, Belgium

To ensure low power consumption while maintaining flexibility and performance, future Systems-on-Chip (SoC) will have to combine several types of processor cores and data memory units of widely different sizes.
To interconnect the IPs of these heterogeneous platforms, Networks-on-Chip (NoC) have been proposed as an efficient and scalable alternative to shared buses.
NoCs can provide throughput and latency guarantees by establishing virtual circuits between source and destination. In most NoCs, IP-blocks are connected to their own router through a network interface. Routers are interconnected to each other by point-to-point links to form a given network topology (e.g. mesh, torus,...). Their role is to forward the data from the source to the destination IP.
In real-time systems, many IP-blocks are subjected to performance/throughput constraints. One very simple way of providing guarantees on throughput and latency between two IP blocks consists of establishing a virtual circuit. This virtual circuit is exclusively dedicated to communication between them. Multiple virtual circuits can share the same physical communication resources (e.g. links). This concept is known as *Switched Virtual Circuit (SVC)*.
The best-known approach to implement SVC is Time Division Multiplexing (TDM). In this scheme, the time is discretized in equally long periods of time called time-slots. During a time-slot, the available bandwidth is exclusively dedicated to a given virtual circuit. Network resources are thus shared consecutively in time among the different circuits.
The main problem with TDM is precisely that the switching configuration of the router has to be updated for each time-slot. Thus, local configuration memories have to be implemented within routers resulting in high area and energy overhead. TDM also imposes tight scheduling constraints on the reservation of circuits.
We propose a solution that implements SVC with *Spatial-Division Multiplexing (SDM)*. This exploits the fact that on-chip network links are physically made of a set of wires. SDM consists of allocating only a sub-set of the link wires to a given virtual circuit. Messages are thus serialized on a group of wires. The switch configuration is set once and for all at the connection set-up. No inside-router configuration memory is therefore needed and the constraints on the reservation of the circuits are relaxed.
The main contribution of this work is to introduce the SDM technique in the context of NoCs and to propose a complete communication architecture exploiting SDM that includes routers and network interfaces. We have validated our technique with an RTL level implementation of the complete network. Simulations have been performed with a realistic case study based on a video application.
This case study clearly illustrates the advantages of our technique over TDM in terms of energy consumption, area overhead, flexibility and reliability. The SDM technique performs better in terms of area overhead and energy consumption than the traditional TDM technique. SDM thus appears as a very valuable alternative to TDM that is worth to be explored in more depth as well as a combination with TDM in a hybrid scheme.

Evaluation of On-chip Networks Using Deflection Routing

Zhonghai Lu, Mingchen Zhong and Axel Jantsch
Department of Electronic, Computer and Software Systems
Royal Institute of Technology, Sweden
{zhonghai,axel}@imit.kth.se, {mingchen@kth.se}

Deflection routing, also called *hot potato* routing, is being proposed for Networks on Chip (NoC) since it is simple and adaptive. A deflection switch can be much smaller and faster than a wormhole or virtual cut-through switch. As projected, a deflection switch can run 2.38 GHz with a gate count of 19370 in 65 nm technology. Determining the architecture of a deflection network is not a simple task because there exist many alternatives concerning network topology, routing algorithm and deflection policy. In addition to *network diameter*, *switch degree*, *link capacity*, the topology impacts two important deflection-network properties, *deflection index* and *don't-care density*. *Deflection index* is the largest number of hops that a single deflection adds to a packet's shortest path. *Don't-care density* is the percentage of destination nodes to which a source node has more than one non-overlapped shortest path to send packets. A packet is termed *don't-care* if it has more than one non-overlapped shortest path from current node to its destination. A routing algorithm determines the favorable path of packets, whereas a deflection policy resolves contentions for links.

Using the **N**ostrum **N**oC **S**imulation **E**nvironment (**NNSE**), we have evaluated those design alternatives that are being proposed for on-chip networks. Specifically, we consider three 2D topologies such as 2D mesh, 2D torus, and the Manhattan Street Network (MSN), four routing algorithms such as random, dimension XY, delta XY and minimum deflection, as well as three deflection policies such as non-priority, weighted priority, and straight-through policies. The MSN has a torus structure but unidirectional links. The dimension XY tries to route packets first along the X axis then the Y axis. The delta XY routes packets according to the difference between a packet's source and destination address, i.e., the Δ steps along the X and Y axis. The minimum deflection minimizes the number of deflections at each hop. The non-priority policy misroutes packets with equal probability. The weighted priority rule makes misrouting decisions in favor of higher priority packets. The priority of a packet is calculated according to a weighted expression involving its age, distance, deflection times and default value,. The straight-through policy favors a straight-through packet against a packet requiring a turn on its current hop. A deflection policy may or may not take the *don't-care* property of packets into account.

The networks are evaluated with random traffic. Packets are injected at a constant rate. Statistics are collected at the steady state. The torus performs best since it has the highest link bandwidth and toroidal boundaries. But if normalizing throughput (packets/cycle) with the link capacity, it turns out the mesh has the highest throughput per link before saturation. The minimal routing algorithms outperform the non-minimal ones. We find that a deflection policy taking advantage of don't-care packets reduces deflections. On the mesh with XY routing, it enhances the saturation throughput by 24% and meanwhile decrements the average network delivery time about 1 cycle by 19% (Figure 1(a)). It also improves the worst-case behavior. As shown in Figure 1(b) for injection rate 0.7, the observed maximum deflection count for the priority without don't-care, priority with don't-care and minimum deflection policies is 36, 6, 19, respectively.

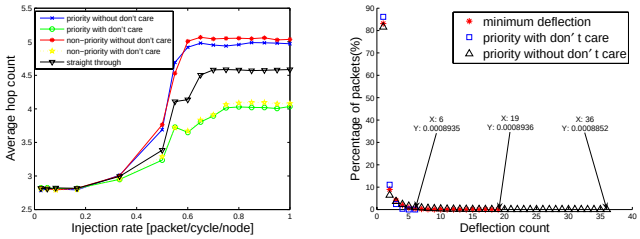

(a) Network delivery time (b) Deflection distribution

Figure 1. Performance considering don't-care

We can conclude from our results that the network topology is the most significant factor than the other two parameters, since it directly determines the deflection index and don't-care density. As long as a routing algorithm chooses the shortest path, the performance of different algorithms is close to each other. To improve the average-case/worst-case performance and resolve livelock, packets should be globally and historically (consider delivery history) prioritized, and a deflection rule should use the priority information and take the don't-care property of packets into account. We believe these conclusions can be guidelines to determine a deflection network architecture, for instance, selecting a routing algorithm or deflection policy which has potentially low cost and high speed for hardware implementation.

Dynamic Time-Slot Allocation for Networks on Chip with TDMA QoS

T. Marescaux[1], B. Brické[1,2], P. Debacker[1,2], V. Nollet[1], H. Corporaal[3]
[1]IMEC V.Z.W., Kapeldreef 75, 3001 Leuven, Belgium
[2]Katholieke Universiteit Leuven, ESAT, Leuven, Belgium
[3]Technical University Eindhoven (TU/e), Eindhoven, The Netherlands
Theodore.Marescaux@imec.be

In order to meet the ever-increasing design complexity, future sub-100nm platforms [1, 2] will consist of a mixture of heterogeneous computing, memory and I/O resources. These Multi-Processor Systems on Chip (MP-SoC) are expected to use flexible and scalable switched communication architectures such as a Network-on-Chip (NoC) [2, 3]. NoCs can provide different classes of communication -more generally called *Quality of Service* (QoS)- with real-time guarantees application designers can rely on.

The class of NoCs we are considering offers hard-guarantees in terms of bandwidth and/or latency coupled to best effort traffic. Providing hard-guaranteed QoS requires exclusive reservations of resources on the communication channels. Reservations happen either in space -typically buffer space is reserved on the routers to a particular guaranteed communication- or in time -communication channels are multiplexed in time and time-slices are allocated to a particular guaranteed communication. The latter technique is called Time Division Multiplexing Access (TDMA) and proves to be more resource-efficient when providing several levels of QoS.

This poster presents, to our knowledge, the first algorithm to dynamically perform routing and allocation of guaranteed communication resources on NoCs that provide QoS with TDMA techniques. We test the efficiency of our algorithm by allocating the communication channels required for an application composed of a 3D pipeline and an MPEG-2 decoder/encoder video chain on a 16 node MP-SoC. Dynamism in the communication is created by the 3D application.

The process of allocating guaranteed communication resources requires finding an optimal or nearly-optimal route through the network from source to destination while ensuring a contention-free time-slot scheduling. We show that the complex allocation problem of spatial and temporal resources that needs to be performed at run-time to establish new guaranteed connections can be reduced to the traversal of the TDMA-NoC graph we introduce. We have studied traversal algorithms, running on a central StrongARM processor clocked at 200 MHz, and found that our extended version of the IDA*algorithm provides the best results. To establish the 40 guaranteed connections of a 3D pipeline to render a complex scene, our algorithm requires between 450 to 900 μs, depending on the slot table size, or about 1000 cycles per hop. The memory foot-print of our algorithm is under 2.5 KB, which makes it practical to implement on an embedded processor such as the StrongARM, that have cache sizes of about 16 KB. Central time-slot allocation algorithms are practical for small-scale MP-SoC systems.

References

[1] H. De Man: On Nanoscale Integration and Gigascale Complexity in the Post .Com World. Proc. DATE 2002, Paris, France, March 2002.

[2] S. Kumar, A. Jantsch, M. Millberg, J. berg, J. Soininen, M. Forsell, K. Tiensyrj, and A. Hemani, "A network on chip architecture and design methodology," in Proceedings, IEEE Computer Society Annual Symposium on VLSI, Apr. 2002.

[3] A. Jantsch and H. Tenhunen, "Will Networks on Chip Close the Productivity Gap", Networks on Chip, Kluwer, 2003.

Application Parallelism Exploitation Using NoC with Multithreaded Processors

Ruxandra Pop and Shashi Kumar

Department of Electronics and Computer Engineering, School of Engineering, Jönköping University,
551 11 Jönköping, Sweden, +46 36 15 77 00, {Ruxandra.Pop, Shashi.Kumar}@ing.hj.se

Increasing the timing performances of concurrent applications which are running on Network on Chip (NoC) architectures is mainly correlated with the ability of mapping and scheduling methodologies for exploiting the Thread Level Parallelism (TLP) of concurrent applications through the available NoC parallelism. Matching the architectural parallelism to the application concurrency for obtaining good performance-cost tradeoffs is another aspect of the problem.

In our previous work [1], we have demonstrated that Multithreaded Processors (MTP) are good NoC resources for exploiting the TLP of concurrent applications, through the overlapped execution of several tasks with the aim of hiding the latencies of accesses to the local memory.

In this work, we explore the parallelism of NoC architectures with MTP resources, aiming to find the right number of MTPs which will provide the best performance-cost tradeoffs for a certain set of concurrent applications. For this purpose, we use the off-line methodology described in [1] which performs the task mapping and scheduling of concurrent applications to NoC with MTP and General-purpose Processor (GP) resources, such that the soft timing constraints are met and the timing performances are maximized. The methodology enforces the utilization of the available on-chip parallelism for exploiting the TLP of concurrent applications along with minimizing the inter-core communication volume. Figure 1 shows the application and the architecture models for this methodology.

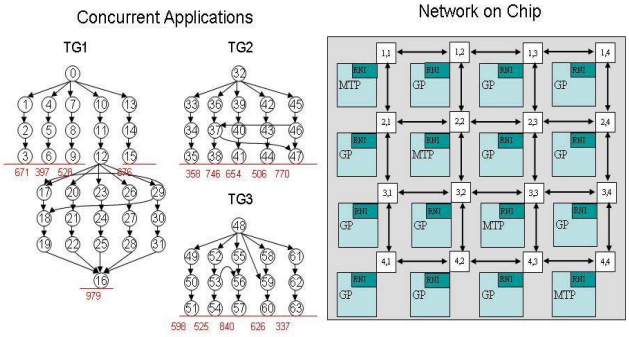

Figure 1 – Application Model and Architecture Model

The set of concurrent applications is specified as a set of task graphs, where each task graph (TG) represents an application. The concurrency is given by the total number of parallel execution paths among all TGs in the TG set. For example, the set of applications with series-parallel structures in Figure 1 has a TLP of 19.

The NoC architecture has MTP and GP resources interconnected in a two-dimensional mesh topology. The MTP model is that described in [1]. The parallelism of NoC with MTPs is given by the total number of thread contexts over all MTPs and single-threaded GPs. For example, the NoC in Figure 1 has a parallelism of 24.

The experiments were performed on eight synthetic benchmarks generated with TGFF 3.0 [2] and a 4×4 NoC with 1, 4, 8, 12 and 16 MTPs respectively, each of 3 thread contexts. The benchmarks have up to 100 tasks, 3 TGs and a maximum TLP of 7 per TG in the TG set. Figure 2 shows the variation of speedup (on average) when increasing the parallelism of NoC by adding more MTPs, for benchmarks with random interconnectivity (dashed line) and with series-parallel structures (solid line) respectively.

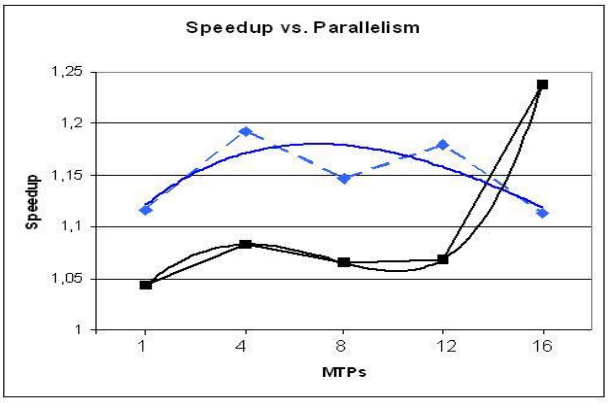

Figure 2 – Speedup vs. Parallelism

As it can be seen, the speedup trendline is almost ideal for general TGs. The speedup is increasing to 19% for 4 MTPs and after that is slowly decreasing to 11% for 16 MTPs, reaching the same value as for 1 MTP.

The speedup trendline for parallel TGs suggests that only homogeneous NoC with solely MTPs is able to fully exploit the TLP of series-parallel structures. Thus, the speedup is around 7% for 1 to 12 MTPs and then it boosts to 23.7% for 16 MTPs. A similar behaviour is yielded for the execution on a 3x3 NoC with 1, 3, 6, 9 MTPs, respectively.

References

[1] R. Pop and S. Kumar, Mapping Applications to NoC Platforms with Multithreaded Processor Resources, in Proc. IEEE Norchip Conference, Oulu, 21-22 Nov 2005.

[2] R. P. Dick, et al., TGFF: Task Graphs For Free, Int. Workshop Hardware/Software Codesign, Mar. 1998.

Curing Hotspots in Wormhole NoCs

Isask'har Walter[1], Israel Cidon[2], Ran Ginosar[2], Avinoam Kolodny[2]
Electrical Engineering Department, Technion, Haifa, Israel
[1]zigi@tx.technion.ac.il [2]{cidon, ran, kolodny}@ee.technion.ac.il

Presenting author: Isask'har Walter

Abstract

In a network on chip (NoC) based system, modules may operate in various utilization rates. A bandwidth-limited system on chip (SoC) module working close to its capacity is termed a *hotspot*. In this situation, the module is unable to consume incoming packets fast enough, and the entire network may be affected. As buffers at the router adjacent to the hotspot module are filled up, new arrivals to this router are blocked. This creates a domino effect, by which the delivery of packets to ports of more distant routers is slowed down, forming a *saturation tree* with the hotspot module as its root, as illustrated in Figure 1. The NoC suffers increased delays in packet delivery and unfair network utilization.

While saturation trees exist in store-and-forward networks, the threat is particularly troublesome in wormhole switching that is commonly used by NoC architectures due to its low buffer requirements communication latency. "Stretching" of packets across several hops and the backpressure mechanism causes hotspot effects to extend network wide instantly. It is important to note that the hotspot phenomenon is independent of links bandwidth, as a saturation tree may build up in a system with infinite capacity links and a single heavily loaded module. Consequently, even carefully designed, largely over-provisioned NoCs may suffer of poor performance if potential hotspots are left unhandled, as a single hotspot may ruin the performance of the entire network. Unfortunately, hotspots are common in real-life SoCs, as internal components (caches, CAMs, special purpose processors) or external DRAM interfaces are bandwidth limited and in high demand by other units.

We propose a novel credit-based resource allocation mechanism for solving hotspot congestion problems in wormhole-based NoCs. A hotspot allocation controller is introduced to arbitrate short, high priority credit requests. The controller regulates hotspot access according to the quality of service requirements of the specific system. Credit requests and grants are transmitted as high-priority packets (grants may be piggybacked on other messages). The mechanism prevents the accumulation of packets destined at a hotspot module in the network, while the NoC routers remain unchanged. Consequently, other traffic remains unaffected even when the hotspot load increases significantly. Using simulations, we show the effectiveness of the suggested mechanism by evaluating the system's performance and fairness.

