

SIGDA DATE 2004

+

×

↺

↻

<https://past.date-conference.com/proceedings-archive/2004/YEAR.HTM>



DATE¹⁹

MARCH 25–29, 2019

FLORENCE, ITALY

FIRENZE FIERA

DESIGN, AUTOMATION AND TEST IN EUROPE THE EUROPEAN EVENT FOR



Design,
Automation
and Test
in Europe

DATE⁰⁴

Proceedings Volume I

Edited by
Georges Gielen
Joan Figueras

Sponsored by



CNIT La Défense, Paris, France

February 16-20, 2004

www.date-conference.com

MAIN PAGE

DATE 2004

Proceedings:

Table of Content

Abstracts

Session Index

Author Index

Designers' Forum:

Table of Content

Abstracts

Session Index

Author Index

Design, Automation
and Test in Europe

Qualcomm Exhibit 1065
Qualcomm v. Network System Technologies - IPR2024-00343

MAIN PAGE

DATE 2004

Proceedings:

Table of Content

Abstracts

Session Index

Author Index

Designers' Forum:

Table of Content

Abstracts

Session Index

Author Index

DATE 19

MARCH 25—29, 2019

FLORENCE, ITALY

FIRENZE FIERA

DESIGN, AUTOMATION AND TEST IN EUROPE THE EUROPEAN EVENT FOR

DATE 2004 TABLE OF CONTENTS

Sessions: [\[Plenary\]](#) [\[1A\]](#) [\[1B\]](#) [\[1C\]](#) [\[1D\]](#) [\[1E\]](#) [\[1F\]](#) [\[2A\]](#) [\[2B\]](#) [\[2C\]](#) [\[2E\]](#) [\[2F\]](#) [\[3A\]](#) [\[3B\]](#) [\[3C\]](#) [\[3E\]](#) [\[3F\]](#) [\[4A\]](#) [\[4B\]](#) [\[4C\]](#) [\[4E\]](#) [\[4F\]](#) [\[4G\]](#) [\[5A\]](#) [\[5B\]](#) [\[5C\]](#) [\[5E\]](#) [\[5F\]](#) [\[5G\]](#) [\[P1\]](#) [\[P2\]](#) [\[P3\]](#) [\[6A\]](#) [\[6B\]](#) [\[6C\]](#) [\[6E\]](#) [\[6F\]](#) [\[6G\]](#) [\[7A\]](#) [\[7B\]](#) [\[7C\]](#) [\[7E\]](#) [\[7F\]](#) [\[7G\]](#) [\[8A\]](#) [\[8B\]](#) [\[8C\]](#) [\[8E\]](#) [\[8G\]](#) [\[9A\]](#) [\[9B\]](#) [\[9C\]](#) [\[9E\]](#) [\[9G\]](#) [\[10A\]](#) [\[10B\]](#) [\[10C\]](#) [\[10E\]](#) [\[10F\]](#) [\[10G\]](#) [\[1P4\]](#) [\[1P5\]](#) [\[1P6\]](#)

Volume I

Plenary : [Keynote Session](#)

Moderator: J. Figueras, UP Catalunya, ES

Opportunities and Challenges in Building Silicon Products in 65nm and Beyond [p. 2]
G. Spirakis

1A: [Architectural-Level Power Management](#)

Moderators: J. Henkel, NEC, US; A. Macii, Politecnico di Torino, IT

[PDF icon](#) Fine-Grained Dynamic Voltage and Frequency Scaling for Precise Energy and Performance Trade-Off Based on the Ratio of Off-Chip Access to On-Chip Computation Times [p. 4]
K. Choi, R. Soma, and M. Pedram

[PDF icon](#) Hybrid Architectural Dynamic Thermal Management [p. 10]
K. Skadron

[PDF icon](#) Value-Conscious Cache: Simple Technique for Reducing Cache Access Power [p. 16]
Y. Chang, C. Yang, and F. Lai

[PDF icon](#) State-Preserving vs. Non-State-Preserving Leakage Control in Caches [p. 22]
D. Parikh, K. Sankaranarayanan, Y. Li, K. Skadron, Y. Zhang, and M. Stan

1B: [Formal Verification Using Functional and Structural Information](#)

Moderators: A. Veneris, Toronto U, CA; K. Winkelmann, Infineon Technologies, DE

[PDF icon](#) Arithmetic Reasoning in DPLL-Based SAT Solving [p. 30]
M. Wedler, D. Stoffel, and W. Kurz

[PDF icon](#) Enhanced Diameter Bounding via Structural Transformation [p. 36]
J. Baumgartner and A. Kuehlmann

[PDF icon](#) Improved Symbolic Simulation by Dynamic Functional Space Partitioning [p. 42]
T. Feng, L. Wang, K. Cheng, and A. Lin

1C: [Power, Timing and Diagnosis Constrained Testing](#)

Moderators: S. Kundu, Intel, US; B. Straube, FHG IIS/EAS Dresden, DE

[PDF icon](#) Using BDDs and ZBDDs for Efficient Identification of Testable Path Delay Faults [p. 50]
S. Padmanaban and S. Tragoudas

[PDF icon](#) Level of Similarity: A Metric for Fault Collapsing [p. 56]
I. Pomeranz and S. Reddy

[PDF icon](#) Design of Routing-Constrained Low Power Scan Chains [p. 62]
Y. Bonhomme, P. Girard, L. Guiller, C. Landrault, S. Pravossoudovitch, and A. Virazel

[PDF icon](#) Z-Sets and Z-Detectons: Circuit Characteristics that Simplify Fault Diagnosis [p. 68]
I. Pomeranz, S. Venkataraman, S. Reddy, and B. Seshadri

1D: [Mixed-Signal Circuits and Systems](#)

Moderators: A. Rodriguez-Vazquez, IMSE-CNM, ES; P. Wambacq, IMEC, BE

[PDF icon](#) A 2.7V 350pW 11-b Algorithmic Analog-to-Digital Converter with Single-Ended Multiplexed Inputs [p. 76]
A. Nagari and G. Nicollini

[PDF icon](#) Digital Background Gain Error Correction in Pipeline ADCs [p. 82]
A. Ginés, E. Peralias, and A. Rueda

[PDF icon](#) Digital Ground Bounce Reduction by Phase Modulation of the Clock [p. 88]
M. Badaroglu, G. Gielen, H. De Man, P. Wambacq, G. Van Der Plas, and S. Donnay

[PDF icon](#) Pseudo-Random Sequence Based Tuning System for Continuous-Time Filters [p. 94]
A. Baschiroto, S. D'Amico, F. Corsi, C. Marzocca, and G. Matarrese

1E: [Communication-Centric and Source-Level Optimisations for High-Level Synthesis](#)

Moderators: J. Teich, Erlangen-Nuremberg U, DE; P. Cheung, Imperial College London, UK

[PDF icon](#) A Crosstalk Aware Interconnect with Variable Cycle Transmission [p. 102]
L. Li, N. Vijaykrishnan, M. Kandemir, and M. Irwin

[PDF icon](#) Layout Conscious Bus Architecture Synthesis for Deep Submicron Systems on Chip [p. 108]
N. Thepavasuwan and A. Doboli

[PDF icon](#) Loop Shifting and Compaction for the High-Level Synthesis of Designs with Complex Control Flow [p. 114]
S. Gupta, N. Dutt, A. Nicolau, and R. Gupta

Page 2

Qualcomm Exhibit 1065

Qualcomm v. Network System Technologies - IPR2024-00343

1F: Panel Session: SystemC and System Verilog: Where do They Fit? Where are They Going?

