# Network-on-Chip Architectures

**Lecture Notes in Electrical Engineering**

Volume 45

For other titles published in the series, go to
www.springer.com/series/7818

Chrysostomos Nicopoulos
Vijaykrishnan Narayanan • Chita R. Das

# Network-on-Chip Architectures

## A Holistic Design Exploration

Springer

Chrysostomos Nicopoulos
Department Electrical and Computer Engineering
University of Cyprus
Cyprus
nicopoulos@ucy.ac.cy

Vijaykrishnan Narayanan
Department Computer Science and Engineering
Pennsylvania State University
University Park PA
USA

Chita R. Das
Department of Computer Science and Engineering
Pennsylvania State University
University Park PA
USA

ISBN 978-90-481-3030-6 e-ISBN 978-90-481-3031-3
DOI 10.1007/978-90-481-3031-3
Springer Dordrecht Heidelberg London New York

Library of Congress Control Number: 2009929392



Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

# Abstract

The continuing reduction of feature sizes into the nanoscale regime has led to dramatic increases in transistor densities. Computer architects are actively pursuing multi-core designs with billions of transistors on a single die. Integration at these levels has highlighted the criticality of the on-chip interconnects; global interconnect delays are dominating gate delays and affecting overall system performance. Packet-based Network-on-Chip (NoC) architectures are viewed as a possible solution to burgeoning global wiring delays in many-core chips, and have recently crystallized into a significant research domain. NoCs are steadily becoming the de facto interconnect solution in complex Systems-on-Chip (SoC), because of their scalability and optimized electrical properties. However, current research also indicates that the chip area and power budgets are increasingly being dominated by the interconnection network. To combat this escalating trend, attention should be paid to the optimization of the interconnect architecture.

Unlike traditional multi-computer macro-networks, on-chip networks instill a new flavor to communication research due to their inherently resource-constrained nature. Scarcity in the area and power budgets devoted to the interconnection fabric necessitates a re-interpretation of the networking paradigm. Furthermore, despite the lightweight character demanded of the NoC components, modern designs require ultra-low communication latencies in order to cope with inflating data bandwidths. These conflicting requirements transform the NoC design process into a grand challenge for the system designer. The work presented in this volume aims to address these issues through a comprehensive and holistic exploration of the design space. To truly appreciate the nuances underlying the NoC realm, the design aspects of the on-chip network are viewed through a penta-faceted prism encompassing five major issues: (1) ***performance***, (2) ***silicon area consumption***, (3) ***power/energy efficiency***, (4) ***reliability***, and (5) ***variability***. These five aspects serve as the fundamental design drivers and critical evaluation metrics in the quest for efficient NoC implementations.

The research described in this volume explores the field by employing a two-pronged approach: (a) ***MICRO-architectural innovations*** within the major NoC components, and (b) ***MACRO-architectural choices*** aiming to seamlessly merge the interconnection backbone with the remaining system modules. These two research threads, along with the aforementioned five key metrics mount a holistic and in-depth attack on most issues surrounding the design and integration of NoCs in modern multi-core architectures. Based on this premise of two complementary core themes, the volume is divided into two corresponding parts; the first part delves into the world of MICRO-architectural exploration of the NoC paradigm, while the second part shifts the focus to a MACRO-architectural abstraction level. Ultimately, both parts work in unison in attacking several pressing issues concerning on-chip interconnects in the new multi/many-core reality.

# Acknowledgments

My immense appreciation and wholehearted gratitude goes to my thesis advisor, Professor Vijaykrishnan Narayanan. I am deeply indebted to him for his constant encouragement, unending support, insightful guidance, and wisdom in all aspects of life. Being a mentor and role model, he exposed me to countless unique opportunities in the field of my studies, and gave me the unique opportunity to interact extensively with world-class researchers in the area of computer engineering. His inexhaustible enthusiasm and inspirational motivation have helped me grow both as a student and as a person. With him at the helm, I have enjoyed the most exhilarating, enjoyable, and rewarding ride of my life; and for that, I will always be eternally grateful. Most importantly, beyond the thrills of a myriad research endeavors, he cared – above all – about my happiness and success in life.

I would also like to thank Professor Ken Jenkins for serving as my thesis co-advisor and for his insightful comments and suggestions.

I am sincerely grateful to Professor Chita R. Das for his invaluable contributions to my work over the last few years. His vast wealth of expertise, broad vision, and perceptive wisdom have led to a hugely successful and fruitful collaboration. It is this collaboration that forms the foundation of this volume.

I am equally indebted to Professor Mary Jane Irwin for her steadfast support and precious advice throughout my presence in the Microsystems Design Lab. I am especially grateful to her for exposing me to a multitude of opportunities and giving me the chance to gain enormous experience and blossom as a researcher.

I am extremely grateful to Professors Yuan Xie and Mahmut Kandemir for their sincere friendship, enlightening discussions, priceless insights, and intellectual contributions during my studies. Furthermore, I would like to thank Professor Ram Narayanan for graciously serving on my Ph.D. thesis committee and for his perceptive and encouraging comments. I also wish to thank Professor Paul S. Weiss for instilling a scientific flavor in the committee and – as the external member – bringing a different, yet invaluable perspective to the table.