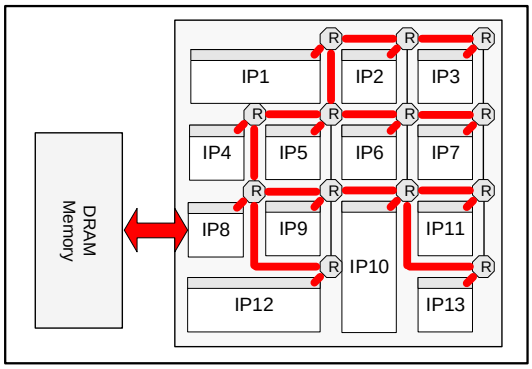

Figure 1: SoC hotspot at external DRAM interface and the resulting NoC saturation Tree

STAR : An Efficient Routing Strategy for NoC with Mixed QoS Requirements

Daniel Andreasson and Shashi Kumar
Department of Electronics and Computer Engineering, School of Engineering, Jönköping University
Daniel.Andreasson@ing.hj.se and Shashi.Kumar@ing.hj.se

Introduction to STAR: Network-on-Chips will be a heterogeneous environment, where processing elements (PE) of different types create the required functionality through communication. Such traffic environments are known to exhibit bursty communication (e.g. multimedia applications). The drawback of such traffic is that the demand for bandwidth will vary greatly. Nevertheless, resources must be able to handle worst-case demand. This means that network resources are seldom fully utilized.

Traffic in NoC can at least be classified into two types: guaranteed traffic and best-effort traffic [3, 4]. The guaranteed traffic reserves bandwidth, and can offer QoS guarantees. The best effort class has no such performance guarantees. However, the cost of guaranteed traffic, beside the extra architectural resources required, is the underutilization of reserved resources. We define the duration where the producer does not use its reservation as the slack-time. To reduce the cost of underutilization we propose a strategy named STAR for Network-on-Chips. STAR stands for **S**lack-**T**ime **A**ware **R**outing, and is an efficient routing strategy and a collection of schemes designed to, by clever planning, make best-effort traffic efficiently use the slack-time of the guaranteed traffic. The guaranteed traffic producers detect and estimate slack-time with the help of a special playout stage buffer [1]. Based on the slack-time information they can take appropriate routing decisions.

The STAR framework: STAR currently lists three ways in which the slack time can be managed and used for improving the best-effort traffic performance.

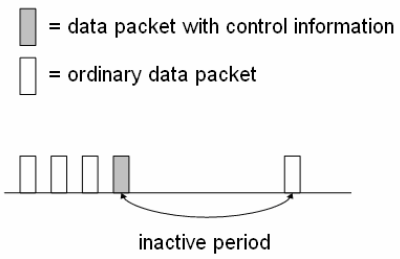

Figure 1. Bursty guaranteed packet production

1) The first is a general scheme that lets the best-effort traffic use the spare capacity of the reservations. The producer of guaranteed packets issues special control packets with the amount of time that packets will not be sent. This is depicted in *Fig.1*. The router can then let the best-effort traffic use the reservation during that time.

2) When the total worst-case capacity is reserved for a producer, the capacity can be divided on several paths. *Fig 2* conceptually shows how the reserved capacity is divided on two paths. As can be seen, when the real traffic rate is lower than the highest rate, it is possible to distribute slack-time in different ways. In *Fig. 2*, the second path is favoured at the expense of the first path. The scheme includes a function to distribute the slack-time over multiple-paths, according to the demand of best-effort traffic. And if necessary, it is possible to change the slack-time distribution during run-time [1].

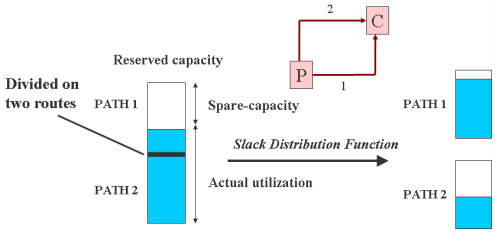

Figure 2. Example of Slack-Time Distribution over Two Paths

3) Slack-time can also be distributed on a per-hop basis over a single-path. The scheme assumes that both the best-effort and the guaranteed traffic are bursty. During the periods of high best-effort traffic load, the guaranteed packets can wait for a maximal defined time, and more best-effort packets can pass a router in the direction of the reserved path. Thus, the jitter of best-effort traffic is improved at the expense of the guaranteed traffic. *Fig. 3* shows an example of such improvement from a simulation of the scheme in a Network-On-Chip [2].

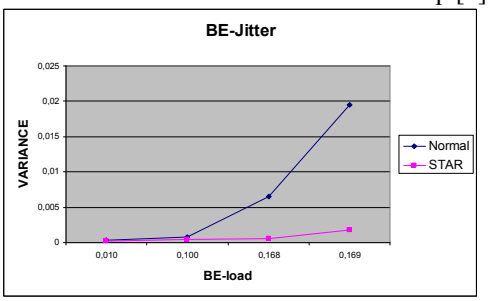

Figure 3. Jitter Improvement of Best Effort Traffic

Conclusion: We have proposed the strategy of STAR for on-chip networks with mixed QoS Requirements. It uses clever planning of the slack-time of guaranteed traffic. Slack-time management involves important issues as the buffering of packets and maintaining and updating available slack-time information. This implies that network routers and resources have special facilities for handling those issues.

References

[1] Daniel Andreasson and Shashi Kumar, "Slack-Time Aware Routing in NoC Systems", Proceeding of International Symposium of Circuits and Systems, May 23-26, Kobe, Japan.

[2] Daniel Andreasson and Shashi Kumar, "Improving BE traffic QoS using GT slack in NoC Systems", Proceeding of 23th IEEE Norchip Conference, 21 - 22 Nov 2005, Oulu, Finland

[3] M. Millberg, E. Nilsson, R. Thid, and A. Jantsch, "Guaranteed bandwidth using looped containers in temporally disjoint networks within the Nostrum network on chip", In DATE of 2004.

[4] E. Rijpkema, K. Goossens, et al, "Trade offs in the design of a router with both guaranteed and best effort services for networks on chip", In DATE of 2003.

Assessing Adaptive Routing Behaviour in NoCs Through Cycle-Accurate Functional Simulation

Federico Angiolini†, Luca Lerin‡,Davide Bertozzi‡,Luca Benini†

† DEIS, University of Bologna, 40136 Bologna, Italy

‡ ENDIF, University of Ferrara, 44100 Ferrara, Italy

fangiolini@deis.unibo.it, llerin@ing.unife.it, dbertozzi@ing.unife.it, lbenini@deis.unibo.it

I. Abstract

Most Network-on-Chip (NoC) prototypes reported in the open literature make use of deterministic routing [6], [5], thus ignoring path diversity exposed by the topology and achieving poor load balancing. There are many reasons behind this design choice: (i) statically-routed networks are easier to make deadlock-free; (ii) no risk of out-of-order packet delivery and no need for large buffering resources to reorder packets; (iii) adaptive routing is associated with increased routing cost, delay and jitter; (iv) unpredictable performance with respect to wise IP cores mapping and routing channel allocation. Overall, the work in [2] advocates for deterministic, deadlock-free, minimal and wormhole-based routing.

However, there are some cases where adaptive routing could be required by the distinctive features of specific NoC implementations. First, it could allow to effectively manage the challenging problem of communication among modules dynamically placed on a reconfigurable device (the dynamic network-on-chip concept [1], [3]). Second, adaptive routing should be used when addressing the problem of broken links or malfunctioning routers [4]. Third, it could be an effective option in presence of unpredictable traffic patterns at design time.

Beyond this intuitive selection of application domains that best match the features of adaptive routing algorithms, no accurate assessment of their behaviour in NoC prototypes has been reported yet. This work sheds light on pitfalls, fallacies and potential benefits of applying adaptive routing to packet-switched NoC architectures, leveraging cycle-accurate functional simulation. We compare deterministic versus adaptive routing schemes from the perspective of application-perceived performance. Parametric traffic generators have been employed to analyze the behaviour of the routing schemes under different traffic patterns, and topologies with varying complexity have been considered.

Since adaptive routing needs routing decisions to be taken dynamically at the switches, we modified the ×pipes [6] NoC architecture to support distributed routing and traffic balancing. In particular, we designed an adaptation mechanism based on a sliding window which exploits the ACK/NACK flow control policy natively supported by ×pipes. Each switch output port has a sliding window, whose content is shifted at each clock cycle. If an ACK is received from that link, a 0 is pushed into the window, otherwise a 1 is written (see Fig. 1(a)). The distributed routing strategy is based on the destination field of the incoming packet and on the content of a local routing table, which indicates the different alternative paths that can be taken to reach that destination. The path is selected by watching the content of the sliding windows of the alternative paths, as depicted in Fig. 1(b), and by forwarding the packet to the less congested link. The congestion metric is relative to the most recent time slice; the length of the latter is proportional to the depth of the monitoring window. Idle cycles are considered to be equivalent to ACKs, to provide a chance to reuse over time a link that had previously been heavily congested.

Since the effectiveness of adaptive routing algorithms depends on the interaction of a number of factors (*e.g.*, relative length of minimal versus non-minimal routing paths, features of the traffic patterns, amount of buffering devoted to the adaptation mechanism and sensitivity to traffic variations, detection of local versus non-local traffic congestion) a first set of experiments was carried out on a simple topology that allowed to isolate each contribution and to explore design parameters. Our analysis points out some interesting trade-offs among length of non-minimal paths and congestion levels on minimal paths.

Then, we applied the adaptive routing scheme to a more realistic mesh topology. In this scenario, well-known problems of adaptive routing schemes (*e.g.*, poor adaptation to traffic conditions occurring far away from the decision point) were quantitatively assessed, and some workarounds (involving higher implementation cost) were devised. Moreover, we pointed out second-order effects that might make the behaviour of the sliding window mechanism dependent on non-congestion related parameters, and hence unpredictable. However, significant savings on execution times were obtained when congestion could be directly sensed by the adaptation mechanism and less congested (even though non-minimal) paths could be taken. We finally came up with considerations about applicability of the simple sliding window mechanism and of fully-adaptive routing to real NoC architectures, identifying the key challenges that still need to be addressed to make this scenario a viable option for NoC routing.

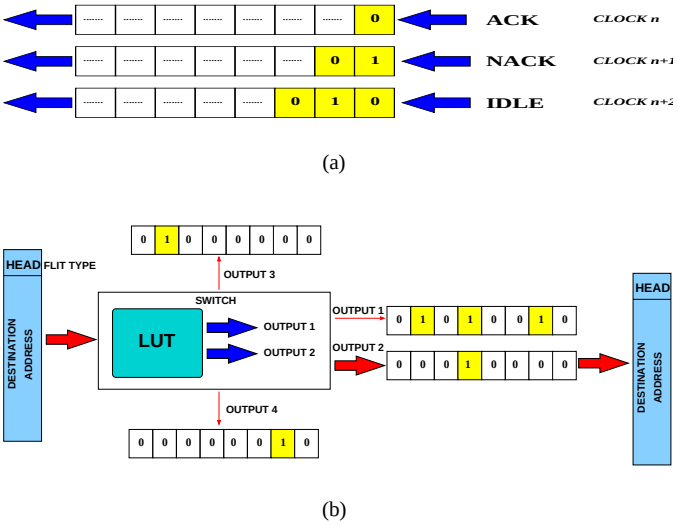

Fig. 1. (a) Sliding window mechanism, (b) LUT-based distributed routing.

References

[1] Christophe Bobda et al. A Dynamic NoC Approach for Communication in Reconfigurable Devices J. Becker, M. Platzner, S. Vernalde (Eds.): FPL 2004, LNCS 3203, pp. 1032-1036, 2004, Springer-Verlag Berlin Heidelberg 2004

[2] J. Hu, R. Marculescu, Exploiting the routing flexibility for energy/ performance aware mapping of regular NoC architectures, *CSSI Technical Report, Carnegie Mellon University, Sept. 2002.*

[3] Ali Ahmadinia et al. A Practical Approach for Circuit Routing on Dynamic Reconfigurable Devices? *International Workshop on Rapid System Prototyping (RSP 2005).*

[4] Muhammad Ali, Michael Welzl, Martin Zwicknagl, Sybille Hellebrand, Considerations for fault-tolerant Network on Chips *IEEE ICM'05, Islamabad, Pakistan, 13-15 December 2005.*

[5] Kees Goossens, John Dielissen, Andrei Radulescu, Aethereal Network on Chip: Concepts, Architectures, and Implementations *IEEE Design and Test of Computers, September/October 2005 (Vol. 22, No. 5) pp. 414-421.*

[6] Davide Bertozzi, Luca Benini, ×pipes: A NoC architecture for gigascale Systems-on-CShip. *IEEE Circuits and Systems Magazine*, 2004.

Performance Enhancement through Early Release and Buffer Optimization in Network-on-Chip Router Architectures *

Jongman Kim *Dongkook Park* *Chrysostomos Nicopoulos* *N. Vijaykrishnan* *Chita R. Das*

Department of Computer Science & Engineering
The Pennsylvania State University
University Park, PA 16802
{*jmkim, dpark, nicopoul, vijay, das*}*@cse.psu.edu*

The crux of System-on-Chip (SoC) architectures embodying hundreds of functional units is the on-chip interconnect, which is becoming a critical bottleneck in meeting the performance and power consumption budgets of the chip design. Although some commercial SoCs use a shared bus for interconnecting the functional units, packet-based interconnection networks are being increasingly adopted in complex SoC designs. These Network-on-Chip (NoC) architectures are required to not only provide ultra-low latency, but also occupy a small footprint and consume as little energy as possible.

Towards the goal of designing area-constrained, low-latency and energy-efficient on-chip communication networks, we propose two novel optimization techniques for on-chip routers, based on look-ahead routing. The first proposed technique employs Early Release - Early Ejection and Link-Ahead Injection - which leads to less actual transfer time and reduced silicon area and energy consumption. The second technique adopts adjustable buffers to enhance average utilization and balance flit occupancy, which, in turn, help mitigate blocking delay.

Early Ejection: A flit destined for the local PE does not traverse the crossbar, but, instead, it is ejected immediately upon arrival (hence the "Early Ejection" mechanism). This technique utilizes the look-ahead routing information to detect if the incoming flit is destined for the local PE, and, accordingly, ejects it after the input DEMUX. Early Ejection reduces routing delay at the destination node by avoiding both switch allocation and switch traversal. It also reduces the load on each crossbar input port.

Link-Ahead Injection: Similar to Early Ejection, a flit leaving the local PE need not go through the crossbar. Instead, it is directly MUXed into the crossbar output links ("Link-Ahead Injection"). This technique eliminates the local input queuing buffer assigned to newly injected messages, since buffering is now offloaded to the network interface unit. Furthermore, the crossbar no longer requires a dedicated input port for the local PE. Control signals run ahead of packet injection to check output link and next-node buffer availabilities. Link-Ahead Injection reduces intra-router latency at the source node by bypassing both input buffering and crossbar traversal.

Both the Early Ejection and Look-Ahead Injection techniques (collectively called the Early Release mechanism) provide significant advantages to nearest-neighbor traffic, and can reap substantial performance and energy benefits when used in conjunction with NoC mapping which places frequently communicating PEs close to each other. Early Release also provides graceful degradation in the event of router failure; the local PE can avoid isolation from neighbors, since the network interface unit is directly connected to both the Early Ejection and Link-Ahead Injection channels.

Adjustable Buffer: Given the relatively limited buffer resources in NoCs, as compared to macro-networks, we closely examined existing buffer systems to determine if there are any possible optimizations. If the routing algorithm used is not suitable for predictable design-time optimizations, a dynamically allocated central buffer can, instead, offer similar performance benefits. In traditional macro-networks, a centralized buffer system allows for greater overall buffer space utilization, by dynamically allocating the space to those ports which require it the most. We investigated a form of centralized buffer system that is amenable to NoC designs. In traditional centralized buffer systems, one large bank of memory is shared using a large number of ports, while a complicated memory management, or pointer system, maintains the location of the data. In contrast, the proposed adjustable buffer system utilizes a low-overhead, dual-port memory architecture to achieve increased buffer utilization and performance with significantly reduced hardware complexity. This makes it suitable for resource-constrained NoC implementations.

In addition to improving performance, the proposed techniques optimize both buffer and crossbar sizes. Reduced buffer and crossbar sizes, in turn, translate into reduced energy consumption in buffer read/write operations and crossbar switching activity. Our simulation results, performed using a cycle-accurate simulator and a synthesized implementation of the router design in 90nm CMOS technology, show that the proposed mechanisms reduce packet latency by 32% and energy consumption by 21% compared to conventional router architectures.