Organiser/Moderator: G. Martin, Cadence Berkeley Labs, US; D. Sciuto, Politecnico di Milano
Panellists:
S. Swan, Cadence, US
F. Ghenassia, STMicroelectronics, FR
P. Flake, Synopsys, US
J. Sraygi, Intel, Israel
W. Rosenstiel, Tübingen U, DE

 [PDF icon](#) SystemC and System Verilog: Where do They Fit? Where are they going? [p. 122]

2A: Low Power Systems and Architectures

Moderators: E. Schmidt, Chip Vision Design Systems, DE; C. Guardiani, PDF Solutions, IT

 [PDF icon](#) Re-Configurable Bus Encoding Scheme for Reducing Power Consumption of the Cross Coupling Capacitance for Deep Sub-Micron Instruction Bus [p. 130]
S. Wong and C. Tsui

 [PDF icon](#) Hierarchical Adaptive Dynamic Power Management [p. 136]
Z. Ren, B. Krogh, and R. Marculescu

 [PDF icon](#) A Self-Tuning Cache Architecture for Embedded Systems [p. 142]
C. Zhang, F. Vahid, and R. Lysecky

 [PDF icon](#) Scheduling Reusable Instructions for Power Reduction [p. 148]
J. Hu, N. Vijayabrahman, S. Kim, M. Kandemir, and M. Irwin

2B: Advanced Formal Verification Techniques

Moderators: R. Drechsler, Bremen U, DE; H. Eveking, TU Darmstadt, DE

 [PDF icon](#) Using Counter Example Guided Abstraction Refinement to Find Complex Bugs [p. 156]
P. Bjesse and J. Kukula

 [PDF icon](#) Cost-Efficient Block Verification for a UMTS Up-Link Chip-Rate Coprocessor [p. 162]
G. Fey, D. Stoffel, H. Trylus, and K. Winkelmann

 [PDF icon](#) Automatic Verification of Safety and Liveness for XScale-Like Processor Models Using WEB Refinements [p. 168]
P. Manolios and S. Srinivasan

2C: New Algorithms for TPG

Moderators: H. Obermeir, Infineon Technologies, DE; M. Hsiao, Virginia Tech., US

 [PDF icon](#) A Probabilistic Method for the Computation of Testability of RTL Constructs [p. 176]
J. Fernandes, M. Santos, A. Oliveira, and J. Teixeira

 [PDF icon](#) Graph-Based Functional Test Program Generation for Pipelined Processors [p. 182]
P. Mishra and N. Dutt

 [PDF icon](#) Automatic Generation of Validation Stimuli for Application-Specific Processors [p. 188]
O. Golubeva, M. Sonza Reorda, and M. Violante

 [PDF icon](#) Efficient Static Compaction of Test Sequence Sets through the Application of Set Covering Techniques [p. 194]
M. Dimopoulos and P. Linardis

2E: Optimisation of Memory Hierarchies

Moderators: R. Bergamaschi, IBM TJ Watson Res. Center, US; R. Hermida, Madrid Complutense U, ES

 [PDF icon](#) Data Reuse Analysis Technique for Software-Controlled Memory Hierarchies [p. 202]
I. Issenin, N. Dutt, E. Brockmeyer, and M. Miranda

 [PDF icon](#) Automatic Tuning of Two-Level Caches to Embedded Applications [p. 208]
A. Gordon-Ross, F. Vahid, and N. Dutt

 [PDF icon](#) Low Static-Power Frequent-Value Data Caches [p. 214]
C. Zhang, J. Yang, and F. Vahid

 [PDF icon](#) Using a Victim Buffer in an Application-Specific Memory Hierarchy [p. 220]
C. Zhang and F. Vahid

2F: Hot Topic -- High Security Smartcards

Organiser/Moderator: M. Renaudin, ITMA Laboratory, FR; F. Bouesse, ITMA Laboratory, FR
Speakers:
P. Proust, Gemplus Corporate R&D Security Technologies, FR
J. Tual, Axalto - Schlumberger, FR
L. Souren, STMicroelectronics, FR
F. Germain, DCSSI - French Government Service on the Security of Information Systems, FR

 [PDF icon](#) High Security Smartcards [p. 228]

3A: New Directions in Low-Power Design

Moderators: M. Miranda, IMEC, BE; W. Nebel, OFFIS, DE

 [PDF icon](#) Energy-Aware Communication and Task Scheduling for Network-on-Chip Architectures Under Real-Time Constraints [p. 234]
J. Hu and R. Marculescu

 [PDF icon](#) A Low Cost Individual-Well Adaptive Body Bias (IWABB) Scheme for Leakage Power Reduction and Performance Enhancement in the Presence of Intra-Die Variations [p. 240]
T. Chen and J. Gregg

 [PDF icon](#) A Logic Level Design Methodology for a Secure DPA Resistant ASIC or FPGA Implementation [p. 246]
K. Turi and I. Verbauwhede

 [PDF icon](#) Power Minimization in a Backlit TFT-LCD Display by Concurrent Brightness and Contrast Scaling [p. 252]
W. Cheng, Y. Hou, and M. Pedram

3B: Advances in SAT

Moderators: P. Bjesse, Synopsys, US; G. Cabodi, Politecnico di Torino, IT

 [PDF icon](#) Managing Don't Cares in Boolean Satisfiability [p. 260]
S. Safarpour, A. Veneris, R. Drechsler, and J. Lee

 [PDF icon](#) Exploiting Signal Unobservability for Efficient Translation to CNF in Formal Verification of Microprocessors [p. 266]
M. Velev

 [PDF icon](#) A Novel SAT All-Solutions Solver for Efficient Preimage Computation [p. 272]
B. Li, M. Hsiao, and S. Sheng

3C: [Analogue and High-Frequency Test](#)

Moderators: *A. Richardson, Lancaster U, UK; F Azais, LIRMM, FR*

-  [PDF icon](#) Efficient Test Strategy for TDMA Power Amplifiers Using Transient Current Measurements: Uses and Benefits [p. 280]
G. Srinivasan, S. Bhattacharya, A. Chatterjee, and S. Cherubal
-  [PDF icon](#) Random Jitter Extraction Technique in a Multi-Gigahertz Signal [p. 286]
C. Ong, D. Hong, K. Cheng, and L. Wang
-  [PDF icon](#) Low Cost Analog Testing of RF Signal Paths [p. 292]
M. Negreiros, L. Carro, and A. Susin
-  [PDF icon](#) A Method for Parameter Extraction of Analog Sine-Wave Signals for Mixed-Signal Built-In-Self-Test Applications [p. 298]
D. Vázquez, G. Huertas, G. Leger, A. Rueda, and J. Huertas

3E: [Energy Efficient Memory Usage](#)

Moderators: *E. de Kock, Philips Research, NL; G. Constantinides, Imperial College, UK*

-  [PDF icon](#) A Novel Implementation of Tile-Based Address Mapping [p. 306]
S. Hettiaratchi and P. Cheung
-  [PDF icon](#) Power Aware Variable Partitioning and Instruction Scheduling for Multiple Memory Banks [p. 312]
Z. Wang and X. Hu
-  [PDF icon](#) Time-Energy Design Space Exploration for Multi-Layer Memory Architectures [p. 318]
R. Szymonek, K. Kuchcinski, and F. Cathoor
-  [PDF icon](#) Breaking Instance-Independent Symmetries in Exact Graph Coloring [p. 324]
A. Ramani, F. Aloul, I. Markov, and K. Sakallah

3F: [Hot Topic -- How Can System Level Design Solve the Interconnect Technology Scaling Problem?](#)

Organiser: *R. Lauwereins, IMEC, BE*

Moderator: *R. Wilson, CMP Media, US*

Panellists:

K. Maex, IMEC, BE

P. Groeneveld, Magma Design Automation, US

G. Martin, Cadence, US

A. Cuomo, STMicroelectronics, IT

F. Cathoor, IMEC, BE

P. van de Steeg, Philips Semiconductors, NL

-  [PDF icon](#) How Can System Level Design Solve the Interconnect Technology Scaling Problem? [p. 332]

4A: [System Level Design Methodology](#)

Moderators: *Y. Mathys, Motorola, FR; H. Hsieh, UC Riverside, US*

-  [PDF icon](#) System Design Using Kahn Process Networks: The Compaan/Laura Approach [p. 340]
E. Deprettere, B. Kienhuis, T. Stefanov, A. Turjan, and C. Zissulescu
-  [PDF icon](#) Microarchitecture Development via Metropolis Successive Platform Refinement [p. 346]
D. Densmore, A. Sangiovanni-Vincentelli, and S. Rekh
-  [PDF icon](#) Fast Exploration of Parameterized Bus Architecture for Communication-Centric SoC Design [p. 352]
C. Shin, Y. Kim, E. Chung, K. Choi, J. Kong, and S. Eo
-  [PDF icon](#) SoftContract: An Assertion-Based Software Development Process that Enables Design-by-Contract [p. 358]
J. Brunel, P. Giusto, M. di Natale, A. Ferrari, and L. Lavagno
-  [PDF icon](#) A System Level Exploration Platform and Methodology for Network Applications Based on Configurable Processors [p. 364]
D. Quim, B. Lavigueur, G. Bois, and M. Aboulhamid