Eternal gratitude and my sincerest appreciation go to my colleagues and dear friends Jongman Kim and Dongkook Park for a very productive, supportive, successful, and enjoyable collaboration. I would also like to thank Thomas D. Richardson and Feihui Li for being great friends and collaborators; their material and intellectual contributions to my work are priceless. Many thanks go to Jungsub Kim, Suresh Srinivasan, Aditya Yanamandra, Reetuparna Das, Feng Wang, and Theocharis

Theocharides for their enormous help and support. I am also profoundly grateful to all past and present members of the Microsystems Design Lab for the amazing and unforgettable memories of the last few years.

Many thanks also go to all my friends and relatives for their unremitting and resolute confidence in me and their constant encouragement and support.

Finally, I owe my deepest gratitude and boundless appreciation to my parents, Andreas and Elpiniki, for their endless love, unwavering support, constant encouragement, and unfaltering and sincere belief in me. Without them, this volume would not have happened…

# Contents

# List of Figures

# List of Tables

# List of Abbreviations

| | |
|---|---|
| AC | Allocation comparator |
| ACK | Acknowledgment |
| ACLV | Across chip line-width variations |
| AD | Adaptive |
| AG | Affinity group |
| ASIC | Application specific integrated circuits |
| BC | Bit-complement |
| BIST | Built-in self-test |
| CB | Connection box |
| CC | Cache compression |
| CDMA | Code-division multiple access |
| CMOS | Complementary metal-oxide semiconductor |
| CMP | Chip multi-processor |
| CPU | Central processing unit |
| DAMQ | Dynamically allocated multi-queue |
| DAMQWR | DAMQ-with-recruit-registers |
| DEMUX | De-multiplexer (i.e. decoder) |
| DimDe | Dimensionally decomposed router |
| DLT | Direction lookup table |
| DN | Direct neighbor |
| DNUCA | Dynamic non-uniform cache architecture |
| DOR | Dimension-order routing |
| DRAM | Dynamic random access memory |
| DT | Deterministic |
| dTDMA | Dynamic time-division multiple access |
| E2E | End-to-end |
| ECC | Error correcting codes |
| EDP | Energy-delay product |
| EM | Electromigration |
| F2B | Face-to-back |
| F2F | Face-to-face |
| FBF | Fullest buffer first |

| | |
|---|---|
| FC-CB | Fully connected circular buffer |
| FCFS | First-come first-served |
| FEC | Forward error correction |
| FIFO | First-in first-out |
| FP | Fast path |
| HBH | Hop-by-hop |
| HCE | Hot carrier effects |
| HDL | Hardware description language |
| HoL | Head-of-line |
| IIL | Inter-intra layer |
| ILP | Instruction-level parallelism |
| IN | Indirect neighbor |
| IP | Intellectual property |
| IS | Information set |
| LCR | Leakage classification register |
| LWF | Longest wait first |
| MEMS | Micro-electro-mechanical systems |
| MEP | Multiple entry point |
| MESI | Modified-exclusive-shared-invalid |
| MLBS | Multi-layer buried structures |
| MPSoC | Multi-processor system-on-chip |
| MSI | Modified-shared-invalid |
| MUX | Multiplexer |
| NACK | Negative acknowledgment |
| NBTI | Negative bias temperature instability |
| NC | Network interface controller compression |
| NetInMem | Network-in-memory |
| NIC | Network interface controller |
| NoC | Network-on-chip |
| NR | Normal random |
| NUCA | Non-uniform cache architecture |
| OPC | Optical proximity correction |
| PA | Proximity-aware |
| PBD | Platform-based design |
| PC | Physical channel |
| PDP | Power-delay product |
| PE | Processing element |
| PEF | Performance, energy and fault-tolerance |
| PFA | Path frequency analyzer |
| PFT | Path frequency table |
| PIM | Parallel iterative matching |

| | |
|---|---|
| PS | Path set |
| PTM | Predictive technology model |
| PV | Process variation |
| QoS | Quality-of-service |
| RC | Routing computation |
| RDF | Random dopant fluctuations |
| RoCo | Row-column decoupled router |
| RT | Routing table |
| RTL | Register-transfer level |
| SA | Switch allocation |
| SEC/DED | Single error correction and double error detection |
| SER | Soft-error rate |
| SIMD | Single-instruction multiple-data |
| SIP | Service information provider |
| SNUCA | Static non-uniform cache architecture |
| SoC | System-on-chip |
| SOI | Silicon-on-insulator |
| SON | Service-oriented networking |
| SS | Self-similar |
| ST | Sliding timeslice |
| STB | Set-top box |
| STI | Sony, Toshiba, and IBM |
| TDDB | Time-dependant dielectric breakdown |
| TDMA | Time-division multiple access |
| TLP | Thread-level parallelism |
| TMR | Triple module redundancy |
| TN | Tornado |
| TPA | Thermal proximity-aware |
| TSMC | Taiwan Semiconductor Manufacturing Company |
| TTM | Time-to-market |
| UBS | Unified buffer structure |
| UCL | Unified control logic |
| UR | Uniform random |
| UTP | Unique-Token protocol |
| VA | Virtual channel arbitration |
| VC | Virtual channel |
| VCDAMQ | Virtual channel dynamically allocated multi-queue |
| ViChaR | Virtual channel regulator |
| XBAR | Crossbar |