*This research was supported in part by NSF grants CCR-0208734, EIA-0202007, CCF-0429631, CNS-0509251, CAREER 0093085, and a grant from DARPA/MARCO GSRC.

A.4 Routing methods:

- **"STAR: An Efficient Routing Strategy for NoC with Mixed QoS Requirements"**
  D. Andreasson and S. Kumar
  Jönköping University, Sweden
- **"Assessing Adaptive Routing Behaviour in NoCs Through Cycle-Accurate Functional Simulation**
  F. Angiolini[1], L. Lerinz[2],D. Bertozzi[2],L. Benini[1]
  1 DEIS, University of Bologna, 40136 Bologna, Italy
  2 ENDIF, University of Ferrara, 44100 Ferrara, Italy
- **"Performance Enhancement through Early Release and Buffer Optimization in Network-on-Chip Router Architectures"**
  J. Kim, D. Park, C. Nicopoulos, N. Vijaykrishnan and C. Das
  The Pennsylvania State University, USA

A.5 NoC Applications:

- **"Emerging Networks on Chip: Lessons Learned to Win the Future"**
  Iyad Al Khatib
  Royal Institute of Technology, Sweden
- **"Interoperability Protocol and Algorithms for Network-on-Chip Autonomous-System Communications for the Next Generation Biomedical Sensor-Networks"**
  Iyad Al Khatib[1], A. Jantsch[1], R. Nabiev[2] and L. Onana Alima[3]
  1 Royal Institute of Technology, Sweden
  2 Karolinska University Hospital Huddinge, Sweden
  3 Université de Mons-Hainaut, Belgium
- **"EEG Biomedical Sensors Network on Chip (EBSNOC)"**
  M. Saleh and A. Jantsch
  Royal Institute of Technology (KTH), Sweden
- **"Design of a simple clockless Network-on-Chip for an audio DSP chip"**
  M. Stensgaard[1], T. Bjerregaard[1], J. Sparsø[1] and J. Pedersen[2]
  1 Technical University of Denmark(DTU)
  2 William Demant Holding
- **"Application-Specific Networks-on-Chips"**
  J. Xu and W. Wolf
  Princeton University, USA

Emerging Networks on Chip: Lessons Learned to Win the Future

Iyad Al Khatib
Department of Microelectronics & IT
Royal Institute of Technology
Stockholm, Sweden
iyad@imit.kth.se

In order for industry in many different disciplines to adopt the NoC solution, NoC must be fit into future-market applications. For example, to try to produce a NoC always addressing multimedia may not prove commercially and industrially favorable, because such applications have already satisfactory market products, and IT investors are still paranoid (since the fall of the IT-market in late 2000) to invest in new fields. Our aim is to show how NoC can be useful and how industry can make successful products. Therefore, lessons learned from different fields related to NoC shall be investigated again in order not to repeat neither the technical nor market mistakes. The market mistakes have been many in different IT fields. However, there is clearly a multi-core trend rising now due mainly to the fact that the industry tries so hard to push for more efficiency and power form single processors. To reach the required efficiency industry has noticed a need for distributed many-cores. The NoC community shall make use of this rising need and push towards solutions in this trend, e.g. for computing and particularly "laptop" companies. Laptops are multi-purpose machines, and their operating systems seem to have a much faster development rate than their processing hardware. Moreover, their market is huge. That is not to forget, fields where size is important like in aerospace engineering. From a scientific viewpoint, a look at what has been done wrong in networking is essential. For instance, the networking field has been growing rapidly since the Internet and wireless communications boom in the last decade, but it suffers from problems like Internet protocol design, which wasn't made originally to carry multimedia but rather to carry text. Now engineers want to use this originally-data-and-text medium for nearly any communication purpose! Thus, it suffers from quality issues. To overcome such problems, we are developing a deep study on such networking protocols and working on listing the lessons to learn. One main result we see is that there is a need for a different addressing scheme, where in a NoC we will need a core-to-application addressing (Fig.1), which doesn't exist in any networking protocol (where addressing is always based on node-to-node). Such studies may save a lot of time and effort spent on experimenting in NoC and would relieve us from the burdens of the trial-and-error since we may have a working theory to deploy directly.

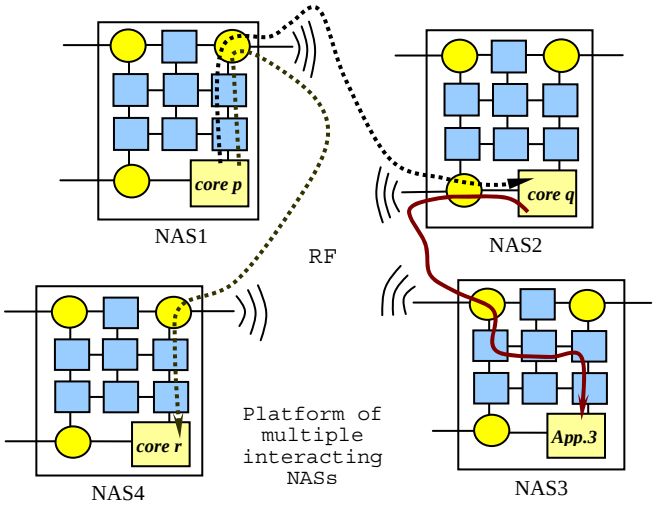

Figure 1. New addressing scheme showing core p sending a message, where there is a need to know to on which core the application resides (in this case Application *App.3*). It is the job of the routing protocol to find it and not the core (or node in Internet).

Interoperability Protocol and Algorithms for Network-on-Chip Autonomous-System Communications for the Next Generation Biomedical Sensor-Networks

Iyad Al Khatib, Axel Jantsch
Dept. of Microelectronics & IT
Royal Institute of Technology
Stockholm, Sweden
{iyad, axel}@imit.kth.se

Rustam Nabiev
Biomedical Engineering Dept.
Karolinska University Hospital
Huddinge, Stockholm, Sweden
rustam.nabiev@karolinska.se

Luc Onana Alima
Dept. of Distributed Systems
Université de Mons-Hainaut
Mons, Belgium
luc.onana@umh.ac.be

The interest in high performance chip architectures for biomedical applications is evident. We present a Network-on-Chip (NoC) inter-communication solution to the ever-growing demand for complex computations requiring huge numbers of processors in sensor networks of biomedical, healthcare, and medical emergency applications. The design of a single NoC and its protocols must take into consideration inter-NoC communications in the early design phase. We define a mechanism to enable many NoCs to interact. The advantages of this mechanism are: the ability to connect multiple different NoCs in a scalable manner, increasing computational capabilities, and NoC compatibility.

We look at NoC as an autonomous component in a sensor network. The future for NoCs may lead many vendors to realize the need to connect different NoCs built for various applications, which is only achieved if there were standard external NoC policies implemented on all interconnected NoCs. With the notion of the NoC as an Autonomous System (NAS), we define the boundary between internal NoC events and external NoC events. We define two main classes for NoC communications: (1) Intra-NoC protocols that deal with internal NoC data transfer via internal routing and switching, and (2) Inter-NoC protocols that define the rules and policies used by a NoC to interact with other NoCs, whether being of the same type and vendor or being of a different type and vendor (Fig. 1). We anticipate many Inter-NoC communication protocols to evolve, and thus propose the ENoP (External NoC Protocol) as the first protocol for inter-NoC communications. The ENoP is "integrated" with the on-chip communication protocol in the sense that individual resources, for instance, in the chip of NAS1 can be visible and addressable from the chip in NAS2. The visibility has to be controlled by the border core/gateway. Designing an Inter-NoC communication mechanism in this phase of the NoC growth will facilitate a faster evolution for NoC interoperability and compatibility. We could anticipate this coming demand when defining projects with biomedical and space industrial partners by evaluating their need for large computational platforms. In this work, we present our evolving mechanism and algorithms for external NoC communication. Our simulations show a millisecond convergence time for the ENoP.

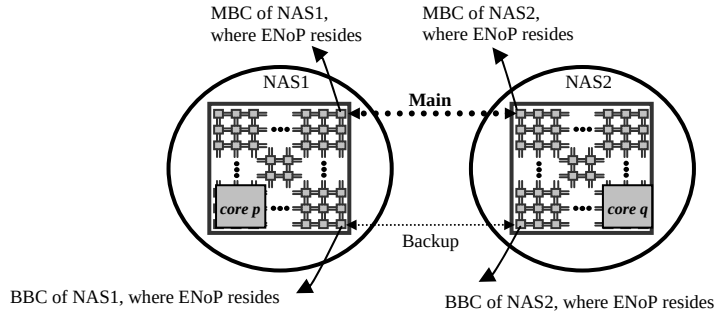

Figure 1. Two NASs interconnected via Main and Backup links. The Backup link is used in case of Main-link failure. The choice of which will be the Main and which will be the Backup is made by the ENoP (based on bandwidth, reliability, and packet loss tests). Main and Backup connections may be wired or wireless. MBC and BBC are the Main Border Core and Backup Border Core, respectively, and they are the cores where the ENoP resides. The ENoP communicates with internal cores using a vendor specific internal NoC protocol, and it aids the NoC to communicate with the neighboring NoCs on the external side. Sending a message from an application in *core p* on the first NoC asking for some processing of results from another application in *core q* on the second NoC has to pass through two Border Cores. For instance, NAS1 may be responsible for pressure/altitude variations (and many other aerospace dynamics) and NAS2 may be responsible for astronaut's heart analysis, where each application requires millions of multiplications per second (accurate heart analysis requires 1.75 million multiplications every 3.5 seconds).

EEG Biomedical Sensors Network on Chip (EBSNoC)

Mohammad Saleh and Axel Jantsch
Department of Microelectronics and Information Technology (IMIT)
Royal Institute of Technology (KTH)
Stockholm, Sweden
mohsaleh@kth.se, axel@imit.kth.se

Abstract- *We presents a novel idea of designing a wearable biomedical sensor network on chip as electroencephalogram (EEG) to monitor patient brain and record the brain activities in specific memory stick attached to the EBSNoC. This can be helpful to monitor brain activities for specific patient's case like epilepsy. The EBSNoC will provide a way to avoid risks and will save lives.*

1. Introduction

One of the brain diseases we focus on is epilepsy. The EBSNoC system will be used to monitor the patient's brain and analyze the brainwaves in real-time in order to alert the patient and send alarm signal when the EBSNoC detects that the brain is generating irregular wave patterns. The aim of this project is to design an inexpensive system to save patients' lives, provide a treatment at the right time before the epileptic seizure occurs.

2. EBSNoC System Architecture

The EBSNoC receives wireless signals from four sensors placed on the scalp. These sensors measure the brain rhythm. The patient can know his status by looking to the EBSNoC analyzer.

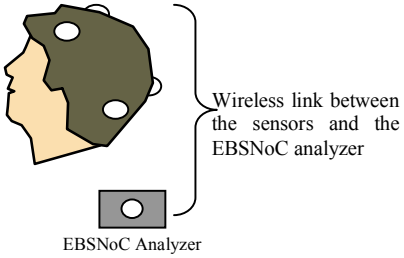

Figure 1: EBSNoC System

2.1 EBSNoC Analyzer

The EBSNoC analyzer is a program running on the NoC and designed to analyze brain signals. It consists of four modules. Each of these modules is implemented on a separate cluster of processors on the NoC. A central processor collects the results from each module and starts its analysis based on a medical algorithm. The most critical point of analysis is the floating point of the epilepsy at which a decision is made by the EBSNoC to send an alarm to the patient. If no floating point was detected, a GREEN light is ON. When a floating point is detected, an alarm is sent to the patient

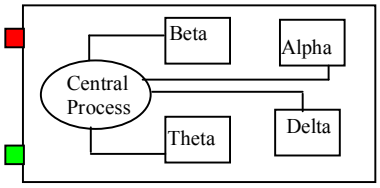

Figure 2: EBSNOC analyzer

2.2 EBSNoC Algorithm

The algorithm runs on four clusters of processors to manage and review the rhythmic wave status of measured data by checking the relative change in the characteristics of the normal brainwaves every time interval. The results of these parallel processes are be transmitted to an epileptic seizure module, which tests the rhythmic waves and compares them with a floating point of the epilepsy. Accordingly, the patient can be alarmed.

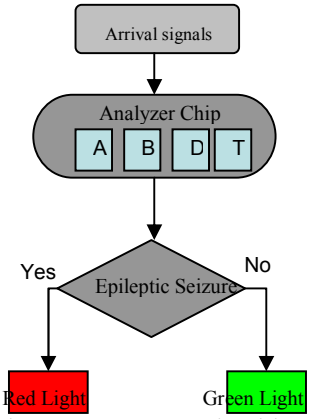

Figure 3: EBSNoC Algorithm

3. Conclusion

A wearable EBSNoC as a generic system for brainwave monitoring, where the functionality of the NoC can be specialized by means of updating a NoC software. Our main focus is on epileptic patients since our EBSNoC design can detect early signs of epilepsy and alarm the patient to take protective measures before epileptic seizure occurs. This measure can save lives by avoiding the risks of seizure, especially in cases of driving.

Design of a simple clockless Network-on-Chip for an audio DSP chip.

Mikkel Bystrup Stensgaard†, Tobias Bjerregaard†, Jens Sparsø†, and Johnny Halkjær Pedersen‡

†Informatics and Mathematical Modelling (IMM)
Technical University of Denmark(DTU)
Email: {mibs,tob,jsp}@imm.dtu.dk

‡Oticon
Email: jhp@oticon.dk

Abstract

We replace the communication infrastructure in an existing audio DSP chip with a packet-switched, clockless, Network-on-Chip (NoC). The two networks are place and routed and compared in terms of area, routing complexity, and power consumption. The study is interesting because it is a commercial chip, and that we are able to compare the NoC with the current communication infrastructure. Besides, it demonstrates that NoCs are feasible solutions, even for small applications with limited bandwidth requirements.

The DSP application consists of a number of audio processing blocks, connected by a circuit-switched network. The network is used to set up a dataflow between the blocks as illustrated in figure 1, and contains 13 inputs and 12 outputs. As only a subset of the blocks need to communicate, the original network is designed as a subset of a fully connected crossbar, using multiplexors at the outputs of the network. The multiplexors have between 2 and 5 inputs, which makes the network extremely small. This solution does not scale well, as the number of wires and size of the multiplexors increase dramatically for an increasing number of communicating blocks and as the communicating subset gets larger. Routing is also complicated due to the many long wires which also require buffers to drive. Last but not least, the network is application specific as not all blocks are able to communicate.

In contrast to most existing research, the NoC is designed as simple as possible due to the limited bandwidth and latency requirements. It is implemented as a bufferless binary tree of simple *router* and *merger* modules as illustrated in figure 2. It employs a parallel, bundled data protocol to minimize the power consumption and also include network adaptors to interface with the blocks, handle multicast and perform synchronization. It has some resemble to FLEETzero [2] and CHAIN[1], but is even simpler.

For this specific application the NoC uses approximately 3-4 times more area and power than the original network. This is still only 1% of the chip area and the total length of wire is decreased by 30%, easing the place and route process. Also, all blocks can communicate, the NoC scales linearly, and as it decouples the communicating blocks they are able to run at different frequencies (GALS). At last, we believe it shows that feature rich NoCs are not always needed, and a large group of applications can benefit from simple NoCs.

Acknowledgment

The authors thank *Oticon* for providing information about the audio DSP application.

References

[1] John Bainbridge and Steve Furber. CHAIN: A delay-insensitive chip area interconnect. *IEEE Micro*, 22:16–23, 2002.

[2] William S. Coates, Jon K. Lexau, Ian W. Jones, Scott M. Fairbanks, and Ivan E. Sutherland. FLEETzero: An asynchronous switch fabric chip experiment. In *Proc. International Symposium on Advanced Research in Asynchronous Circuits and Systems*, pages 173–182. IEEE Computer Society Press, March 2001.

[3] M. B. Stensgaard. Design of an asynchronous communication network for an audio DSP chip. http://www2.imm.dtu.dk/pubdb/p.php?3974, Master thesis 2005.

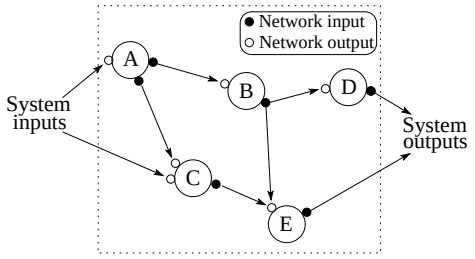

Figure 1: The network sets up a dataflow between the different audio processing blocks.