4B: [System Level Modelling and Analysis](#)

Moderators: *S. Singh, Xilinx, US; A. Jantsch, Royal Inst. of Tech., SE*

-  [PDF icon](#) Refinement of Mixed-Signal Systems with Affine Arithmetic [p. 372]
C. Grimm, W. Heupke, and K. Waldschmidt
-  [PDF icon](#) System-Level Performance Analysis in SystemC [p. 378]
H. Posadas, F. Herrera, P. Sánchez, E. Villar, and F. Blasco
-  [PDF icon](#) Modeling and Validating Globally Asynchronous Design in Synchronous Frameworks [p. 384]
M. Mousavi, P. Le Guernic, J. Talpin, S. Shukla, and T. Basten
-  [PDF icon](#) Synchronous Protocol Automata: A Framework for Modelling and Verification of SoC Communication Architectures [p. 390]
V. D'Silva, S. Ramesh, and A. Sowmya
-  [PDF icon](#) Aspects of Formal and Graphical Design of a Bus System [p. 396]
T. Seceleanu and T. Westerlund

4C: [Advances in SoC Testing](#)

Moderators: *R. Dorsch, IBM Deutschland Entwicklung, DE; E. Larsson, Linköping U, SE*

-  [PDF icon](#) Scan Power Minimization through Stimulus and Response Transformations [p. 404]
O. Sinanoglu and A. Orailoglu
-  [PDF icon](#) Synchro-Tokens: Eliminating Nondeterminism to Enable Chip-Level Test of Globally Asynchronous Locally-Synchronous SoC's [p. 410]
M. Heath, W. Burleson, and I. Harris
-  [PDF icon](#) Wrapper Design for Testing IP Cores with Multiple Clock Domains [p. 416]
Q. Xu and N. Nicollet
-  [PDF icon](#) Efficient Modular Testing of SOC's Using Dual-Speed TAM Architectures [p. 422]
A. Sehgal and K. Chakrabarty
-  [PDF icon](#) An Arithmetic Structure for Test Data Horizontal Compression [p. 428]
M. Flottes, R. Poirier, and B. Rouzeyre

4E: New Issues in Analogue System- and Circuit-Level Performance Modelling

Moderators: F. Fernandez, IMSE-CNM, ES; R. Schwenker, Infineon Technologies, DE

-  [PDF icon](#) A Phase-Frequency Transfer Description of Analog and Mixed-Signal Front-End Architectures for System-Level Design [p. 436]
E. Mariens and G. Gielen
-  [PDF icon](#) Hierarchical Automatic Behavioral Model Generation of Nonlinear Analog Circuits Based on Nonlinear Symbolic Techniques [p. 442]
L. Naithe, V. Burkhardt, L. Hedrich, and E. Barka
-  [PDF icon](#) Performance Modeling of Analog Integrated Circuits Using Least-Squares Support Vector Machines [p. 448]
T. Kiely and G. Gielen
-  [PDF icon](#) Extended Subspace Identification of Improper Linear Systems [p. 454]
G. Vandersteen, D. Linten, R. Pintelon, and S. Donnay
-  [PDF icon](#) Identification and Modeling of Nonlinear Dynamic Behavior in Analog Circuits [p. 460]
X. Huang and H. Mantooth

4F: Fabrics and Scheduling for Reconfigurable Computing

Moderators: G. Koch, Micronas GmbH, DE; C. Passerone, Politecnico di Torino, IT

-  [PDF icon](#) Exploring Logic Block Granularity for Regular Fabrics [p. 468]
A. Koorapaty, V. Khaterpal, P. Gopalakrishnan, M. Fu, and L. Pileggi
-  [PDF icon](#) Network Topology Exploration of Mesh-Based Coarse-Grain Reconfigurable Architectures [p. 474]
N. Bansal, S. Gupta, N. Dutt, A. Nicolau, and R. Gupta
-  [PDF icon](#) A Configurable Logic Architecture for Dynamic Hardware/Software Partitioning [p. 480]
R. Lysecky and F. Vahid
-  [PDF icon](#) Configuration-Sensitive Process Scheduling for FPGA-Based Computing Platforms [p. 486]
G. Chen, M. Kandemir, and U. Sesser

4G: Power Aware Design and Synthesis

Moderators: R. Zafalon, STMicroelectronics, IT; K. Roy, Purdue U, US

-  [PDF icon](#) Simultaneous State, Vt and Tox Assignment for Total Standby Power Minimization [p. 494]
D. Lee, H. Deogun, D. Blaauw, and D. Sylvester
-  [PDF icon](#) A Scalable ODC-Based Algorithm for RTL Insertion of Gated Clocks [p. 500]
P. Babighian, E. Macii, and L. Benini
-  [PDF icon](#) Impact of Data Transformations on Memory Bank Locality [p. 506]
M. Kandemir
-  [PDF icon](#) Why Transition Coding for Power Minimization of On-Chip Buses Does Not Work [p. 512]
C. Kretschmar, D. Müller, and A. Nieuwland
-  [PDF icon](#) Overhead-Conscious Voltage Selection for Dynamic and Leakage Energy Reduction of Time-Constrained Systems [p. 518]
A. Andrei, M. Schmitt, P. Eles, Z. Peng, and B. Al-Hashimi

5A: System Level Design: Case Studies, Exploration and Optimisation

Moderators: B. Kienhuis, LACS, Leiden U, NL; F. Petrot, Pierre et Marie Curie U, FR

-  [PDF icon](#) Dynamic Power Management Using Data Buffers [p. 526]
Y. Lu and L. Cai
-  [PDF icon](#) Dynamic Memory Management Design Methodology for Reduced Memory Footprint in Multimedia and Wireless Network Applications [p. 532]
D. Atienza, J. Mendias, S. Managakis, D. Soudris, and F. Catthoor
-  [PDF icon](#) High-Level System Modeling and Architecture Exploration with SystemC on a Network SoC: S3C2510 Case Study [p. 538]
H. Jang, M. Kang, K. Shim, M. Lee, K. Chae, and K. Lee
-  [PDF icon](#) A SystemC-Based Verification Methodology for Complex Wireless Software IP [p. 544]
G. Post, P. Venkataraghavan, T. Roy, and D. Seetharaman

5B: Recent Advances in Digital Systems Simulation

Moderators: M. Zwolinski, Southampton U, UK; M. Lajolo, NEC Laboratories, US

-  [PDF icon](#) A New Optimized Implementation of the SystemC Engine Using Acyclic Scheduling [p. 552]
D. Pérez, O. Temam, and G. Mouchard
-  [PDF icon](#) Stimuli Generation with Late Binding of Values [p. 558]
A. Ziv
-  [PDF icon](#) Native ISS-SystemC Integration for the Co-Simulation of Multi-Processor SoC [p. 564]
F. Fummi, S. Martini, G. Perbellini, and M. Poncino
-  [PDF icon](#) Extraction of Schematic Array Models for Memory Circuits [p. 570]
S. Bose and A. Nandi

5C: On-Line Testing and Reliability for Nanometric Technology

Moderators: T. Mak, Intel Corp., US; Y. Tsiatouhas, Ioannina U, GR

-  [PDF icon](#) Effective Software-Based Self-Test Strategies for On-Line Periodic Testing of Embedded Processors [p. 578]
A. Paschalis and D. Gizopoulos
-  [PDF icon](#) Evaluating the Effects of SEUs Affecting the Configuration Memory of an SRAM-Based FPGA [p. 584]
M. Bellato, A. Bernardi, A. Candelari, M. Rebaudengo, M. Sonza Reorda, M. Violante, M. Ceschia, D. Bortolato, A. Paccagnella, and P. Zambolin
-  [PDF icon](#) Early SEU Fault Injection in Digital, Analog and Mixed Signal Circuits: A Global Flow [p. 590]
R. Leveugle and A. Ammar
-  [PDF icon](#) On Concurrent Error Detection with Bounded Latency in FSMs [p. 596]
S. Almukhatim, Y. Makris, and P. Drineas