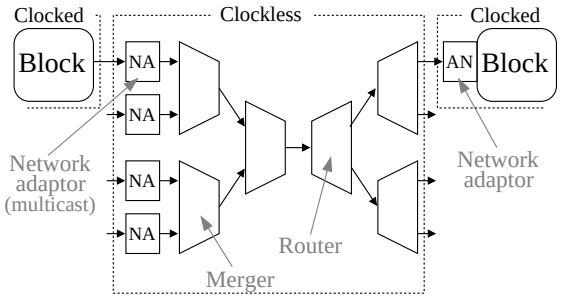

Figure 2: The NoC is constructed from simple *router* and *merger* modules. The communicating blocks interface the NoC through network adaptors which handle multicast and synchronization.

Application-Specific Networks-on-Chips

Jiang Xu (jiangxu@princeton.edu), Wayne Wolf (wolf@princeton.edu)
Dept. of Electrical Engineering, Princeton University
Joerg Henkel
University of Karlsruhe
Srimat Chakradhar
NEC Research

Many networks-on-chips are designed separately from the multiprocessors in which they will be used. We believe that application-specific networks, designed for a particular architecture, can provide much more efficient networks. Application-specific, irregular topologies can save energy and area while providing performance at least as good as that provided by regular-topology networks. This poster will cover some results discussed in a paper accepted to the MPSOC special issue of ACM TECS.

This poster will discuss our results on two design examples: a smart camera design and an H.264 decoder. In both cases we used simulation to characterize the traffic in the architectures. Here is a block diagram for H.264:

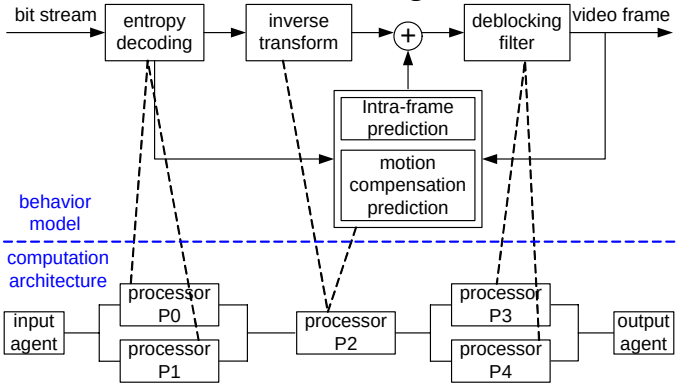

Here are some workload characteristics for H.264:

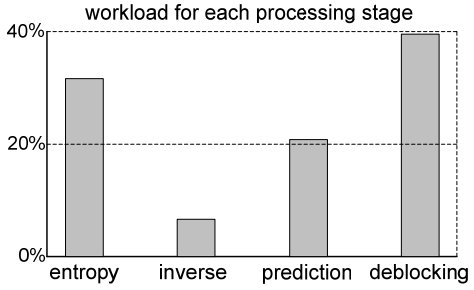

Based on those characteristics, we designed specialized topologies to match the application data requirements. We used floorplans to evaluate physical characteristics. We compared these application-specific networks to RAW-style mesh networks.

These results compare application-specific and mesh networks for the H.264 decoder:

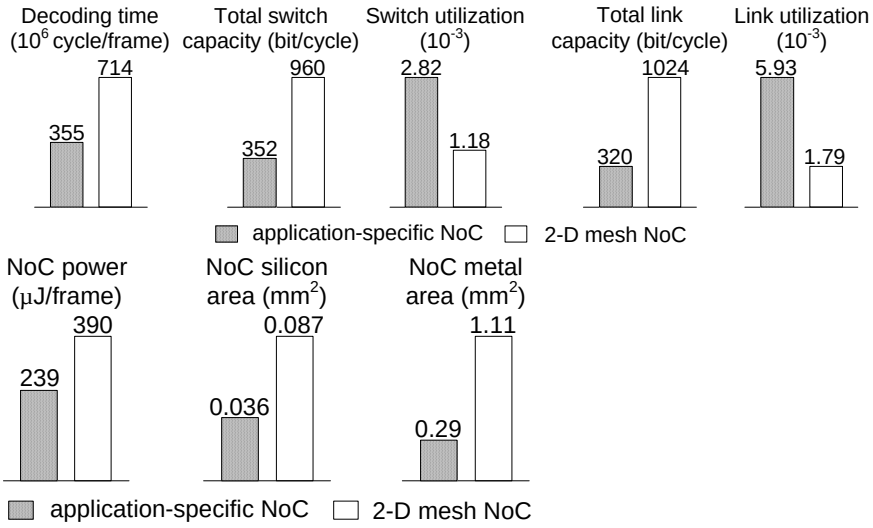
Decoding time (10^6 cycle/frame)
355
714
Total switch capacity (bit/cycle)
352
960
Switch utilization (10^{-3})
2.82
1.18
Total link capacity (bit/cycle)
320
1024
Link utilization (10^{-3})
5.93
1.79
application-specific NoC
2-D mesh NoC
NoC power (μJ/frame)
239
390
NoC silicon area (mm^2)
0.036
0.087
NoC metal area (mm^2)
0.29
1.11
application-specific NoC
2-D mesh NoC

B.1 Interconnection Technology: NanoCMOS, Optical, PICMOS

- **"Extended Global Interconnect Architecture for Nano-CMOS Technologies",**
  J.Balachandran[1], S.Brebels[1], G.Carchon[1], M.Kuijk[2], W.De Raedt[1], B.Nauwelaers[3] and E.Beyne[1],
  1 IMEC Belgium
  2 VUB Belgium
  3 KUL Belgium
- **"Novel Performance of Three-dimensional (3D) On-chip Crossbar Bus using non-Silicon Transistors"**
  S. Fujita[1/2], K. Abe[1], K. Nomura[1] and T. Lee[3],
  1 Frontier Research Laboratory[1], Toshiba Corporation, Komukai-Toshiba-cho, Saiwai-ku, Kawasaki, Japan
  2 Toshiba America Research, USA
  3 Stanford University, USA
- **"Reconfigurable Optical Networks for On-Chip Multiprocessors"**
  W. Heirman[1], J. Dambre[1], I. O'Connor[2], J. Van Campenhout[1]
  1 ELIS, Ghent University, Belgium;
  2 LEOM, Ecole Centrale de Lyon, France
- **"Optical interconnect for on-chip data communication"**
  I. O'Connor[1], F. Tissafi-Drissi[1], D. Navarro[1], F. Mieyeville[1], F. Gaffiot[1], J. Dambre[2], M. De Wilde[2], D. Stroobandt[2] and D. Van Thourhout[3]
  1 Ecole Centrale de Lyon, France
  2 Ghent University, Belgium
  3 IMEC/Ghent University, Belgium
- **"Heterogeneous integration of light sources on an SOI waveguide platform for photonic interconnects on CMOS (PICMOS)"**
  G. Roelkens[1], J. Van Campenhout[1], D. Van Thourhout[1], R. Baets[1], P. Rojo-Romeo[2], C. Seassal[2], P. Regreny[2], P. Viktorovitch[2], I. O'Connor[2], L. Di Cioccio[3] and J. Fedeli[3]
  1 Ghent University Belgium
  2 Ecole Centrale de Lyon France
  3 CEA-DRT/LETI Franc

Extended Global Interconnect Architecture for Nano-CMOS Technologies

J.Balachandran, S.Brebels, G.Carchon, M.Kuijk[1], W.De Raedt, B.Nauwelaers [2] and E.Beyne

Microwave and RF Systems Group, IMEC vzw, Kapeldreef 75, 3001, Leuven, Belgium

Phone: +32 16 288700 Fax: +32 16 281501 e-mail: bchandra@imec.be

[1] *Vrije Universiteit Brussel,ETRO, Pleinlaan 2,1050 Brussel, Belgium*

[2] *Katholieke Universiteit Leuven, ESAT, Kasteelpark Arenberg 10, 3001 Leuven, Belgium.*

Evolution of integrated circuits to nanometer technologies with integration densities on the order of billion transistors and GHz range frequencies makes communication instead of computation a bottleneck. This communication problem is due to long wires termed as global interconnects, connecting the functional blocks on the chip. While transistor performance improves with scaling, the performance of global interconnects deteoriates. As a result, communication bandwidth and latencies are severely constrained. Furthermore, power/energy per communication bit has also increased dramatically.

In addition, power distribution within the chip is becoming detrimental to the overall performance in nanometer regime. With Scaling, V_{dd} drops down while supply current (I_{dd}) constantly increases. Consequently, even a small voltage drop across power distribution network can be a significant fraction of supply voltage, hurting the performance.

Technological advances such as Cu/low-k and circuit techniques such as repeaters do not cope with performance needs of advanced nano-CMOS technologies. While emerging technologies such as 3D integration and optical interconnects have manufacturing and other technical limitations, reverse scaling is an effective approach, perhaps the only solution to simultaneously address signal transmission and power distribution issues. ITRS roadmap recommends reverse scaling in three hierarchical classes namely local, semi-global and global with increasing metal pitches to balance wiring density and performance needs. A more attractive alternative to further extend on-chip multi-level wire stacking is to exploit Wafer Level Package (WLP) redistribution layers.

The WLP is a processing technology where wiring layers of typical packaging dimensions are post processed directly on top of an active silicon wafer. Compared to on-chip technologies, the WLP process permits a relatively easy fabrication of wide, less resistive yet dense copper lines with thick low-k dielectrics. The WLP interconnects behave inherently as LC transmission lines with near speed of light propagation. Further more they don't necessitate repeaters as their delay is linear with length.

Fig1. shows the extended global interconnect architecture with WLP techniques. It consists of two Cu layers for signal transmission and a dedicated PG planes. The signaling layers are separated by a thick low-k BCB (k =2.7) dielectric while the PG planes are separated by a thin high-k dielectric Ta_2O_5 (k=25), forming integrated decoupling capacitor of value as high as 10nF/mm^2. The PG planes not only ease power distribution but also provide good return path thereby enabling high quality data transmission at high speeds in the signaling layers.

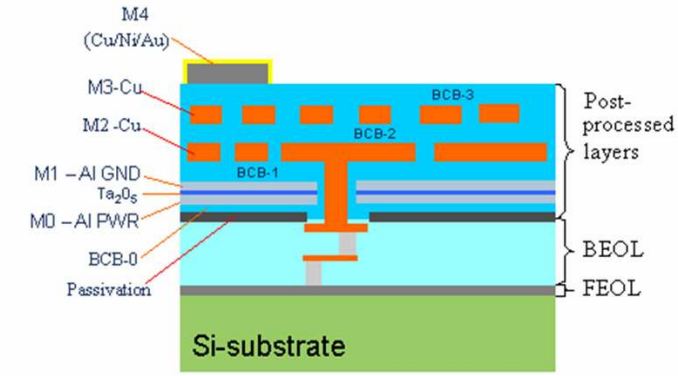

Fig1. Extended Global Interconnect architecture using WLP techniques

In this paper, we show how the proposed architecture can solve some of critical interconnect issues in the nanometer technologies. First, we compare the performance of the extended global interconnects with that of on-chip wires by considering relevant metrics such as bandwidth, latency and power, with the results derived from a test chip. We then highlight suitable signaling and wiring schemes for this extended architecture. With a proposed differential signaling scheme, we show that the extended global interconnects can support data rates as high as 20Gbps/ch for less than 1pJ per bit. We also discuss the applicability of this extended global interconnect architecture in L2 cache memory interconnections, clock distribution and in on-chip networks.

Novel Performance of Three-dimensional (3D) On-chip Crossbar Bus using non- Silicon Transistors

Shinobu Fujita[1,2], Keiko Abe[1], Kumiko Nomura[1], Thomas H. Lee[3]

[1] Frontier Research Laboratory, Toshiba Corporation, Komukai-Toshiba-cho, Saiwai-ku, Kawasaki, 212-8582 Japan. ;

[2]Toshiba America Research, Inc., 2590 Orchard Parkway, San Jose, CA 95131 USA *E-mail:* **shinobu.fujita@toshiba.co.jp**

[3]Stanford University Center for Integrated Systems, CIS-205 MC, 4070Stanford, CA 94305-4070 USA;

It is known that the performance of high-performance microprocessor is recently limited by the bottleneck of data transfer through bus circuits. This trend will become more marked for multi-core processor architecture. We have proposed that the bottleneck can be solved by increasing parallelism by adopting 3D bus based on 3D circuits using non-silicon devices. In our poster, we present the dramatic improvement of latency of bus by adopting 3D bus architectures, and future prospect to increase bandwidth using non-silicon transistors suitable for process integrations of 3D circuits.

3D circuit integration technology can be categorized into three generations in terms of the alignment technology of each circuit layer, as shown in Fig. 1. Density of vertical interconnection, a key parameter to estimate the capability of parallel data transfer or bandwidth of 3D circuits, increases with a shift from each generation to the next.

Latency of 2D bus is dominated by wire delay of interconnect. Considering wire delay of global interconnects with repeaters calculated for 65nm, 45nm or 32nm technology, it has been confirmed that the clock frequency for 2D bus will saturate at less than 10 GHz, even though system clock can continue to be increased. On the other hand, wire delay for 3D bus can be decreased drastically with increasing the number of 3D circuit layers, and reach the negligible value compared with gate delay, as shown in Fig.2.

We also calculated gate delay associated with 3D cross-bar bus (Fig. 3) based on bulk silicon 32nm CMOS technology and using SPICE simulation. Also, we estimated that of some non-silicon devices. In addition, order of maximum bandwidth was estimated in the case of each generation of 3D circuit integration technology. Figure 3 shows comparison of bandwidth and latency for each generation, where chip area is $10cm^2$. The 3D active device stacking technology has a potentiality to realize 3D on-chip crossbar bus with a bandwidth over *10P(10 Peta, 10^{16})bit/s* and a latency of order of 10 psec.

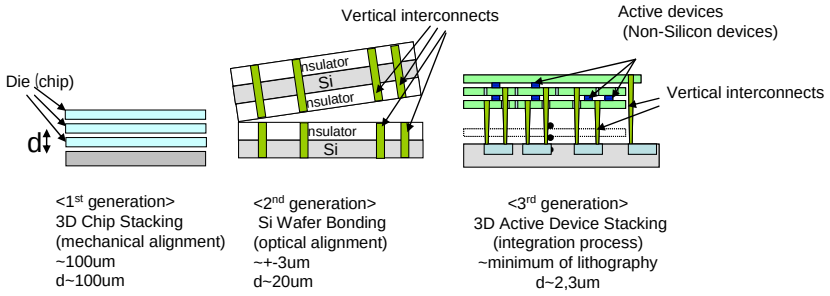

Fig.1 Generation of 3D circuit integration technology

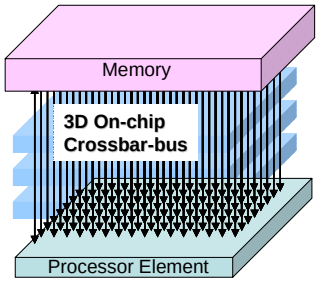

Fig.3 3D on-chip crossbar bus

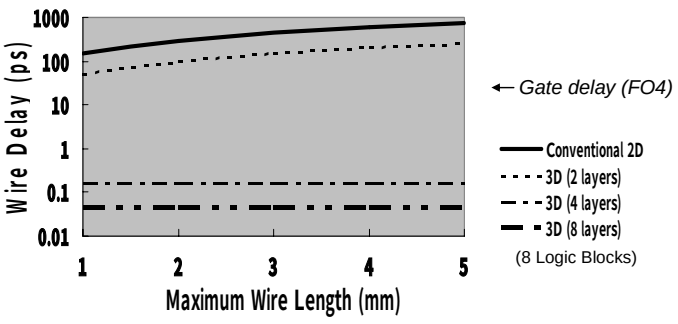

Fig.2 Wire delay of bus circuits (32nm technology)

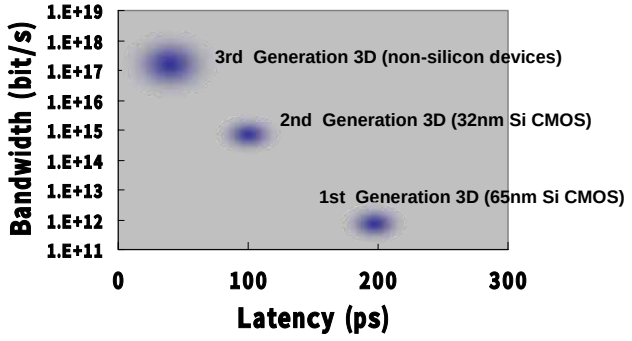

Fig.4 Maximum bandwidth and latency of 3D crossbar bus

Reconfigurable Optical Networks for On-Chip Multiprocessors

W. Heirman[a], J. Dambre[a], I. O'Connor[b], J. Van Campenhout[a]

[a]ELIS, Ghent University, Sint-Pietersnieuwstraat 41, B-9000 Gent, Belgium;
[b]LEOM, Ecole Centrale de Lyon, 36 Avenue Guy de Collongue, F-69134 Ecully, France
Presenting author: Wim Heirman, wim.heirman@elis.ugent.be

In previous work,[1,2] we looked at the feasibility of adding an optical, reconfigurable interconnection network to a server or supercomputer class shared memory multiprocessor machine. By allowing the interconnection network to adapt to the changing requirements made by the application executing on the multiprocessor machine, a significant speedup can be provided. Reconfiguration was to be implemented using tunable VCSEL light sources and a free-space optical selective broadcast device. Since the VCSEL tuning speed is limited to around 1 ms, we are only able to exploit low-frequency dynamics of the workload. Also, the reconfigurable network links have to be taken off line during reconfiguration. We therefore divided the network into a fixed base network (with regular topology), and a set of *extra links* that can be reconfigured at runtime. The extra links are placed between network nodes that communicate intensely, providing a short path (measured as the number of intermediate nodes where routing, arbitration and buffering is performed) for most of the traffic. It also redistributes the available network bandwidth to where it is most needed, minimizing contention and again reducing packet latency. Using full-system simulations of the processors and interconnection network, we predicted application speedups of up to 50%.

Following Moore's law, single chips already incorporate multiple processing cores (IBM Power4, Dual Opteron and Xeon, ...), and are scaling to dozens of cores in the near future. This allows one to place a large part of the 16-processor shared-memory machine that was studied in our previous work on a single chip, effectively transforming the interconnection network we studied to a Network-on-Chip. With the right technology, such as optical add-drop filters,[3] it is possible to implement reconfigurable on-chip networks,[4] providing the same general environment as the off-chip networks we studied before. We therefore propose to port our previous work to on-chip networks.

Of course, some assumptions made in the off-chip case are no longer valid when moving to on-chip networks. For instance, on-chip optical waveguides cannot easily cross each other. This restricts connectivity compared to a free-space implementation. We will therefore construct a network composed of 2 signaling layers, each providing waveguides in orthogonal directions, creating a mesh network. When compared to a regular packet-switched mesh network, reconfiguration will allow some packets to 'skip' certain intermediate hops. This avoids arbitration and buffering, again yielding a lower end-to-end packet latency.

REFERENCES

1. W. Heirman, I. Artundo, L. Desmet, J. Dambre, C. Debaes, H. Thienpont, and J. Van Campenhout. Speeding up multiprocessor machines with reconfigurable optical interconnects. In *Proc. SPIE Vol. 6124, Optoelectronic Integrated Circuits X, Photonics West*, San Jose, California, January 2006.
2. W. Heirman, J. Dambre, D. Stroobandt, J. Van Campenhout, C. Debaes, and H. Thienpont. Prediction model for evaluation of reconfigurable interconnects in distributed shared-memory systems. In *Proceedings of SLIP'05*, pages 51–58, San Francisco, California, April 2005. ACM Press.
3. A. Kazmierczak, M. Briere, E. Drouard, P. Bontoux, P. Rojo-Romeo, I. O'Connor, X. Letartre, F. Gaffiot, R. Orobtchouk, and T. Benyattou. Design, simulation, and characterization of a passive optical add-drop filter in silicon-on-insulator technology. *IEEE Photonics Technology Letters*, 17:1447–1449, July 2005.
4. I. O'Connor et al. Towards reconfigurable optical networks on chip. In *Proceedings of ReCoSoc'05*, Montpellier, France, June 2005.

Optical interconnect for on-chip data communication

I. O'Connor[1], F. Tissafi-Drissi[1], D. Navarro[1], F. Mieyeville[1], F. Gaffiot[1],

J. Dambre[2], M. De Wilde[2], D. Stroobandt[2], D. Van Thourhout[3]

Contact: ian.oconnor@ec-lyon.fr

[1]LEOM
Ecole Centrale de Lyon
36 avenue Guy de Collongue
F-69134 Ecully, France

[2]Electronics and Information Systems
Ghent University
St. Pietersnieuwstraat 41
B-9000 Ghent, Belgium

[3]Photonics Research Group
IMEC/Ghent University
St. Pietersnieuwstraat 41
B-9000 Ghent, Belgium

Abstract

It is believed that the concept of integrated optical interconnect is a potential technological solution to alleviate some of the ever more pressing issues involved in exchanging data between cores in SoC architectures (inter-line crosstalk, latency, global throughput, connectivity and power consumption). This abstract summarises work carried out in the framework of the EU-funded PICMOS project on the quantitative comparison of optical interconnect to electrical interconnect in the context of on-chip data communication.

The investigation strategy (evaluation approach, performance metrics) to quantitatively analyse the impact of specifications on both optical and electrical interconnect performance aims at accurately comparing link performance. To enable complete link simulation with transistor-level interface circuits, Verilog-A behavioural models were developed for optical sources, detectors and waveguides. To automatically design both interconnect types following this simulation-based approach, two optimisation toolsets were used (RuneII toolset by LEOM for optical links and ELSA toolset by IMEC/PARIS for electrical links).

In the RuneII toolset, analog design automation techniques were exploited to develop hierarchical synthesis methods for the links and for transistor-level transimpedance amplifier and current-mode source driver interface circuits. The synthesis method is capable of operating over the technology nodes of interest and with varying link-level specification sets, summarised in Table 1. The optimisation strategy that has been chosen minimises the power consumption for a given data rate.

Specification	Value
CMOS technology	{65; 45; 32} nm
BER	1.00E-18 s^{-1}
ITRS max frequency	{2.98; 5.2; 11} GHz
Length	{2500um,20000um}
Activity rate	1
Ambient temperature	70 °C
V_{dd} (CMOS)	{1.2; 1.1; 1.0} V

Table 1 Specifications for optical link synthesis

The performance metrics taken into account were *static and dynamic power*, *link delay* (measured from electrical input to electrical output), *CMOS area* and *communication density* for optical and electrical links with equal pitch and equal single-link data rate. Electrical links were evaluated in a parallel-wire configuration, including capacitive coupling and crosstalk effects in the power and delay evaluations. The comparison was performed for a number of technologies: predictive technology models for several technology generations (gate lengths down to 32nm) and commercial technologies from STMicroelectronics. For each technology, link simulations were performed for link lengths from 2.5mm up to 2 cm. Figure 1 shows some typical results (here concerning total power comparison).

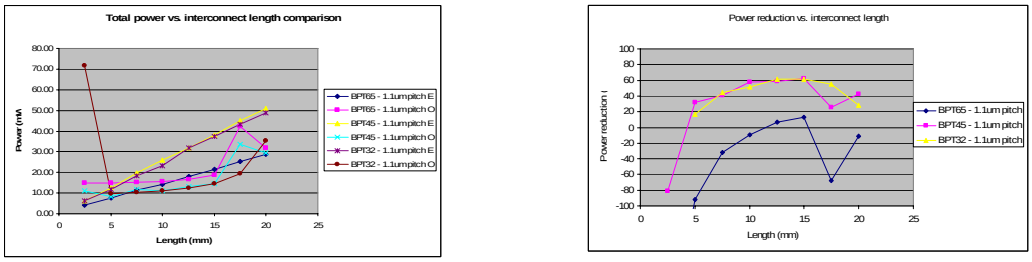

Figure 1 Total power vs interconnect length comparison for 65-45-32nm gate length technologies

Heterogeneous integration of light sources on an SOI waveguide platform for photonic interconnects on CMOS (PICMOS)

Günther Roelkens, Joris Van Campenhout, Dries Van Thourhout, Roel Baets
Ghent University – IMEC, Sint-Pietersnieuwstraat 41, B-9000 Ghent – Belgium
e-mail: gunther.roelkens@intec.ugent.be

Pedro Rojo-Romeo, Christian Seassal, Phillipe Regreny, Pierre Viktorovitch, Ian O'Connor
Laboratoire d'Electronique, Microelectronique et Micro-systèmes, Ecole Centrale de Lyon, 36 Avenue Guy de Collongue, 69134 Ecully cedex-France

Lea Di Cioccio, Jean-Marc Fedeli
CEA-DRT/LETI, 17 Rue des Martyrs, 38054 Grenoble cedex 9-France

In future generations of electronic circuits a severe bottleneck is expected on the global interconnection level. With decreasing device dimensions, it is increasingly difficult to keep propagation delays acceptable and even with the most optimistic estimates for conductor resistivity and dielectric permittivity, the projected performance roadmap will not be met. Therefore there is a need for radically different interconnect approaches and, as indicated by the ITRS-roadmap, one of the most promising solutions is the use of an optical interconnect layer. An optical interconnect layer could allow for an enormous bandwidth increase, for immunity to electromagnetic noise, for a decrease in the power consumption, the possibility for synchronic operation within the circuit and with other circuits and for a reduced immunity to temperature changes.

To achieve an optical interconnection layer on top of a CMOS electronic circuit we propose to bond a Silicon-On-Insulator (SOI) optical waveguide layer on top of the electronic circuit. In recent years, SOI has emerged as a promising platform for high density passive integrated optics fabricated on a wafer scale. For active opto-electronic components like laser diodes, modulators and detectors, III-V material remains the workhorse of telecom industry due to its superior performance. Therefore, there is also the need to integrate III-V active components on top of the CMOS circuits.

We focus on the integration of InP/InGaAsP light sources on an SOI optical waveguide platform. Depending on the application, different requirements for the sources are needed. When a signal needs to be broadcasted over the chip a high power laser diode is needed, which may require a large footprint. For a point-to-point datalink less optical power is required, while the footprint has to be smaller to allow a dense integration.

For the case of a high power laser diode we focus on integrating a standard Fabry-Perot type laser diode on top of an SOI waveguide circuit as shown in figure 1. An efficient coupling scheme was designed to couple light from the laser diode into the SOI waveguide. A threshold current of 10mA is envisioned (footprint 1000um^2). For the case of a low power laser diode, an evanescently coupled microdisk laser was designed and fabricated as shown in figure 2. Threshold current of a stand alone laser diode was 1.5mA for a microdisk of 8um wide. Ultimately, threshold currents down to 10μA are expected. Fabrication details will be given in the poster presentation.

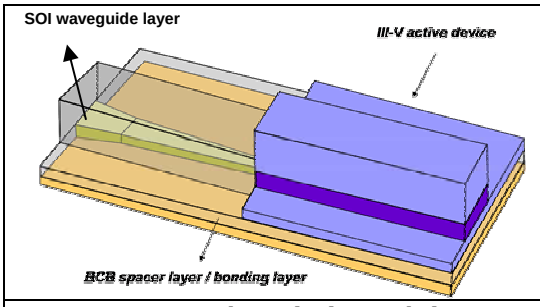

Figure 1: FP laser diode coupled to SOI

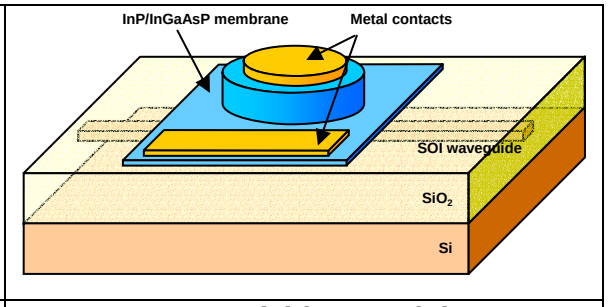

Figure 2:Microdisk laser coupled to SOI

B.2 Signalling and Components: Transmission Modes, Protocols, Encoding, Timing and Synchronisation Issues, Components for NoC

- **"A 20Gb/s Transceiver for Network-On-Chip"**
  Z. Asgar, J. Zou, P. Jain, R. Kamath and R. Harjani
  University of Minnesota, USA
- **"Buffer output capacitance effect on NoC line drivers performance"**
  G. Cappuccino, A.Pugliese, G. Cocorullo
  DEIS-University of Calabria, Italy
- **"Multiple-rail phase-encoding for NoC"**
  C. D'Alessandro[1], D. Shang[1], A. Bystrov[1], A. Yakovlev[1], O. Maevsky [2]
  1 University of Newcastle upon Tyne, UK
  2 Intel Labs, Moscow
- **"High-Speed Serial Interconnect for NoC "**
  Rostislav (Reuven) Dobkin, Ran Ginosar and Avinoam Kolodny
  Electrical Engineering Department, Technion, Haifa, Israel
- **"Dynamic Multi-grain Pipelined Interconnect"**
  M. Imai, T. Azuma, K. Watanabe, M. Kondo, H. Nakamura and T. Nanya
  Research Center for Advanced Science and Technology, The University of Tokyo
- **"Clock Synchronization in NoCs using Distributed FIFOs"**
  P. Pande and J. Nyathi
  Washington State University, USA
- **"A Distributed, Interleaved FIFO for SoC Interconnect"**
  S. Sood, M. Greenstreet and R. Saleh
  University of British Columbia
- **"Design of fast and reliable synchronisers for NoCs "**
  J. Zhou, D. Kinniment, G. Russell, and A. Yakovlev
  University of Newcastle upon Tyne, UK

A 20Gb/s Transceiver for Network on Chip

Zain Asgar*, Jia Zou, Priyesh Jain, Raghavendra Kamath and Ramesh Harjani
Department of Electrical and Computer Engineering,
University of Minnesota, 200 Union Street SE, Minneapolis, MN 55455
{asga0002*,zoux0024, jainx042, kama0076, harjani}@ece.umn.edu

Abstract

With the increasing integration of ICs, the International Technology Roadmap for Semiconductors (ITRS) expects system-on-chips (SOCs) to reach 10 GHz and 4 billion transistors by the end of the decade. System level integration is prohibitively difficult using conventional bus based approaches as skew and wire delays dominate. To overcome this issue we present a fully-asynchronous variable speed (up to 20 Gb/s) per channel transceiver using the 0.13um UMC technology for networks on chip. The total power consumption for the transceiver including amplifiers is 50 mW. The transceiver, shown in Figure 1, is designed to send and receive data through an on-chip transmission line over 1cm in length, without repeaters. The transceiver requires accurate multi-phase clocks, which are generated using a multiplying delay locked loop (MDLL), to serialize/deserialize the data. The incoming data is de-serialized using multi-phase parallel sampling and semi-digital data recovery. A single clock generator can be utilized for both the transmitter and receiver as the clock phases are never altered by the incoming data, as is typically done in a traditional analog deserializer, thereby allowing us to reduce power. Additionally, as the clocks are shared across channels, multi-channel communications can be implemented with a low overhead. The receive amplifier uses an active feedback, active inductor circuit to detect the low input voltage levels due to channel attenuation and to drive the large output capacitance presented by the 8 samplers. The samplers are built using a clocked sample-and-hold with regenerative amplification to sample the 20 Gb/s incoming data. The data is then retimed by interpolating the clock edge in a closed loop form. The samplers are used as phase detectors and a hybrid analog/digital filter scheme is used for the loop filter. The digital filter is used for coarse adjustments of the interpolators while the analog scheme independently fine tunes each interpolator. The analog fine tuning performs the accurate phase alignment needed for the samplers and the interpolators. The data from the samplers is then word-wise realigned and decoded before exiting the block. We present the overall architecture and simulation data for the NOC transceiver including the particulars for the transmission line model used to guide the design of the front-end transmit and receive amplifiers.

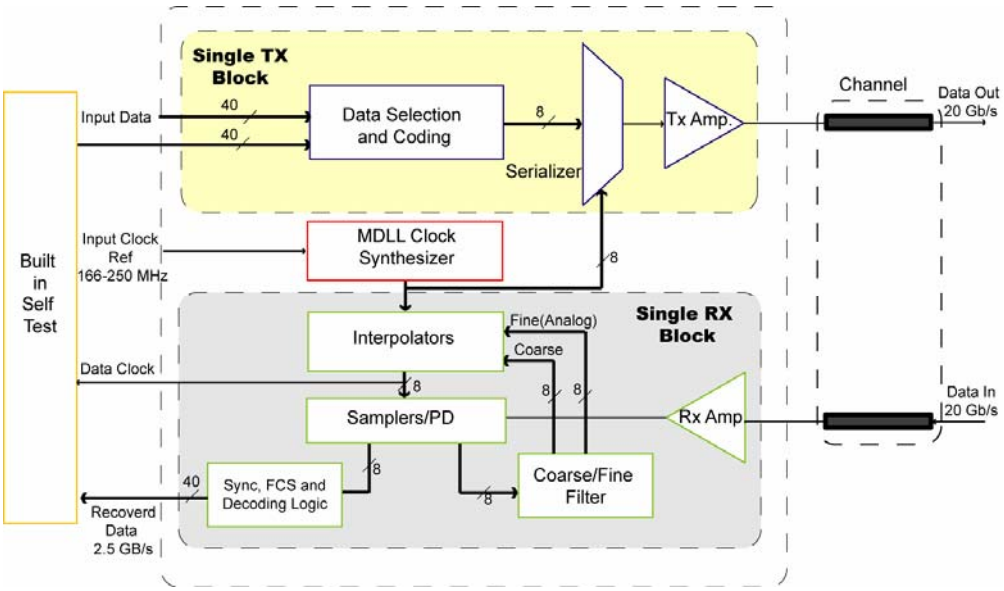

Figure 1: Transceiver Model

* Denotes Contact Author

Buffer output capacitance effect on NoC line drivers performance

Presenter: Gregorio Cappuccino
Electronics, Computer Science and Systems Department,
DEIS-University of Calabria
87036 Rende, Italy
cappuccino@deis.unical.it

Other Authors: Andrea Pugliese, , Giuseppe Cocorullo

Abstract:
Network-on-Chip (NoC) architectures have emerged as a hopeful alternative to overcome the problems associated with long on-chip buses. NoCs are based on functional blocks linked by a router network, these latter interconnected by long wires.
Costs, speed, power and Area performance are the principal concerns that NoC designers are learning to cope with.
Although the concept of NoC is relatively new, it has already attracted significant attention in the research community because of its potential in solving the on-chip communication problem. Most of solutions proposed in literature are based on previous studies and analysis done for on-chip interconnects. One of the well exploited techniques to optimise performances borrowed from the past is the repeater insertion. However the particular area and power requirements imposed by complex circuits dictate the need for a more accurate design of repeaters and in some cases, it implies to take in to account important phenomena previously considered as" secondary" or unimportant. Thus optimisation techniques require effective models to enable a true optimisation of overall performances.
Recent studies point out the role of these "secondary" phenomena, as threshold variations, current leakage [1, 2] or device channel length modulation [3], all phenomena that may affect significantly the operations of modern nanometer CMOS . However to the best of our knowledge, actual role played by buffer output capacitance on the buffer performance remains quite unexplored. The poster try to present part of the work done and some results carried out about the actual effect of drain MOSFET capacitance on the transient behaviour of the line drivers. The effect on the actual effectiveness of some well known sizing criterions is also presented.