5E: Parasitic-Aware Analogue Design

Moderators: H. Graeb, TU Munich, DE; G. Vandersteen, IMEC, BE

-  [PDF icon](#) Fast, Layout-Inclusive Analog Circuit Synthesis Using Pre-Compiled Parasitic-Aware Symbolic Performance Models [p. 604]
M. Rangan, A. Agarwal, H. Sampath, R. Vemuri, W. Verhaegen, and G. Gielen
-  [PDF icon](#) Sensitivity-Based Modeling and Methodology for Full-Chip Substrate Noise Analysis [p. 610]
R. Murgai, S. Reddy, T. Miyoshi, T. Horie, and M. Tahoori
-  [PDF icon](#) SubCALM: A Program for Hierarchical Substrate Coupling Simulation on Floorplan Level [p. 616]
T. Brandtner and R. Weigel
-  [PDF icon](#) Optimization of Integrated Spiral Inductors Using Sequential Quadratic Programming [p. 622]
Y. Zhan and S. Sapatnekar

5F: Hardware/Software System Design and Architecture Exploration

Moderators: B. Juurlink, TU Delft, NL; R. Leupers, RWTH Aachen, DE

-  [PDF icon](#) System Design for DSP Applications Using the MASIC Methodology [p. 630]
A. K. Deb, A. Jantsch, and J. Öberg
-  [PDF icon](#) Flexible Software Protection Using Hardware/Software Codesign Techniques [p. 636]
J. Zambreno, A. Choudhary, R. Simha, and B. Narahari
-  [PDF icon](#) Interactive Cosimulation with Partial Evaluation [p. 642]
P. Schaumont and I. Verbauwhede
-  [PDF icon](#) Communication Analysis for System on Chip Design [p. 648]
A. Siebenborn, O. Bringmann, and W. Rosenstiel

5G: Hot Topic -- Extremely Low-Power Logic

Organizer/Moderator: C. Piguet, CSEM, CH

Presenters:

J. Gautier, CEA-LETI, FR
C. Heer, Infineon Technologies, DE
I. O'Connor, Ecole centrale de Lyon, FR
U. Schlichtmann, Technical University of Munich, DE

-  [PDF icon](#) Extremely Low-Power Logic [p. 656]

IP1: Interactive Presentations

-  [PDF icon](#) Decomposition of Instruction Decoder for Low Power Design [p. 664]
W. Kuo, T. Hwang, and A. Wu
-  [PDF icon](#) Functional Level Power Analysis: An Efficient Approach for Modeling the Power Consumption of Complex Processors [p. 666]
J. Laurent, N. Julien, E. Senn, and E. Martin
-  [PDF icon](#) Formal Verification Coverage: Are the RTL-Properties Covering the Design's Architectural Intent? [p. 668]
P. Basu, S. Das, P. Dasgupta, P. Chakrabarti, C. Mohan, and L. Fx
-  [PDF icon](#) Functional Coverage Metric Generation from Temporal Event Relation Graph [p. 670]
Y. Kwon and C. Kyung
-  [PDF icon](#) Automatic Scan Insertion and Pattern Generation for Asynchronous Circuits [p. 672]
A. Efthymiou, D. Edwards, and C. Sotiriou
-  [PDF icon](#) Automatic Synthesis and Simulation of Continuous-Time $\Sigma\Delta$ Modulators [p. 674]
H. Aboushady, L. de Lamarre, N. Beillean, and M. Louérat
-  [PDF icon](#) A Methodology for System-Level Analog Design Space Exploration [p. 676]
F. De Bernardinis and A. Sangiovanni-Vincentelli
-  [PDF icon](#) Systematic Design for Optimization of High-Resolution Pipelined ADCs [p. 678]
R. Lotfi, M. Taherzadeh-Sani, and O. Shoaei
-  [PDF icon](#) A Direct Bootstrapped CMOS Large Capacitive-Load Driver Circuit [p. 680]
J. Garcia, J. Montiel-Nelson, J. Sosa, and H. Navarro
-  [PDF icon](#) Co-Processor Synthesis: A New Methodology for Embedded Software Acceleration [p. 682]
B. Hounsell and R. Taylor
-  [PDF icon](#) Behavioural Bitwise Scheduling Based on Computational Effort Balancing [p. 684]
M. Molina, R. Ruiz-Santua, J. Mendias, and R. Hermida
-  [PDF icon](#) A Tool for Automatic Generation of RTL-Level VHDL Description of RNS FIR Filters [p. 686]
A. Nannarelli, A. Del Re, and M. Re
-  [PDF icon](#) On Transfer Function and Power Consumption Transient Response [p. 688]
L. Cao
-  [PDF icon](#) Polynomial Abstraction for Verification of Sequentially Implemented Combinational Circuits [p. 690]
T. Raudvere, A. Singh, I. Sander, and A. Jantsch
-  [PDF icon](#) Regression Simulation: Applying Path-Based Learning in Delay Test and Post-Silicon Validation [p. 692]
L. Wang

IP2: Interactive Presentations

-  [PDF icon](#) A Game Theoretic Approach to Low Energy Wireless Video Streaming [p. 696]
A. Iranli, K. Choi, and M. Pedram
-  [PDF icon](#) Block-Enabled Memory Macros: Design Space Exploration and Application-Specific Tuning [p. 698]
A. Ivaldi, A. Macii, E. Macii, and L. Benini
-  [PDF icon](#) Synthesis of Partitioned Shared Memory Architectures for Energy-Efficient Multi-Processor SoC [p. 700]
E. Macii, K. Patel, and M. Poncino
-  [PDF icon](#) A Low Power Strategy for Future Mobile Terminals [p. 702]
M. Nikitovic and M. Brorsson
-  [PDF icon](#) A 0.18 μm CMOS Implementation of On-Chip Analogue Test Signal Generation from Digital Test Patterns [p. 704]
L. Rolindez, S. Mir, G. Prenat, and A. Bounceur
-  [PDF icon](#) A Digital Test for First-Order $\Sigma\Delta$ Modulators [p. 706]
G. Leger and A. Rueda
-  [PDF icon](#) Trim Bit Setting of Analog Filters Using Wavelet-Based Supply Current Analysis [p. 708]
S. Bhunia, A. Raychowdhury, and K. Roy
-  [PDF icon](#) SoC Test Scheduling with Power-Time Tradeoff and Hot Spot Avoidance [p. 710]
J. Chin and M. Nourani
-  [PDF icon](#) STEPS: Experimenting a New Software-Based Strategy for Testing SoCs Containing P1500-Compliant IP Cores [p. 712]
M. Benabdenbi, F. Pécheux, A. Greiner, M. Tuna and E. Viaud
-  [PDF icon](#) Are Our Designs for Testability Features Fault Secure? [p. 714]
C. Metra, M. Omana, and T. Mak
-  [PDF icon](#) Test Compression and Hardware Decompression for Scan-Based SoCs [p. 716]
F. Wolff, C. Papachristou, and D. McIntyre
-  [PDF icon](#) Concurrent Sizing, V_{dd} and V_{th} Assignment for Low-Power Design [p. 718]
A. Srivastava, D. Sylvester, and D. Blaauw
-  [PDF icon](#) Sizing and Characterization of Leakage-Control Cells for Layout-Aware Distributed Power-Gating [p. 720]
P. Babighian, E. Macii, and L. Benini

IP3: [Interactive Presentations](#)