References:

[1] A. Morgenshtein, I.Cidon, A. Kolodny and R. Ginosar, "Low-Leakage Repeaters for NoC Interconnects," ISCAS 2005.

[2] F. Fallah, S. Nassif, M. Pedram, "Technology, Circuit and Design Automation Techniques for Leakage Minimization of CMOS VLSI Circuits," Design Automation and Test in Europe (DATE), Mar 2005 (Master Course).

[3] Ashok Narasimhan, Shantanu Divekar, Praveen Elakkumanan, Ramalingam Sridhar "A Low-Power Current-Mode Clock Distribution Scheme for Multi-GHz NoC-Based SoCs" Proc. of the 18th International Conference on VLSI Design held jointly with 4th International Conference on Embedded Systems Design (VLSID'05);

Multiple-rail phase-encoding for NoC

Crescenzo D'Alessandro[1], Delong Shang[1], Alex Bystrov[1],
Alex Yakovlev[1], Oleg Maevsky[2]
[1]University of Newcastle upon Tyne, UK
[2]Intel Labs, Moscow

Abstract

The successful use of Systems-on-Chip (SoC) and in particular of Networks-on-Chip (NoC) requires the deployment of a reliable on-chip communication fabric. Design of such interconnections often revolves around the use of serial links, which improve reliability being less sensitive to skew between the lines; these links also employ the clocked paradigm in order to simply design. In this context performance improvement can only be achieved by the use of faster clocks. In order to maximise performance the use of parallel link architecture can be explored: however, the ever-increasing intra-chip variability predicted by the ITRS [3] makes it difficult if not impossible to satisfactorily design long high-speed on-chip interconnect, due to the resulting skew between the lines. Clocked links (serial or otherwise) can suffer from the effects of Single-Event Upsets (SEUs), caused primarily by radiations; these effects consist in one of the lines being hit by a charged particle which generates a pulse on the line. Cross-talk effects can also degrade the performance of the links in similar ways. Taking into account the predictions expressed by Dupont in [2], next generation chips will become more sensitive to such effects even at sea level. Traditional asynchronous techniques help to mitigate these problems by introducing redundancy, necessary to indicate the presence of valid data at a stage of a data path. Dual-rail and more generally $1 - of - m$ techniques [1] can be employed to eliminate the effects of variability, as they ensure correct communication independently from a centralised clock; the latter also provides a higher data-rate compared to the dual-rail protocol. However, these techniques still suffer from radiation effects and cross-talk. A different approach is proposed which allows higher data-rates and higher resistance to SEUs: the data is sent encoded into the phase relationship between n wires. When a symbol is transmitted all the lines toggle monotonically and the receiver recovers the data by observing the order of events on the lines. This approach resembles the Rank Order Coding proposed by Thorpe [4]. The receiver identifies and in some cases filters out transient faults by analysing the behaviour of the lines. Variability is accounted for by synchronising the lines; this task is simplified thanks to the monotonicity of the symbol transmission. A novel encoding algorithm to obtain a sequence-based symbol from a one-hot encoded symbol is also proposed together with theoretical results for the link. Thanks to the use of n wires each symbol encodes $n!$ states and therefore $\lfloor \log_2(n!) \rfloor$ bits can be sent in each symbol, resulting in significant increase of data-rate. The higher bit-rate per symbol also results in less transitions per packet, eventually resulting in lower power consumption.

[1] W.J. Bainbridge. *Asynhcronous System-on-Chip Interconnect.* PhD thesis, University of Manchester, UK, March 2000.

[2] E. Dupont, M. Nicolaidis, and P. Rohr. Embedded robustness IPs. In *Proceedings of the 2002 Design, Automation and Test in Europe Conference and Exhibition (DATE'02)*, 2002.

[3] International Technology Roadmap for Semiconductors (ITRS-2003), 2003.

[4] SJ Thorpe and J Gautrais. Rank Order Coding. In J Bower, editor, *Computational Neuroscience: Trends in Research*, pages 113–118, New York, 1998. Plenum Press.

High-Speed Serial Interconnect for NoC

Rostislav (Reuven) Dobkin, Ran Ginosar, Avinoam Kolodny
EE Department, Technion - Israel Institute of Technology

Delays over global wires in large systems on chip (SoC) do not scale with technology. This problem challenges on-chip data communications in terms of latency, throughput and power.

Network on Chip (NoC) is advocated as a solution for the SoC interconnect problem. In NoC, numerous multiple wires, required for on-chip bit-parallel interconnect, provide high data rates at the cost of large chip area, routing difficulty, and high power. In addition, such links are often utilized only a small portion of the time, but dissipate leakage power at all times. Leakage is incurred at the line drivers and also at the repeaters, which are often necessary for long interconnects.

Bit-serial communication offers an alternative to bit-parallel interconnects, mitigating the issues of area, routability, and leakage power, since there are fewer wires, fewer line drivers, and fewer repeaters. However, in order to provide the same throughput as an N-bit parallel interconnect, the serial link must operate N times faster.

Conventional fast differential-signaling serial link circuits, typically employed for chip-to-chip I/O communications, perform data sampling at the receiver side and are inappropriate for on-chip serial links, since they require complex clock and data recovery PLL-based circuits, which consume excessive power and area. We investigate a low power asynchronous data transfer technique, based on differential transition signaling. The circuits present a lower bound on the bit time of one gate delay.

We propose a novel fast bit-serial communication scheme (Figure 1) that employs low-latency synchronizers at the source and sink, two-phase NRZ Level Encoded Dual Rail (LEDR) asynchronous protocol (allowing non-uniform delay intervals between successive bits), serializer and deserializer (SERDES), line drivers and receivers, and differential channel encoding. Acknowledgment is returned only once per word, rather than bit by bit, enabling multiple bits in a wave-pipelined manner over the serial channel. The wires should be designed as wave-guides, enabling multiple traveling signals. At signal propagation velocity of at least c/10 on a well-designed wave-guide, and at a desired data rate of one bit per 15ps (the expected FO4 inverter delay at 65nm), a 1mm wire may carry at least two successive bits simultaneously.

The targeted high rate requirement of one gate-delay data cycle poses stringent requirements in all components of the communication scheme, calling for investigation and construction of novel high-speed circuits, even for simple logical functions. In this work we present a novel high-rate shift-register architecture that supports the required rate. The shift register is used for serializer and de-serializer construction.

The shift-register was found to be robust in terms of in-die variations. It has been simulated successfully on 65nm process with up to 10σ process variations, and over 24 PVT corners. The fast SERDES, made possible by this shift-register, is useful for high-bandwidth long NoC interconnects, where bit serial communication is preferred thanks to reduced area, easier routing and reduced leakage.

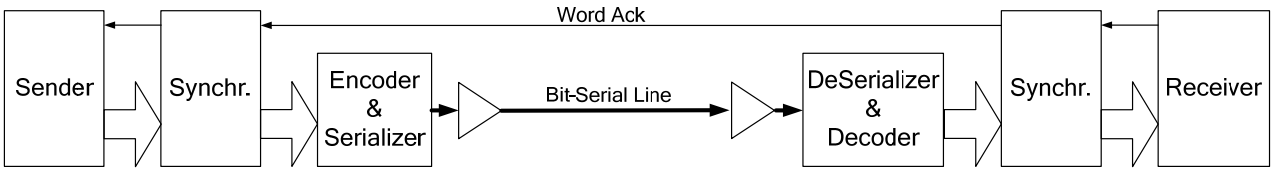

Figure 1: Fast Bit-Serial Communication Scheme

Dynamic Multi-grain Pipelined Interconnect

Masashi Imai Takuya Azuma Kouichi Watanabe Masaaki Kondo Hiroshi Nakamura Takashi Nanya
Research Center for Advanced Science and Technology, The University of Tokyo
{miyabi, takuya, kouichi, kondo, nakamura, nanya}@hal.rcast.u-tokyo.ac.jp

As the semiconductor technology advances, system-on-chips which contain many processing units, memory units, and other peripheral units have been implemented. In such SoCs, the activity of each unit and the rate of data flow between units may be different in accordance with applications and processing data. The performance requirements of data transfer between units are changed dynamically. However, in traditional bus-based architecture, all the units which share the same bus have the common throughput for the data transfer, resulting in redundant energy consumption. Thus, if the most appropriate throughput can be selected dynamically for each data transfer, the energy consumption can be reduced. We present an asymmetric c-element based pipeline latch (Figure 1) that can be used as both a latch-gate and a repeater-gate. For the throughput-oriented applications, $s = 0$ in Figure 1, this latch works as a c-element based pipeline latch. On the contrary, when it is not required to work as high-throughput pipeline, $s = 1$ in Figure 1, it works as a buffer gate. They can change dynamically in accordance with the requirement of processing applications. In addition, when all the latches work as repeater gates, we can use synchronous latches in both sender units and receiver units.

As the first discussion of this work, in this poster, we show some evaluation results of throughput, latency, and energy consumption of interconnects in 90nm process technologies. We evaluate the following circuit structures using HSPICE simulator; buffer inserted interconnect lines using normal BUF gates and the proposed asymmetric C-element based buffers, pipelined interconnect lines using synchronous latches, asynchronous C-element based latches, and the proposed latches. We use the W-element model (Figure 2) in HSPICE model libraries which can handle frequency-independent (RLGC) and lossless (LC) lines as interconnect lines. Figure 3 shows throughput, delay, and power consumption of 15mm interconnect lines, respectively. From Figure 3, it can be seen that the delays of asynchronous interconnects are smaller than those of synchronous interconnects since delay variations due to process variations and the crosstalk can be reduced. It can be also seen that the performance overheads of the proposed circuits are small.

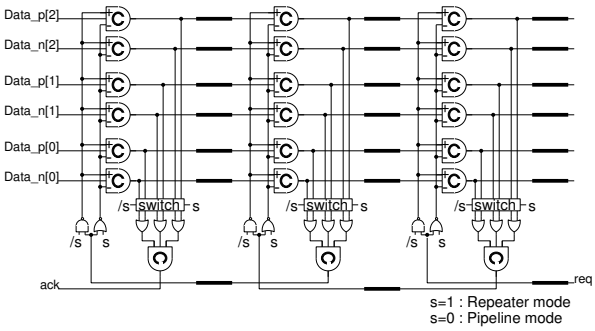

Figure 1. Asymmetric C-element based Bi-mode interconnect.

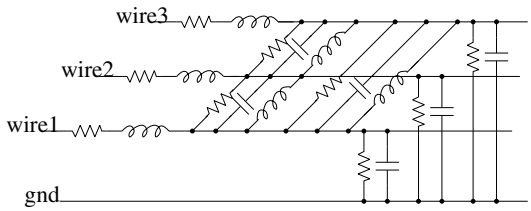

Figure 2. W-element model of interconnect.

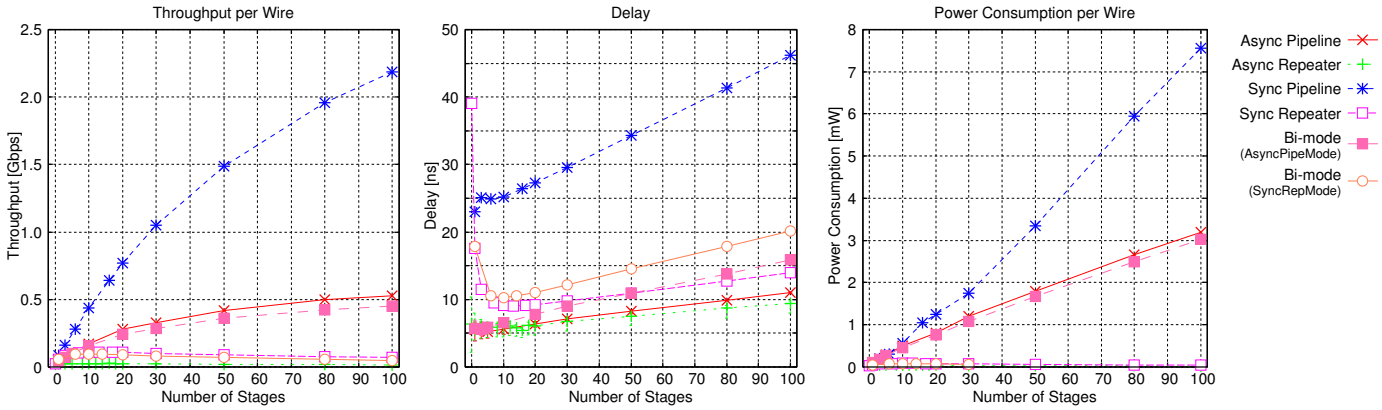

Figure 3. Throughput, delay, and power consumption of interconnect.

Clock Synchronization in NoCs using Distributed FIFOs

Partha Pande and Jabulani Nyathi
Presenter: Partha Pande
School of Electrical Engineering & Computer Science
Washington State University PO Box 642752
Pullman, WA 99164-2752, USA
(pande, jabu)@eecs.wsu.edu

The network-on-Chip (NoC) design paradigm is viewed as an enabling solution for the integration of exceedingly high number of computational and storage blocks in a single chip. Though NoC research has gained significant momentum, the aspect of clock distribution for NoCs has not received noticeable attention. Researchers predominantly suggested utilization of existing clock synchronization methodologies in NoC domain, where they assume to distribute a single synchronous clock with differing phases all along the chip. It is well-known that it can take several clock cycles for a global signal to travel from one end of a chip to the other. Consequently synchronization of future chips with a single clock source and negligible skew will be extremely difficult, if not impossible. As a result of the cross-chip signaling constraints instead of trying to distribute synchronous clock, the whole SoC needs to be divided into multiple functional islands with independent clocks. This in turn leads to the synchronization issues, which has been addressed by a number of research groups in different contexts. The inherent characteristics of NoC architectures allow the whole SoC to be divided into multiple clusters according to the interconnect infrastructure as shown in Fig 1 (a). As these clusters have different frequencies, communication among them gives rise to new challenges in terms of data integrity and energy dissipation. One of the principal characteristics of the NoC architectures is that the functional blocks communicate with one another with the help of intelligent switches. Switches have storage buffers either at the input or output and we propose to re-use these buffers to manage multiple clock domain synchronization.

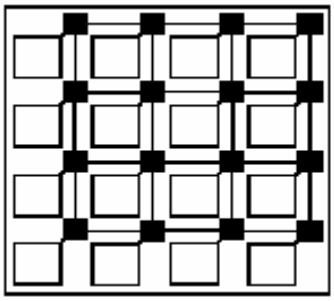

Fig. 1 (a): MESH-based NoC

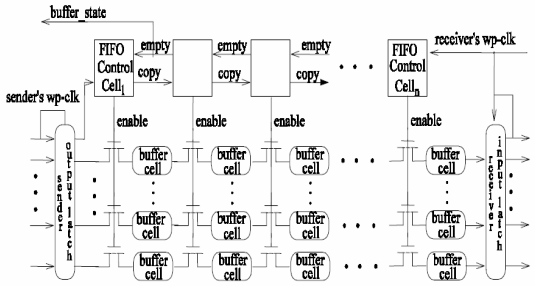

Fig. 1 (b) Distributed FIFO scheme

It is expected that each message injected to the NoC will have a clock signal associated with it. This clock in conjunction with a buffer status signal which has information about data rate transfer will be used to generate a local clock that potentially is at the rate of the receiving IP's clock rate. The basic concept of the scheme is that the status of the receiving module's storage elements is always known and can thus be used to synchronize with the sender. The information about the channel status is easy to manipulate in the event that the sender is slower than the receiver, but introduces a significant challenge when the receiver is slower than the sender. The latter case has traditionally been addressed by allowing the sender to stretch its clock i.e. by using flip-flops to slow the clock frequency of the sender. This leads to a potential problem of metastability. Our approach allows for the generation of a local clock that drives the channel and informs the sender of the channel capacity. We distribute the switch buffers along the inter-switch wire segments as shown in Fig. 1 (b). This will reduce the power dissipation arising out of the exchange of control signals among the communicating modules. Power dissipation is a significant concern and NoC design paradigm addresses this concern as evidenced by energy aware mapping of NoC architectures. Our scheme exploits these benefits by using the inherent NoC buffers to address single cycle and multiple cycle inter-module communication of IPs operating at different frequencies.

A Distributed, Interleaved FIFO for SoC Interconnect

Santosh Sood, Mark R. Greenstreet and Resve Saleh *
University of British Columbia
santoshs@cs.ubc.ca, mrg@cs.ubc.ca, *res@ece.ubc.ca*

January 15, 2006

Abstract

We present a distributed, interleaved FIFO for high performance, on-chip communication. Our approach builds on Stevens's [Ste03] comparison of interconnect techniques. We note that synchronous, source-synchronous, and asynchronous links all are constructed of channels where long buses are broken into shorter segments using buffers and latches or flip-flops. Each approach brings its own limitations. Synchronous communication requires distributing clock signals throughout the interconnect; this can add substantial overhead, especially in designs where interfaces can operate at multiple different frequencies. Furthermore, it is difficult to exploit wave-pipelining in synchronous interconnect because changes in the clock frequency can violate the resulting two-sided timing constraints. For applications where low latency is required, source-synchronous methods suffer from the extra overhead of the skew-compensating FIFO at the receiver. Finally, existing asynchronous techniques are limited by the round-trip required for each handshake cycle: wire delays are incurred twice, and wave-pipelining is precluded.

Our new design starts with the twin-control, asynchronous interconnect proposed by Ho *et al.* [HGD04]. Their design requires a latch at each repeater stage and provides no timing slack when used in a source-synchronous context. Our design overcomes these limitations by increasing the degree of interleaving of the FIFO. We observe that the use of two control paths can be generalized to designs with three or more handshaking paths. Furthermore, we replace the GasP [SF01] shared control wire with separate wires for request and acknowledge. The resulting design is our interleaved, distributed FIFO.

We compare our FIFO with time borrowing two-phase and flip-flop based synchronous design as well as with Ho *el al.*'s original design. We use velocity (distance/time), and energy-delay (energy $*$ time$/$distance^2) metrics to compare the approaches. For practical data rates, the two-phase, time-borrowing method achieves the best performance if it is operated at a single frequency. The distributed, interleaved approach outperforms flip-flop based synchronous communication and Ho *el al.*'s original design. Furthermore, the distributed, interleaved approach maintains its performance over a wide range of throughput, and does not require global clock distribution. These qualities make our design well-suited for solving SoC interconnect challenges.

References

[HGD04] Ron Ho, Jonathan Gainsley, and Robert Drost. Long wires and asynchronous control. In *Proceedings of the Tenth International Symposium on Asynchronous Circuits and Systems*, pages 240–249, April 2004.

[SF01] Ivan Sutherland and Scott Fairbanks. GasP: A minimal FIFO control. In *Proceedings of the Seventh International Symposium on Asynchronous Circuits and Systems*, pages 46–53, April 2001.

[Ste03] Kenneth S. Stevens. Energy and performance models for clocked and asynchronous communication. In *Proceedings of the Ninth International Symposium on Asynchronous Circuits and Systems*, pages 56–66, May 2003.

*This work has been funded in part by a Commonwealth Graduate Fellowship and grants from the BC Innovation Council, Intel, NSERC, and SUN Microsystems.

Design of fast and reliable synchronisers for NoCs

Jun Zhou, David Kinniment, Gordon Russell, and Alex Yakovlev
Newcastle University, UK
{Jun.Zhou,David.Kinniment,G.Russell,Alex.Yakovlev}@ncl.ac.uk

Abstract

As the size of systems on chip (SoC) have increased, it has become difficult or impossible to accurately distribute a single global clock across the entire system [1]. Future systems are likely therefore to consist of many independently, or semi-independently clocked IP blocks, connected by a network that allows data to pass from one block to another with minimal latency, [2]. Each pair of blocks must synchronize the data passing between them, and consequently there will be many more synchronizers whose performance is crucial to the performance of the system itself. An important effect of scaling is the increase in both dynamic and static power dissipation. Currently proposed solutions to this problem include dynamic lowering the voltage in selected blocks when high performance is not required. Unfortunately, reduced power supplies usually disproportionately affect the performance of synchronizers since the synchronizer τ depends on the small signal parameters in metastability rather than large signal switching times, and long latencies can result. Typically jamb latches are used for synchronization, and a 50% reduction in power supply voltage may result in over 100% increase in τ. In line with many other synchronizer circuits the jamb latch has a metastable level that can cause both p and n-type transistors to have low transconductance, particularly at low temperatures where Vth is high. In this poster we describe a new latch circuit designed to give a high performance in low voltage synchronizer applications. By increasing the latch current, and hence the device transconductances only during metastability, we can more than maintain the value of the metastability time constant, τ without significantly increasing the power. At the same time our circuit reduces the variation of τ with Vdd and temperature, so that it has a lower τ at 0.9V (50% nominal Vdd) than the conventional jamb latch has at 1.1V (Over 60% nominal Vdd). The resulting circuit is not only more robust to Vdd variation, it also gives lower failure rates at normal Vdd in applications where long synchronization times are required.

[1] N.A Kurd, J.S. Barkatullah, R.O.Dizon, T.D.Fletcher, and P.D.Madland “Multi-GHz Clocking Schemes for Intel Pentium 4 Microprocessors” Proc. ISSCC 2001 Feb 2001 pp404-405.

[2] Tobias Bjerregaard, Jens Sparsoe, "A Scheduling Discipline for Latency and Bandwidth Guarantees in Asynchronous Network-on-Chip" Proceedings, ASYNC 2005 conference, New York, March 2005

B.3 NoC Infrastructures: NoC Testing, Fault-tolerance, Power management, Testing

- **"An improved method for delay fault testing of NoC interconnections"**
  T. Bengtsson[1], S. Kumar[1], A. Jutman[2] and R. Ubar[2]
  1 Jönköping University, Sweden
  2 Tallinn University of Technology, Estonia
- **"A dedicated communication layer for Power Management of NoC based SoCs"**
  N. Genko[1], D. Atienza[1/2], A. Aquaviva[1] and G. De Micheli[1]
  1 EPFL, Switzerland
  2 DACYA/UCM, Spain.
- **"BIST for NoC Interconnects"**
  C. Grecu[1], P. Pande[2], A. Ivanov[1] and R. Saleh[1]
  1 University of British Columbia, Canada;
  2 Washington State University, USA
- **"Fault-Diagnosis-And-Repair System for Improving the Fault-Tolerance and Manufacturability of MPSoCs"**
  H. Kariniemi and J. Nurmi
  University of Technology, Finland
- **"Versatile XGFT Network-On-Chip with Improved Fault-Tolerance for Multi-Processor Systems-on-Chip "**
  H, Kariniemi and J. Nurmi
  University of Technology, Finland
- **"Utilizing NoC Switches as BIST-structures in 2D Mesh Network-on-Chip"**
  K. Petersén and J. Öberg
  Royal Institute of Technology (KTH), Sweden

An improved method for delay fault testing of NoC interconnections

T. Bengtsson[1], S. Kumar[1], A. Jutman[2], R. Ubar[2]

[1]Jönköping University, Sweden
[2]Tallinn University of Technology, Estonia

Testing of high density SoCs operating at high clock speeds for delay faults is an important but difficult problem. Delay faults are more likely in long wires constituting buses connecting different clock domains in SoCs built using GALS principles. In [1], we presented a method for detection of delay faults in a bus connecting two NoC switches. In this paper, we describe an improvement of this method to make it more efficient.

In NoC systems built using GALS principles, switches communicate using handshaking based asynchronous communication. We assume that a handshaking protocol with two control signals, *Write* and *Ready To Receive*, is used as shown in Fig 1. In such a system, it is difficult to directly measure delay as absolute values in the lines. The method presented in [1] and improved in this paper instead relies on indirect measurements. The proposed method requires only digital components clocked with the existing clock generator.

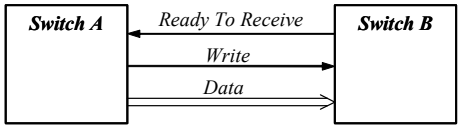

Fig 1. Link from one switch to another

Delay faults might affect one or several wires. Some combinations result only in degradation of the channel throughput. There is also one combination of faults that can result in erroneous data. This fault causes the signal *Write* arrive at the receiver side before *Data*. The receiver reads the data on its next active clock edge after signal *Write* has become high. The time period between *Write* rises until the receiver actually reads the data lines is random in the interval [0, T_R], where T_R is the clock period of the receiver. Let t_l be the time gap between the *Data* and *Write* signals. A delay fault is present if $t_l < 0$ but if $-T_R < t_l < 0$ then data will still be read correctly with probability $(T_R - |t_l|)/ T_R$.

In the method presented in [1], the data was read one clock before the arrival of signal *Write*. In this case the data will never be read correctly if the delay fault described above is present. If there is no such delay fault, the data will be read correctly with some probability [1]. In this method, measurement is repeated a large number of times for every link. If we get a correct value at least once during these measurements we know that this link does not have this delay fault. If we do not get a correct value after a certain number of measurements we interpret that the link is faulty.

This method can be made more efficient if we read the data twice, once just before signal *Write* arrives and once just after. Fig 2. shows timings for *Write* and *Data* at the receiver for correct link. Somewhere in each of the shaded areas in the diagram for clock, one active clock edge will appear. As soon as we get a correct/faulty measurement at the clock before/after *Write* has arrived, we know that the fault is absent/present and we can abort testing. After a finite number of tests there is some probability that some chips can neither be labeled as faulty nor as correct.

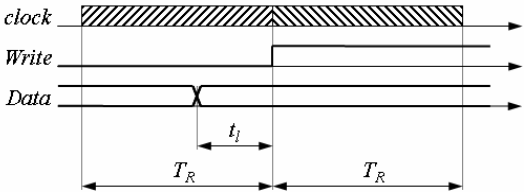

Fig 2. Timing diagram on receiver side

To analyze the performance of this methodology, we consider one data wire and the signal *Write*. We assume that t_l has a normal distribution. Given a probability for fault and given a nominal value on t_l, the variance of t_l can be computed. Let r_n be the percentage of chips for which judgment can not be made. Let n_t be number of tests required on average. It can then be proved that $\lim_{r_n \to 0} n_t(r_t) = \infty$. Given r_n, the maximum and the average number of tests can be computed. The first step in the calculations is to build a function which gives the probability of one test detecting the link fault. From this function, first the maximum number of tests and then r_n can be computed. Table 1 shows sample results from this computation. This is done for different fault probabilities P_f and for different expected t_l. Parameter r_n is chosen to be one tenth of P_f.

P_f / t_l/T_R	Average no tests required n_t			Max number of tests		
	0.1	0.01	0.001	0.1	0.01	0.001
1	2.0 (2.9)	1.8 (2.8)	1.4 (1.7)	13	124	273
0.5	4.8 (13)	3.3 (5.2)	2.6 (3.3)	91	223	619
0.1	24 (62)	16 (27)	13 (15)	448	1178	2244

Table. 1. Performance analysis of the new method

Figures in brackets shows number of tests needed if method in [1] is used. The new method on average terminates faster for bad chips. The improvement is larger for higher fault probabilities. The average number of tests is much smaller than the maximum number.

The testing of delay faults in buses using this new method can be combined with testing of stuck-at and other faults.

REFERENCES

[1] T. Bengtsson, A. Jutman, S. Kumar and R. Ubar, "Delay testing of asynchronous NoC interconnects", 12th International Conference Mixed Design of Integrated Circuits and Systems, June 2005.

A dedicated communication layer for Power Management of NoC based SoCs

N. Genko[1], D. Atienza[1,2], A. Aquaviva[1], G. De Micheli[1]
[1] Laboratoire des Systemes Integres (LSI)/EPFL, 1015 – Lausanne (Switzerland)
[2] DACYA/UCM, 28040 – Madrid (Spain).

Abstract :

In the semiconductor industry, the transistors size is scaling down following Moore's law. This enables the integration of more and more devices on a single chip. In particular it has been observed that forthcoming SoCs will need to include many types of on-chip memories, which will take a very significant part of the die to minimize off-chip communication.

Power breakdown for deep-submicron technology shows that leakage is becoming a major contributor of the total power consumption in SoCs. As memories already occupy most of the on-chip area, leakage of memories is more and more a key element in the overall leakage consumption in SoCs. As a consequence, the design of power manageable memories with leakage reduction support is in great need. Several researchers recently proposed memory designs supporting a leakage-free (or *Sleep*) state.

On this poster we propose a distributed micro-architectural solution to memory leakage power reduction in multiprocessor systems on chip with performance guarantees. The solution is based on a dedicated low-latency interconnect through which distributed power management blocks (Watchers and Controllers) communicate with the purpose of efficiently controlling memory power states. By applying this concept to leakage power management of memories, proactive wake-up signals can be issued from the initiator to the memory to bring the memory out of sleep state just in time to be ready (in idle state) to handle the arriving request without additional latency. Information about incoming transactions is collected at the initiator side by the associated Watcher and proactively propagated on the dedicated communication channel to the Controller of the specified targets.

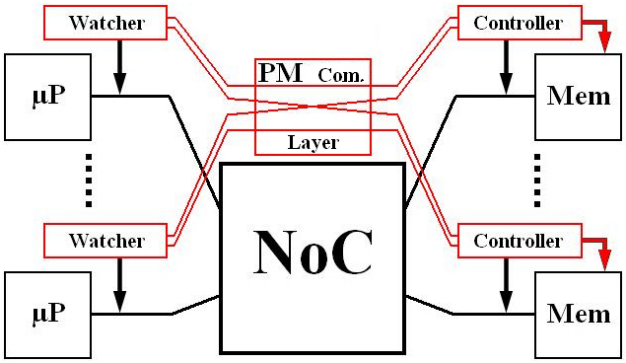

Figure 1: Dual interconnect design for Power management

Results show that the proposed distributed leakage management infrastructure allows us to save up to 55% of leakage energy and more than 30% with respect to local policies with performance guarantees. Moreover, the proposed policy provides savings close to the ideal upper bound. Finally, since power overhead of the synthesized architecture is negligible (10% of the memory leakage power consumption in idle mode), our system represents a viable solution to the problem of the growing contribution of leakage to overall chip power budget.

BIST for NoC Interconnects

Cristian Grecu[1], Partha Pande[2], André Ivanov[1], Res Saleh[1]

[1] SoC Research Laboratory
Department of Electrical and Computer Engineering
University of British Columbia
2356 Main Mall Vancouver, BC, V6T 1Z4 Canada
{grecuc, ivanov, res}@ece.ubc.ca

[2]School of Electrical Engineering & Computer Science
Washington State University PO Box 642752
Pullman, WA 99164-2752, USA
pande@eecs.wsu.edu

Abstract

To date, the Network on Chip (NoC) [1] test methodologies are mainly concerned with the testing of the functional IP cores, using the communication infrastructure as a Test Access Mechanism (TAM) [2]. Our work complements this approach by developing the test strategy for the interconnect infrastructure itself.

We present a built-in self-test methodology for testing the inter-switch links of network-on-chip (NoC)-based chips that uses a high-level fault model that accounts for crosstalk effects due to inter-wire coupling. The novelty of our approach lies in the progressive reuse of the NoC infrastructure to transport test data to its own components under test in a bootstrap manner [3], and in extensively exploiting the inherent parallelism of the data transport mechanism to reduce the test time and implicitly the test cost.

The *Maximal Aggressor Fault* [4] model is an abstract representation of the set of all defects that can lead to one of the six crosstalk errors: rising/falling delay, positive/negative glitch, and rising/falling speed-up. This abstraction covers a wide range of defects including design errors, design rules violations, process variations and physical defects. For a link consisting of N wires, the MAF model assumes the worst-case situation with one victim line and (N-1) aggressors.