-  [PDF icon](#) An Asynchronous Synthesis Toolset Using Verilog [p. 724]
F. Burns, D. Shang, A. Koelmans, and A. Yakovlev
-  [PDF icon](#) Organizing Libraries of DFG Patterns [p. 726]
G. Dittmann
-  [PDF icon](#) Compositional Memory Systems for Data Intensive Applications [p. 728]
A. Molnos, M. Heijltjers, J. Van Eijndhoven, and S. Cotozana
-  [PDF icon](#) Scalar Metric for Temporal Locality and Estimation of Cache Performance [p. 730]
J. Alakarhu and J. Nittylahti
-  [PDF icon](#) NET Framework -- A Solution for the Next Generation Tools for System-Level Modeling and Simulation [p. 732]
J. Lapalme, E. Aboulhamid, G. Nicolescu, L. Charest, J. David, F. Boyer, and G. Bois
-  [PDF icon](#) Modeling and Simulating Memory Hierarchies in a Platform-Based Design Methodology [p. 734]
P. Viana, E. Barros, S. Rigo, R. Azevedo, and G. Araujo
-  [PDF icon](#) Integrating the Synchronous Dataflow Model with UML [p. 736]
P. Green and S. Essa
-  [PDF icon](#) Design and Behavioral Modeling Tools for Optical Network-On-Chip [p. 738]
M. Briere, L. Carrel, T. Michalke, F. Mueyville, I. O'Connor, and F. Gaffiot
-  [PDF icon](#) Hierarchical Modeling and Simulation of Large Analog Circuits [p. 740]
S. Tan, Z. Qi, and H. Li
-  [PDF icon](#) Efficient Mixed-Domain Behavioural Modeling of Ferromagnetic Hysteresis Implemented in VHDL-AMS [p. 742]
P. Wilson, J. Ross, A. Brown, T. Karmierski, and J. Baranowski
-  [PDF icon](#) A Fast Algorithm for Finding Maximal Empty Rectangles for Dynamic FPGA Placement [p. 744]
M. Handa and R. Vemuri
-  [PDF icon](#) Enhancing Reliability of Operational Interconnections in FPGAs [p. 746]
A. Fit-Florea, M. Halas, and F. Kocan
-  [PDF icon](#) Operating System Support for Interface Virtualization of Reconfigurable Coprocessors [p. 748]
M. Vuletic, L. Righetti, L. Pozzi, and P. lenne

Volume II

6A: [Performances Analysis for MPSoC](#)

Moderators: R. Ernst, TU Braunschweig, DE; A. Jantsch, Royal Inst. of Tech., SE

-  [PDF icon](#) Analyzing On-Chip Communication in a MPSoC Environment [p. 752]
F. Angiolini, D. Bertozzi, L. Benini, M. Loghi, and R. Zafalon
-  [PDF icon](#) A Mapping Strategy for Resource-Efficient Network Processing on Multiprocessor SoCs [p. 758]
M. Grünewald, J. Niemann, M. Porrmann, and U. Rückert
-  [PDF icon](#) Cost-Performance Trade-Offs in Networks on Chip: A Simulation-Based Approach [p. 764]
S. Pestana, E. Ruykema, A. Radulescu, K. Goossens, and O. Gangwal
-  [PDF icon](#) A Case Study in Networks-on-Chip Design for Embedded Video [p. 770]
J. Xu, W. Wolf, T. Lv, J. Henkel, and S. Chakrathar

6B: [Synthesis for Noise and Manufacturability](#)

Moderators: T. Villa, Udma U, IT; T. Shiple, Synopsis, FR

-  [PDF icon](#) Exploiting Crosstalk to Speed up On-Chip Buses [p. 778]
C. Duan and S. Khatri
-  [PDF icon](#) False-Noise Analysis for Domino Circuits [p. 784]
A. Glebov, S. Gavrilov, V. Zolotov, M. Becker, C. Oh, and R. Panda
-  [PDF icon](#) Crosstalk Minimization in Logic Synthesis for PLA [p. 790]
Y. Liu, T. Hwang, and K. Wang
-  [PDF icon](#) Synthesis for Manufacturability: A Sanity Check [p. 796]
A. Nardi and A. Sangiovanni-Vincentelli

6C: [Support for BIST](#)

Moderators: G. Carlsson, Ericsson Telecom, SE; K. Chakrabarty, Duke U, US

-  [PDF icon](#) Design of Sub-10-Picoseconds On-Chip Time Measurement Circuit [p. 804]
M. Abbas, D. Kinniment, and G. Russell
-  [PDF icon](#) Impact of Test Point Insertion on Silicon Area and Timing during Layout [p. 810]
H. Vranken, H. Wunderlich, and F. Sapei
-  [PDF icon](#) Designing Self Test Programs for Embedded DSP Cores [p. 816]
H. Rizk, C. Papachristou, and F. Wolff

6E: [Modelling, Simulation and Optimisation in Power/Ground/Substrate](#)

Moderators: P. Feldmann, IBM T.J. Watson Res. Center, US; L. Silveira, INESC ID/IST - TU Lisbon, PT

-  [PDF icon](#) Automated, Accurate Macromodelling of Digital Aggressors for Power/Ground Substrate Noise Prediction [p. 824]
Z. Wang, J. Roychowdhury, and R. Murgai
-  [PDF icon](#) Thermal and Power Integrity Based Power/Ground Networks Optimization [p. 830]
T. Wang, J. Tsai, and C. Chen
-  [PDF icon](#) Synthesized Compact Models (SCM) of Substrate Noise Coupling Analysis and Synthesis in Mixed-Signal ICs [p. 836]
H. Lan and R. Dutton

6F: [Panel Session -- Chips of the Future: Soft, Crunchy or Hard?](#)

Organizer/Moderator: P. Paulin, STMicroelectronics, FR

Panelists:

R. Bramley, STMicroelectronics
A. Silburt, Cisco, CAN
J. Balzano, Alcatel, FR
K. van Berkel, Philips Research, NL
N. Wehn, Kaiserslautern U, DE

-  [PDF icon](#) Chips of the Future: Soft, Crunchy or Hard? [p. 844]

6G: Power-Aware Networks and Interfaces (Low Power Special Day)

Moderators: *M. Pedram, Southern California U, US; A. Amara, ISEP, FR*

-  [PDF icon](#) Tuning In-Sensor Data Filtering to Reduce Energy Consumption in Wireless Sensor Networks [p. 852]
I. Kadayif and M. Kandemir
-  [PDF icon](#) Power-Aware Network Swapping for Wireless Palmtop PCs [p. 858]
A. Acquaviva, E. Lattanzi, and A. Bogliolo
-  [PDF icon](#) Power Aware Interface Synthesis for Bus-Based SoC Design [p. 864]
N. Liversi and P. Banerjee
-  [PDF icon](#) Asynchronous Design by Conversion: Converting Synchronous Circuits into Asynchronous Ones [p. 870]
A. Branover, R. Kol, and R. Ginosar

7A: Networks on Chip Design

Moderators: *G. Nicosescu, Ecole Polytechnique de Montreal, CA; M. Coppola, STMicroelectronics, FR*

-  [PDF icon](#) An Efficient On-Chip Network Interface Offering Guaranteed Services, Shared-Memory Abstraction, and Flexible Network Configuration [p. 878]
A. Radulescu, J. Dieltjen, K. Goossens, E. Rijpkema, and P. Wielage
-  [PDF icon](#) *pipesCompiler: A Tool for Instantiating Application Specific Networks-on-Chip [p. 884]
S. Murali, G. De Micheli, A. Jalabert, and L. Benini
-  [PDF icon](#) Guaranteed Bandwidth Using Looped Containers in Temporally Disjoint Networks within the Nostrum Network on Chip [p. 890]
M. Millberg, E. Nilsson, R. Thid, and A. Jantsch
-  [PDF icon](#) Bandwidth-Constrained Mapping of Cores onto NoC Architectures [p. 896]
S. Murali and G. De Micheli

7B: Advances in Technology Mapping and Circuit Sizing

Moderators: *T. Kuttischebauch, Magma Design Automation, US; L. Stok, IBM, US*

-  [PDF icon](#) Synthesis and Optimization of Threshold Logic Networks with Application to Nanotechnologies [p. 904]
R. Zhang, P. Gupta, L. Zhong, and N. Jha
-  [PDF icon](#) Fast Comparisons of Circuit Implementations [p. 910]
S. Karandikar and S. Sapatnekar
-  [PDF icon](#) Saving Power by Mapping Finite-State Machines into Embedded Memory Blocks in FPGAs [p. 916]
A. Trwari and K. Tomko
-  [PDF icon](#) MemMap: Technology Mapping Algorithm for Area Reduction in FPGAs with Embedded Memory Arrays Using Reconvergence Analysis [p. 922]
M. Kumar, J. Bobba, and V. Kamakoti

7C: Panel Session -- Nanometer Design: What are the Requirements for Manufacturing Test?

Organizer: *K. Thapar, Mentor Graphics Europe, UK*

Moderator: *J. Rajski, Mentor Graphics, US*

Panelists:

M. Vorenkamp, STMicroelectronics, FR
S. Laskar, S. Vradkova, S. Chopra, and K. Srinivasan

-  [PDF icon](#) Local Decisions and Triggering Mechanisms for Dynamic Fault Tolerance [p. 968]
P. Stanley-Marbell and D. Marculescu
-  [PDF icon](#) An Algorithm for Nano-Pipelining of Circuits and Architectures for a Nanotechnology [p. 974]
P. Gupta and N. Jha
-  [PDF icon](#) Smaller Two-Qubit Circuits for Quantum Communication and Computation [p. 980]
V. Shende, I. Markov, and S. Bullock

7G: Embedded Tutorial -- Architectures and Design Techniques for Energy Efficient Embedded DSP and Multimedia Processing

Organizer: *E. Macii, Politecnico di Torino, IT*

Moderator: *N. Chang, Seoul National U, KR*

Speakers:

J. Verbaanwade, UCLA, US

C. Piquet, CSEM, CH

P. Schaumont, UCLA, US

B. Kienhuis, Leiden U, NL

-  [PDF icon](#) Architectures and Design Techniques for Energy Efficient Embedded DSP and Multimedia Processing [p. 988]

8A: Platform-Based Design and VC Reuse Methods

Moderators: *R. Seepold, Carlos III de Madrid U, ES; T. Riesgo, UP Madrid, ES*

-  [PDF icon](#) Measurement of IP Qualification Costs and Benefits [p. 996]
A. Varg, W. Rosenstiel, and M. Radetzki
-  [PDF icon](#) Architecture-Level Performance Estimation for IP-Based Embedded Systems [p. 1002]
K. Ueda, K. Sakamushi, Y. Takeuchi, and M. Imai
-  [PDF icon](#) Generalized Latency-Insensitive Systems for Single-Clock and Multi-Clock Architectures [p. 1008]
M. Singh and M. Theobald
-  [PDF icon](#) Platform Based on Open-Source Cores for Industrial Applications [p. 1014]
M. Bolado, J. Castillo, P. Huerta, H. Posadas, P. Sánchez, C. Sánchez, F. Blasco, and H. Fournen
-  [PDF icon](#) MINCE: Matching Instructions with Combinational Equivalence for Extensible Processor [p. 1020]
N. Cheung, S. Parameswaran, J. Henkel, and J. Chan

8B: Real-Time Issues in Embedded Systems

Moderators: *S. Hu, Notre Dame U, US; F. Wolf, Volkswagen, DE*

-  [PDF icon](#) Design Optimization of Multi-Cluster Embedded Systems for Real-Time Applications [p. 1028]
P. Pop, P. Eles, Z. Peng, V. Izosimov, M. Hellring, and O. Bridal
-  [PDF icon](#) Tuning Analysis for Preemptive Multi-Tasking Real-Time Systems with Caches [p. 1034]
Y. Tan and V. Mooney
-  [PDF icon](#) Workload Characterization Model for Tasks with Variable Execution Demand [p. 1040]
A. Maxiaoguine, S. Künzli, and L. Thiele
-  [PDF icon](#) Context-Aware Performance Analysis for Efficient Embedded System Design [p. 1046]
M. Jersak, R. Henia, and R. Ernst
-  [PDF icon](#) Compact Binaries with Code Compression in a Software Dynamic Translator [p. 1052]
S. Shogan and B. Childers

8C: [Real-Life Defect Modelling and Detection](#)

Moderators: *R. Aitken, Artisan, US; H. Manhaeve, Q-star Test, BE*

-  [PDF icon](#) Pattern Selection for Testing of Deep Sub-Micron Timing Defects [p. 1060]
M. C. Chao, L. Wang, and K. Cheng
-  [PDF icon](#) Balanced Excitation and its Effect on the Fortuitous Detection of Dynamic Defects [p. 1066]
J. Dvorak, B. Cobb, J. Wingfield, and M. Mercer
-  [PDF icon](#) Intermittent Scan Chain Fault Diagnosis Based on Signal Probability Analysis [p. 1072]
Y. Huang, W. Cheng, C. Hsieh, H. Tseng, A. Huang, and Y. Hung
-  [PDF icon](#) A Modeling Approach for Addressing Power Supply Switching Noise Related Failures of Integrated Circuits [p. 1078]
C. Tirumurti, S. Kundu, S. Sur-Kolay, and Y. Chang
-  [PDF icon](#) Soft Faults and the Importance of Stresses in Memory Testing [p. 1084]
Z. Al-Ars and A. van De Goor

8E: [Optimisation in Physical Design](#)

Moderators: *J. Lienig, TU Drasden, DE; R. Otten, TU Eindhoven, NL*

-  [PDF icon](#) Wire Retiming for System-On-Chip by Fixpoint Computation [p. 1092]
C. Lin and H. Zhou
-  [PDF icon](#) Boosting: Min-Cut Placement with Improved Signal Delay [p. 1098]
A. Kahng, S. Reda, and I. Markov
-  [PDF icon](#) Optimal Algorithm for Minimizing the Number of Twists in an On-Chip Bus [p. 1104]
L. Deng and M. Wong
-  [PDF icon](#) A Fast Word-Level Statistical Estimator of Intra-Bus Crosstalk [p. 1110]
S. Gupta and S. Katkooi
-  [PDF icon](#) Full-Chip Multilevel Routing for Power and Signal Integrity [p. 1116]
J. Xiong and L. He

8G: [Hot Topic -- Platforms and Tools for Energy-Efficient Design of Multimedia Systems](#)

Organizer/Moderator: *E. Pol, Philips Research, NL*

Speakers:

H. Van Antwerpen, Philips Research, NL
R. von Vignau, Philips Research, NL
R. Gupta, UC San Diego, US
N. Dutt, UC Irvine, US
N. Venkatasubramanian, UC Irvine, US
S. Mohapatra, UC California Irvine, US
C. Pereira, UC California San Diego, US

-  [PDF icon](#) Energy-Aware System Design for Wireless Multimedia [p. 1124]

9A: [Communication Design for MPSoC](#)

Moderators: *K. Goossens, Philips Research, NL; L. Benini, Bologna U, IT*

-  [PDF icon](#) Unified Component Integration Flow for Multi-Processor SoC Design and Validation [p. 1132]
M. Dzirri, W. Cesário, A. Jerraya, and F. Wagner
-  [PDF icon](#) An Interconnect Channel Design Methodology for High Performance Integrated Circuits [p. 1138]
V. Chandra, A. Xu, H. Schmit, and L. Pileggi
-  [PDF icon](#) Modeling Shared Resource Contention Using a Hybrid Simulation/Analytical Approach [p. 1144]
A. Bobrek, J. Pieper, J. Nelson, J. Paul, and D. Thomas
-  [PDF icon](#) Supporting Cache Coherence in Heterogeneous Multiprocessor Systems [p. 1150]
T. Suh, D. Blough, and H. Lee

9B: [Combining Static and Dynamic Software Optimisation](#)

Moderators: *R. Ernst, TU Braunschweig; P. Kalfasz, Thales Communications, FR*

-  [PDF icon](#) Exploiting Processor Workload Heterogeneity for Reducing Energy Consumption in Chip Multiprocessors [p. 1158]
I. Kadayif, M. Kandemir, and I. Kolcu
-  [PDF icon](#) Fault-Tolerant Deployment of Embedded Software for Cost-Sensitive Real-Time Feedback-Control Applications [p. 1164]
C. Pinello, L. Carloni, and A. Sangiovanni-Vincentelli
-  [PDF icon](#) Task Feasibility Analysis and Dynamic Voltage Scaling in Fault-Tolerant Real-Time Embedded Systems [p. 1170]
Y. Zhang and K. Chakrabarty
-  [PDF icon](#) Quasi-Static Scheduling for Real-Time Systems with Hard and Soft Tasks [p. 1176]
L. Cortés, P. Eles, and Z. Peng

9C: [Hot Topic -- The Status of the New IEEE Test Standards](#)

Organizer/Moderator: *B. Bennetts, Bennetts Associates, UK*

-  [PDF icon](#) Status of IEEE Testability Standards 1149.4, 1532 and 1149.6 [p. 1184]
S. Sunter, A. Osseiran, A. Cron, N. Jacobson, D. Bonnett, B. Eklov, C. Barnhart, and B. Bennetts

9E: [Modelling and Estimation in Circuit Layout](#)