It is advantageous to combine the testing of the NoC inter-switch links with that of the other NoC components (i.e., the switch blocks) in order to reduce the test time and total silicon area overhead. In this view, we suggest that the MAF test vectors be organized in test packets consisting of the actual patterns, control signals, and a header containing the routing information. This can be extended by combining the MAF test sequences with the test packets directed to the switches [5]. As shown in Fig. 1, a Global Test Controller (GTC) block injects the test packets intended to the NoC switches, and interleaves those with the MAF test packets produced by the Test Data Generator (TDG) block. The MAF test data is organized in test packets and a header is appended that identifies the link to be tested. The header also contains a field for carrying the routing information, which describes a path along previously tested, fault-free switches and links.

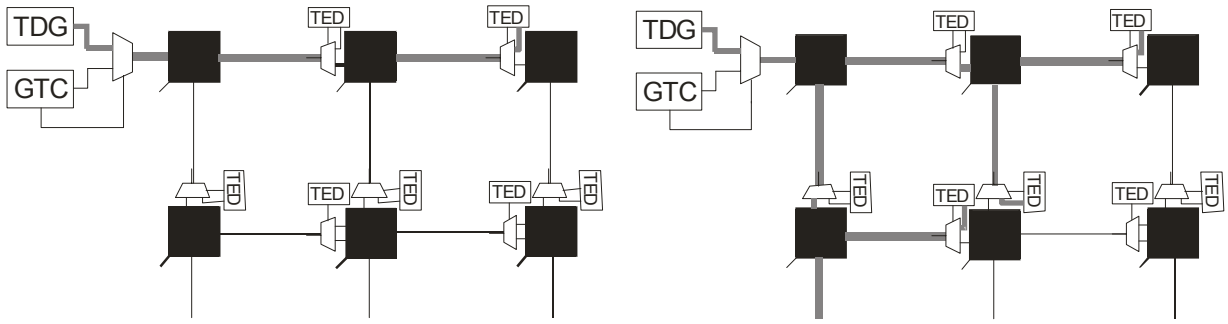

Figure 1: (a) Interleaved unicast MAF test; (b) Interleaved multicast MAF test

We implemented three BIST schemes (point-to-point, unicast, multicast) on an 8-by-8 mesh NoC architecture, similar to the one in Fig. 1a. In each case, we designed and synthesized the hardware blocks used by the corresponding BIST scheme using a standard cell library in a CMOS 130 nm TSMC technology and Synopsys's synthesis tools. The efficiency of each approach was evaluated with respect to the test time and silicon area required. The multicast test approach shows the best trade-off between test time and silicon area required for implementation.

References:

[1] P. P. Pande, C. Grecu, A. Ivanov, R. Saleh, G. De Micheli. "Design, Synthesis, and Test of Networks on Chips," *IEEE Design and Test of Computers*, Vol. 22, No. 5, pp. 404-413, September/October, 2005.

[2] Y. Zorian, E. J. Marinissen, S. Dey, "Testing Embedded-core-based System Chips", IEEE Computer, Vol. 32, No. 6, June 1999, pp. 52-60.

[3] C. Grecu, P. P. Pande, B. Wang, A. Ivanov, R. Saleh. "Methodologies and Algorithms for Testing Switch-Based NoC Interconnects," 20th IEEE International Symposium on Defect and Fault Tolerance in VLSI Systems (DFT'05), 2005, pp. 238-246.

[4] X. Bai, S. Dey, "High-level Crosstalk Defect Simulation for System-on-Chip Interconnects", *19th IEEE VLSI Test Symposium*, Marina del Rey, CA, April 2001, pp. 169-175

[5] C. Grecu, P. P. Pande, A. Ivanov, R. Saleh, "BIST for Network-on-Chip Interconnect Infrastructures", *24th IEEE VLSI Test Symposium*, Berkeley, CA, May 2006.

Fault-Diagnosis-And-Repair System for Improving the Fault-Tolerance and Manufacturability of MPSoCs

Heikki Kariniemi and Jari Nurmi
Institute of Digital and Computers Systems, Tampere University of Technology
P.O.BOX 553, 33101 Tampere, Finland
E-mail: heikki.kariniemi@tut.fi

The number of processors is increasing in the large System-on-Chip circuits and these circuits will resemble distributed computer systems. The processors communicate with packet switched Networks-On-Chip (NOC) which operate like telecommunication or computer networks. The fault-tolerance and reliability of these Multi-Processor SoCs (MPSoC) is influenced by the fault-tolerance of their NOCs. There are redundant processors on the MPSoCs and therefore, faulty processors can be replaced by the faultless processors. However, even a single fault in their NOC may make the whole circuit useless unless the faults can be invalidated and the circuits can be reconfigured to operate correctly. In regard to the manufacturability of the circuits they must be self-diagnosable and self-reconfigurable, because owing to the increasing transistor and wire densities the circuits are more difficult to test and diagnose with traditional test methods. Therefore, systems like the new Fault-Diagnosis-And-Repair (FDAR) system [1], which are able detect and repair static and dynamic faults and which can reconfigure the faulty switch nodes to work correctly despite the faults, are becoming necessary.

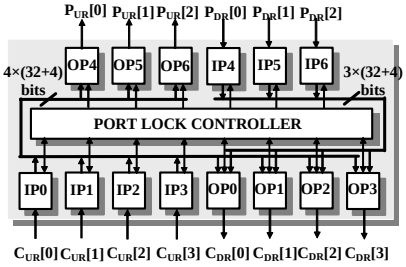

Figure 1. 3x4-switch with port lock controller.

The operation of the FDAR is based on the usage of watchdog timers which measure how long the arbitration or the transfer of the packets takes time in the input ports. If these operations take too much time, the watchdogs launch timeouts. If the switch is operating in a diagnosis mode, it can reconfigure itself by locking the faulty output port - input port pairs which block packets, which invalidates the faults. In a normal operation mode it just removes the packet. The operation mode is determined by the packets' *Mode*-field [1]. Figure 1 depicts a 3×4-switch of an eXtended Generalized Fat Tree (XGFT) NOC with a port lock controller which contains the locks. The switches can self-repair themselves or alternatively the processors probe and reconfigure the network [1]. Version two of the FDAR also allows the processors to reconfigure the switches according to e.g. the results of scan-based tests which show the locations of the faults.

The operation of the FDAR is depicted in Figure 2. The FDAR operates in *Normal* and in *Diagnosis* modes and the timeout levels of the modes are determined with parameters *TIMEOUT_NORMAL* and *TIMEOUT_DIAG* respectively. The value '*0*' of *Blocked* signal shows that the switch is the last one on the blocked packet's routing path and enables the switch to remove the packet. Packet sources have also watchdog timers for controlling the progression of the routing of the packet and they can reconfigure the network by asserting *Remove*-signal if the *repair*-packets are blocked for a longer time. The great advantage of the FDAR is that it repairs switches in such a way that the faultless parts of the switches are still usable after the reconfigurations. The faulty XGFTs can produce good performance despite faults if deterministic routing is used in the faulty parts of the network while adaptive Turn-Back (TB) routing is used in the other parts [1]. The FDAR can be used on-line during systems operation and it is also usable in 2-D mesh networks.

```
IF "Mode = "11"" OR "Mode = "10"" THEN  --"NORMAL MODE"
  IF "Watchdog > TIMEOUT_NORMAL" AND "Blocked = '0'" THEN
    "Remove packet";
  ELSE
    "Route packet forward when it becomes possible";
  END IF;
ELSIF "Mode = "01"" OR "Mode = "00"" THEN  --"DIAG. MODE"
  IF "Watchdog > TIMEOUT_DIAG" OR "Remove = '1'" THEN
    IF "Blocked = '0'" THEN
      IF "Mode = "00"" THEN
        "Reconfigure switch and remove packet";
      ELSE
        "Remove packet";
      END IF;
    ELSE
      "Route packet forward when it becomes possible";
    END IF;
  END IF;
END IF;
```

Figure 2. The operation of the FDAR.

[1] H. Kariniemi, and J. Nurmi, "Fault-Tolerant XGFT Network-On-Chip for Multi-Processor System-on-Chip Circuits," *Proceedings of the 15th International Conference on Field Programmable Circuits and Applications (FPL 2005)*, Tampere, Finland, 2005, pp. 203-210.

Versatile XGFT Network-On-Chip with Improved Fault-Tolerance for Multi-Processor Systems-on-Chip

Heikki Kariniemi and Jari Nurmi
Institute of Digital and Computers Systems, Tampere University of Technology
P.O.BOX 553, 33101 Tampere, Finland
E-mail: heikki.kariniemi@tut.fi

Communication infrastructures of large System-on-Chip (SoC) circuits will be implemented with packet switched networks-on-chip (NOC) which will be used as reusable intellectual property (IP) blocks in platform based design flows used in SoC design. Because different embedded software applications run by the processors of the SoCs will produce different traffic patterns due to the different mappings of their communicating processes to the processors, the NOCs should be able to produce good performance with different traffic patterns. They should also be scalable for different system sizes and performance requirements. Furthermore, they should be fault-tolerant to produce reliable operation and to allow the manufacturing of SoCs with feasible chip yield and costs as soon as possible after the manufacturing is started.

XGFTs are hierarchical multistage interconnection networks where switches in different switch stages can be of different degrees. In addition, their height can be chosen flexibly. For this reason, they are more scalable for different system sizes, performance requirements, and reliability requirements than Fat Trees which they have evolved from. They produce good performance with different traffic patterns, which is shown with simulations [1]. They are very competitive with two-dimensional mesh networks, because they can produce similar performance with smaller resource consumptions. The 2-D mesh networks are not as scalable for different performance requirements, because the number of their switch nodes is always the same as the number of the processors. The shaping of the traffic patterns can be used for reducing the number of switches of the XGFTs, because more local traffic patterns do not need as many switches on top of the network as random traffic.

The scalability of the XGFTs is illustrated in Figure 1. Subfigure A depicts XGFT(3,3,4,3,3,2,1) of height 3 with 36 leaf nodes and a management interface, which could also be constructed from more than one switch node. Subfigure B depicts XGFT(2,6,6,4,1) of height 2 with the same number of leaf nodes. Finally, subfigure C depicts a compact Backbone layout scheme suitable for both of the networks. The lowest switches are in the horizontal channels between the processors and the upper switches surrounded with a dashed line in the vertical channel in the middle of the layout.

Recently, the fault-tolerance of the XGFTs has been improved by a new fault-diagnosis-and-repair (FDAR) system [2], which is able to detect and repair static and dynamic faults during system's operation. It reconfigures the faulty network to work correctly despite the faults. Adaptive Turn-Back (TB) routing makes it possible to achieve high performance in the faultless part of the XGFTs while deterministic or oblivious routing can be used in the faulty parts of the system. The transient bit errors are tackled with parity bit checks which are performed in every link and for every packet separately.

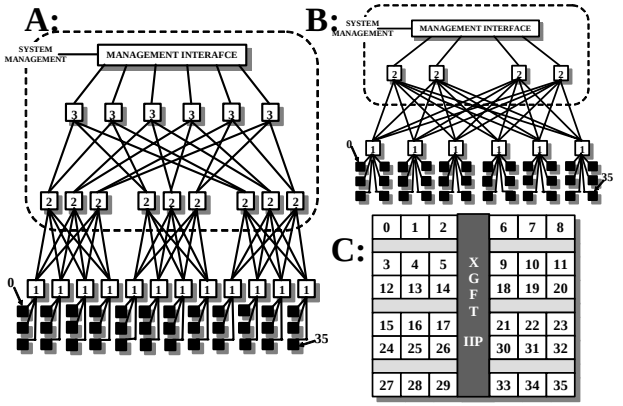

Figure 1. XGFT(3, 3, 4, 3, 3, 2, 1) (A), XGFT(2, 6, 6, 4, 1) (B), and Backbone Layout (C).

The problem of the networks like the Fat Tree and the XGFT is the length of the links between the switch nodes. However, the links can always be pipelined or implemented with low-voltage-swing or asynchronous signaling to reduce delays and power consumption.

[1] H. Kariniemi, and J. Nurmi, "Reusable XGFT Interconnect IP for Network-On-Chip Implementations," *Proc. of the International Symposium on System-on-Chip 2004 (SoC 2004)*, Tampere, Finland, 2004, pp. 95-102.

[2] H. Kariniemi, and J. Nurmi, "Fault-Tolerant XGFT Network-On-Chip for Multi-Processor System-on-Chip Circuits," *Proc. of the 15th International Conference on Field Programmable Circuits and Applications (FPL 2005)*, Tampere, Finland, 2005, pp. 203-210.

Utilizing NoC Switches as BIST-structures in 2D Mesh Network-on-Chip

Kim Petersén, Johnny Öberg

Royal Institute of Technology (KTH), School of Information and Communication Technology (ICT), Dept. of Electronic, Computer and Software Systems (ECS), Email: {kimp, johnny}@imit.kth.se

Abstract[1]

This poster proposes a test methodology and presents a method for carrying out an automatic go/no-go BIST operation at start up of a 2D-mesh NoC Network. It executes in functional mode at full clock speed. Only minor area penalty is introduced in the NoC-network itself; the BIST is placed in the Network-Interface inside the computational resources.

1. Introduction

The Nostrum NoC from KTH is a packet switched Network-on-Chip (NOC) architecture designed to be scalable from a few dozens to several hundreds or even thousands of heterogeneous resources of different kinds, for instance processor cores, DSP cores, FPGA blocks, dedicated HW blocks, memory blocks, etc. To be able to structurally test such a vast and complex chip in a reasonable amount of time, a test methodology for NoCs has to be designed that heavily relies on the usage of Built-In-Self-Test (BIST) structures. This poster outlines a test methodology for carrying out an automatic BIST operation at start-up of a 2-D Mesh NoC. The methodology is general enough to be portable to any NoC structure.

2. Structure to explore

The Nostrum NoC Backbone consists of Resources (R) and Switches (S) organised in a 2D Mesh Manhattan-like structure. A switch is connected to a single one resource in a one-to-one correspondence. The Resources are logic modules that implement various functionality. The signalling between any two Switches and between Switch and Resource is packet based. Each packet is 128 bits wide and consists of Header (32 bits) and User Data (96 bits). The Switches makes an independent routing decision for each individual packet in a single clock cycle. No internal buffering in the switch is used.

All Resources are equipped with a Network Interface (NI) to communicate between the Resource core and the Network. The NI handles the incoming packets from the Switch. A Resource-Network-Interface (RNI) is included for adapting the network to various different local communication protocol standards like, for instance, the AMBA-bus.

3. Start-up test methodology

The boot strap start up test sequence for the Nostrum NoC goes through the following phases: 1) Reset; 2) The BIST sequence start; 3) The NIs log results of its own Switch-Link test; 4a) The (Test) Operating System ((T)OS) Node starts collect test result from all nodes; 4b) NIs wait for initiation from the node containing the (T)OS-node; 5a) NIs respond to the (T)OS request by sending its test log; 5b) the OS-node collect results from the NIs; 6) Normal Operation.

After the BIST sequence has been completed, the (T)OS-node starts booting the system by sending out a query to each node. This query function as a time marker to synchronise the NI's and mark which of the clock cycle windows that should function as a control channel. Each NI answers with its test log. The (T)OS starts querying its neighbourhood, and move its way outward detecting and circumventing broken links on the way. In case all four links out from the (T)OS-node is broken, i.e., the (T)OS-node is not reachable, the chip is considered broken beyond repair.

4. Results

The BIST sequence is comprised of two main steps, called phase 1 and 2. During phase 1 the Switches themselves work as BIST engines; test packets are sent into the switches by choosing an appropriate register reset value. The NIs collects the test result when packets arrive. During phase 2, the NI sends values to the Switch and collects test data.

The total length of the BIST execution was varied to see how fast a good test coverage could be achieved:

	Whole Switch	Datapath
- After 20 clock cycles	46%	68%
- After 100 clock cycles	67%	87%
- After 213 clock cycles	80%	97%

To improve the controllability of the datapath and the observability of the controller, five 2to1-muxes were added to the switch. This makes it possible to achieve a reasonably good fault-coverage very fast for the datapath. A deeper analysis of the controller in the switch and its testpatterns needs to be performed in order to achieve a better total fault coverage. This will be done in the future.

For the links, no extra test logic is added. The switch is executed in functional mode. For the NIs, test logic is added, i.e. a signature register and some logic to keep track of test time and packets to be loaded during phase 2. The NI is running in test mode during execution of the BIST. The result from the BIST is saved as a signature, collected by the NI.

Phase 1 and 2 create a fast go/ no-go test of the NoC. A repair functionality can be added by introducing a phase 3 of the BIST. Phase 1 and 2 can be used to decide if a phase 3 is needed or not and can be triggered by the (T)OS when needed. The presented BIST assumes that reset can be released in the same clock cycle for all switches and NIs.

5. Conclusion and Future Work

The proposed BIST methodology enables a fast go/no-go BIST, with minor extra area in the NoC itself. A test coverage of 97% was obtained for the datapath, and 80% for the whole switch, in only 213 clock cycles.

1.This work is supported by Industrial Research School in Electronic Design.