Moderators: *I. Markov, Michigan U, US; J. Lienig, TU Dresden, DE*

-  [PDF icon](#) Eliminating False Positives in Crosstalk Noise Analysis [p. 1192]
Y. Ran, M. Marek-Sadowska, A. Kondratyev, and Y. Watanabe
-  [PDF icon](#) A New Approach to Timing Analysis Using Event Propagation and Temporal Logic [p. 1198]
A. Mondal, P. Chakrabarti, and C. Mandal
-  [PDF icon](#) A New Effective Congestion Model in Floorplan Design [p. 1204]
Y. Hsieh and T. Hsieh
-  [PDF icon](#) ULSI Interconnect Length Distribution Model Considering Core Utilization [p. 1210]
H. Nakashima, J. Inoue, K. Okada, and K. Masu

9G: Applications of Reconfigurability

Moderators: Y. Tanurhan, Actel, US; W. Rosenstiel, Tuebingen U and FZI Karlsruhe, DE

 [PDF icon](#) Implementation of a UMTS Turbo-decoder on a Dynamically Reconfigurable Platform [p. 1218]
A. La Rosa, C. Passerone, F. Gregoretti, and L. Lavagno

 [PDF icon](#) Design Methodology for a Tightly Coupled VLIW/Reconfigurable Matrix Architecture: A Case Study [p. 1224]
B. Mei, R. Lauwereins, S. Vernalde, and D. Verkest

 [PDF icon](#) Efficient Implementations of Mobile Video Computations on Domain-Specific Reconfigurable Arrays [p. 1230]
I. Ahmed, S. Baloch, A. Pai, T. Arslan, N. Aydin, S. Khawam, and F. Westall

 [PDF icon](#) Mapping Multi-Million Gate SoCs on FPGAs: Industrial Methodology and Experience [p. 1236]
H. Krupnova

10A: Interconnect Modelling for MPSoC

Moderators: B. Candaele, Thales, FR; A. Jerraya, ITMA Laboratory, FR

 [PDF icon](#) Using a Communication Architecture Specification in an Application-Driven Retargetable Prototyping Platform for Multiprocessing [p. 1244]
X. Zhu and S. Malik

 [PDF icon](#) A Power and Performance Model for Network-on-Chip Architectures [p. 1250]
N. Banerjee, P. Vellanki, and K. Chatha

 [PDF icon](#) A System Level Processor/Communication Co-Exploration Methodology for Multi-Processor System-on-Chip Platforms [p. 1256]
A. Wiefierink, T. Kogel, R. Leupers, G. Ascheid, H. Meyer, G. Braun, and A. Nohl

10B: Embedded Software Generation and Optimisation

Moderators: F. Rousseau, ITMA Laboratory, FR; J. Madsen, TU Denmark, DK

 [PDF icon](#) Cache-Aware Scratchpad Allocation Algorithm [p. 1264]
M. Verma, L. Wehmeyer, and P. Marwedel

 [PDF icon](#) Phase Coupled Code Generation for DSPs Using a Genetic Algorithm [p. 1270]
M. Lorenz and P. Marwedel

 [PDF icon](#) A Methodology and Tool Suite for C Compiler Generation from ADL Processor Models [p. 1276]
M. Hohenauer, H. Scharwaechter, K. Karuri, O. Wahlen, T. Kogel, R. Leupers, G. Ascheid, H. Meyer, G. Braun, and H. van Someren

10C: Scan-Based Testing

Moderators: H. Vranken, Philips Research, NL; C. Papachristou, Case Western Reserve U, US

 [PDF icon](#) Nine-Coded Compression Technique with Application to Reduced Pin-Count Testing and Flexible On-Chip Decompression [p. 1284]
M. Tehranipour, M. Nourani, and K. Chakrabarty

 [PDF icon](#) CircularScan: A Scan Architecture for Test Cost Reduction [p. 1290]
B. Arslan and A. Orailoglu

 [PDF icon](#) Hybrid Delay Scan: A Low Hardware Overhead Scan-Based Delay Test Technique for High Fault Coverage and Compact Test Sets [p. 1296]
S. Wang, S. Chakradhar, and X. Liu

 [PDF icon](#) Diagnosis of Scan-Chains by Use of a Configurable Signature Register and Error-Correcting Codes [p. 1302]
A. Leininger, P. Muhmenthaler, and M. Goessel

M. Pedram, Southern California U, US

 [PDF icon](#) Distributed Multimedia System Design: A Holistic Perspective [p. 1342]

IP4: Interactive Presentations

 [PDF icon](#) Adaptive Prefetching for Multimedia Applications in Embedded Systems [p. 1350]
H. Shetty, S. Niar, and L. Eeckhout

 [PDF icon](#) Data Windows: A Data-Centric Approach for Query Execution in Memory-Resident Databases [p. 1352]
J. Pisharath, A. Choudhary, and M. Kandemir

 [PDF icon](#) High-Performance QuIDD-Based Simulation of Quantum Circuits [p. 1354]
G. Viamontes, I. Markov, and J. Hayes

 [PDF icon](#) An Application of Parallel Discrete Event Simulation Algorithms to Mixed Domain System Simulation [p. 1356]
D. Reed, S. Levitan, J. Boles, J. Martinez, and D. Chiarulli

 [PDF icon](#) Fault Tolerance of Programmable Switch Blocks [p. 1358]
J. Huang, M. Tahoori, and F. Lombardi

 [PDF icon](#) A New Self-Checking Sum-Bit Duplicated Carry-Select Adder [p. 1360]
E. Sogomonyan, D. Marienfeld, V. Ocheretnij, and M. Gössel

 [PDF icon](#) A Macromodelling Methodology for Efficient High-Level Simulation of Substrate Noise Generation [p. 1362]
L. Elvira, F. Martorell, X. Aragonés, and J. González

 [PDF icon](#) Accurate Estimation of Parasitic Capacitances in Analog Circuits [p. 1364]
A. Agarwal, H. Sampath, V. Yelamanchili, and R. Vemuri

 [PDF icon](#) GRAAL – A Development Framework for Embedded Graphics Accelerators [p. 1366]
D. Crisu, S. Cotaş, S. Vassiliadis, and P. Liuha

 [PDF icon](#) From Synchronous to Asynchronous: An Automatic Approach [p. 1368]
J. Cortadella, A. Kondratyev, L. Lavagno, K. Lwin, and C. Sotiropoulos

 [PDF icon](#) Enhancing Testability of System on Chips Using Network Management Protocols [p. 1370]
O. Laouamri and C. Aktouf

 [PDF icon](#) Minimization of Crosstalk Noise, Delay and Power Using a Modified Bus Invert Technique [p. 1372]
M. Lampropoulos, B. Al-Hashimi, and P. Rostinger

 [PDF icon](#) Energy-Efficient Design for Highly Associative Instruction Caches in Next-Generation Embedded Processors [p. 1374]
J. Aragon, D. Nicolaescu, A. Veldenbaum, and A. Badulescu

 [PDF icon](#) Dynamic Voltage and Cache Reconfiguration for Low Power [p. 1376]
A. Nacu and A. Givargis

IP5: Interactive Presentations

- [!\[\]\(f9c8920940b4337cbbeec527f8ac9321_img.jpg\) PDF icon](#) Overhead-free Polymorphism in Network-on-Chip Implementation of Object-Oriented Models [p. 1380]
M. Goudarzi, S. Hessabi, and A. .Meyroft
- [!\[\]\(845e9749bae62df6369d34b1189931bf_img.jpg\) PDF icon](#) Multi-Processor SoC Design Methodology Using a Concept of Two-Layer Hardware-Dependent Software [p. 1382]
S. Yoo, M. Youssef, A. Bouchhima, A. Jerraya, and M. Diaz-Nava
- [!\[\]\(938d0015fba187a43362a21bd1933ad2_img.jpg\) PDF icon](#) Synthesis of Reversible Logic [p. 1384]
A. Agrawal and N. Jha
- [!\[\]\(05617ec18e9328b0f1a36e60ffcf2bcf_img.jpg\) PDF icon](#) A Unified Design Space for Regular Parallel Prefix Adders [p. 1386]
M. Ziegler and M. Stan
- [!\[\]\(e26d5d357d1bae1ed2c584ddce77119b_img.jpg\) PDF icon](#) MODD: A New Decision Diagram and Representation for Multiple Output Binary Functions [p. 1388]
A. Jabir and D. Pradhan
- [!\[\]\(05408b9c1c15ead4463a073c4f1f227f_img.jpg\) PDF icon](#) Issues in Implementing Latency Insensitive Protocols [p. 1390]
M. Casu and L. Macchiarulo
- [!\[\]\(75379a268b2c36a9e181d363de73226b_img.jpg\) PDF icon](#) Model-Based Specification and Execution of Embedded Real-Time Systems [p. 1392]
T. Schattkowsky and W. Mheller
- [!\[\]\(1c56f8316f4230d3b8983f61403d9b3c_img.jpg\) PDF icon](#) A Demonstration of Co-Design and Co-Verification in a Synchronous Language [p. 1394]
S. Singh
- [!\[\]\(21fc661b153d496d682392368cc1e65e_img.jpg\) PDF icon](#) Profile Guided Management of Code Partitions for Embedded Systems [p. 1396]
S. Zhou, B. Childers, and N. Kumar

IP6: Interactive Presentations

- [!\[\]\(73937a8f461f9a87bc32687d38d0d956_img.jpg\) PDF icon](#) Realizable Reduction for Electromagnetically Coupled RLMC Interconnects [p. 1400]
R. Jiang and C. Chen
- [!\[\]\(da858e3d6103141af517539e8660fcef_img.jpg\) PDF icon](#) Statistically Aware Buffer Planning [p. 1402]
G. Garcea, N. van der Meijis, and R. Otten
- [!\[\]\(267484f6664d6c3f0c4478b0438e85ff_img.jpg\) PDF icon](#) A Tunneling Model for Gate Oxide Failure in Deep Sub-Micron Technology [p. 1404]
S. Bernadini, J. Portal, and P. Masson
- [!\[\]\(c70f03af5fe29a85b0bb6fde7f153b85_img.jpg\) PDF icon](#) Power Supply Noise Monitor for Signal Integrity Faults [p. 1406]
J. Vázquez and J. de Gyvez
- [!\[\]\(440c728b7cb4f92432de1960e4ca4b7a_img.jpg\) PDF icon](#) Testing of Quantum Dot Cellular Automata Based Designs [p. 1408]
M. Tahoori and F. Lombardi
- [!\[\]\(191d38854a7a5458579942ab824c7428_img.jpg\) PDF icon](#) Net and Pin Distribution for 3D Package Global Routing [p. 1410]
J. Minz, M. Pathak, and S. Lim
- [!\[\]\(f9b65b52d931c2e5aacc5a7bc03d2445_img.jpg\) PDF icon](#) Placement Using a Localization Probability Model (LPM) [p. 1412]
M. Olbrich and E. Barke
- [!\[\]\(bac38977ca5403e65246607a19fbac9e_img.jpg\) PDF icon](#) CMOS Structures Suitable for Secured Hardware [p. 1414]
S. Guillevy, P. Hoogvorst, Y. Mathieu, R. Pacalet, and J. Provost
- [!\[\]\(6f3b477f998143ad0e57eeef366aa8ac_img.jpg\) PDF icon](#) Timing Correction and Optimization with Adaptive Delay Sequential Elements [p. 1416]
K. Rahimi, S. Bridges, and C. Diorio

SIGDA DATE 2004

https://past.date-conference.com/proceedings-archive/2004/YEAR.HTM

90%

DATE 19

MARCH 25-29, 2019

FLORENCE, ITALY

FIRENZE FIERA

DESIGN, AUTOMATION AND TEST IN EUROPE THE EUROPEAN EVENT FOR



MAIN PAGE

DATE 2004

Proceedings:

- Table of Content
- Abstracts
- Session Index
- Author Index

Designers' Forum:

- Table of Content
- Abstracts
- Session Index
- Author Index

Abstracts

Session Index

Author Index

Designers' Forum:

- Table of Content
- Abstracts
- Session Index
- Author Index

Proceedings:

- Table of Content
- Abstracts
- Session Index
- Author Index

Designers' Forum:

- Table of Content
- Abstracts
- Session Index
- Author Index

Foreword

On behalf of the Executive and Technical Program Committees, we take great pleasure in welcoming you to the proceedings of **DATE 04**, the most comprehensive European Conference and Exhibition event bringing together academic researchers, industry specialists, users and vendors in the fields of Design, Automation and Test of electronic circuits and systems.

The endless quest for faster, less consuming, cheaper and safer electronic products, particularly for the growing consumer and communications markets, dictate that increasingly complex designs be generated in continuously shrinking time scales. Design automation is a strategic technology for modern electronic systems, from simple ASICs via embedded IP cores to large systems on chip made of heterogeneous processors communicating via on-chip networks. The use of embedded processors in SoCs requires an ever-increasing emphasis on embedded software development techniques as part of the design process of state-of-the-art electronic systems. Testing of these complex electronic systems is becoming a key factor in enabling the final quality goals and has a growing impact in the overall cost of product development.

DATE has become the pre-eminent European conference addressing research and development activities in the field of design technology. World leading companies in multimedia, wireless communications, and automotive industries have led to the emergence of recognised European strengths in System-on-Chip Technologies, particularly Embedded Systems with a strong analogue and mixed-signal content. **DATE** is a unique event where the latest scientific and technological developments are brought to the public from the leading players in the field. **DATE**, therefore, offers a once-in-a-year occasion to find out more about the technical and architectural trends of circuits and systems, small and large, and the design automation and test challenges they bring to users, tool providers and PCB companies as well as semiconductor and IP providers alike.

A total of 780 submissions were received for **DATE 04**, a significant increase over the previous edition. After careful consideration, the Program Committee has selected 181 papers for regular oral presentation in 53 sessions — an acceptance rate of 27 percent — and 73 papers as interactive presentations (a new format started at DATE 04 to present novel ideas of "ongoing work" of high quality in an interactive one-to-one format). An additional 52 regular papers and 5 interactive presentations are presented in the Designers' Forum, a technical program track devoted to the presentation and discussion of practical, industry-oriented design experiences and methodologies using state of the art design tools. In addition, 12 Special Sessions (embedded tutorials, panels, etc.) have been chosen to bring stimulating insights and vibrant discussions in state-of-the-art topics. In complement to the main technical and scientific program, there are a number of valuable education activities in the days before and after the conference. These are the Pre-Conference Tutorials and Industrial Tutorials on Monday and 3 full-day Master Courses and 4 Workshops on Friday. The organisation of a PCB Symposium, a University Booth, fringe meetings and social events during the conference offer a wide variety of extra opportunities to meet and exchange information on relevant issues for the design, automation and test communities.

DATE continues its tradition of grouping presentations and activities related to Special Themes on specific focus days. This year's Special Days are dedicated to **ULTRA LOW POWER** (Wednesday) and **TRENDS IN MULTIMEDIA** (Thursday). In addition to the keynote address presented at the Tuesday Opening Plenary this year we will have two keynotes related to these Special Days. The Wednesday lunch plenary keynote is devoted to Ultra Low Power and the Thursday lunch keynote to Trends in Multimedia.

The one-day Executive Track, comprising of three sessions and one executive panel, is one of the new initiatives in **DATE 04** designed to offer the technical/business views of high-level executives from EDA and related fields. Presentations by leading company CEOs, CTOs, VPs and other senior executives from EDA vendors, IP suppliers, semiconductor and system houses and test technology suppliers takes place on Tuesday.

Also as a new initiative, **DATE 04** offers Workshops on Friday as a complement to the regular conference running throughout the week. Four workshops will run in parallel showing research results, keynotes on advanced technologies and discussion panels of a different nature. The four Workshops will deal with aspects and technologies not directly covered by the technical program.

xxvii

Following its successful implementation last year, the **DATE Focus on Business and Industry**, a presentation theatre located on the exhibition floor, offers exhibition visitors and conference delegates a fresh view on some key technical and strategic issues in the current electronic systems design market.

Many volunteers have given their best efforts to make this Conference and Exhibition an outstanding event. We would like to thank all the members of the Sponsors Committee, the Executive Committee, the Technical Program Committee, the Vendors Committee and the Exhibitors, and specially our recognition goes to all the Authors, Speakers, Session Organisers, Session Chairs, and Reviewers for their continued interest, energy and support to **DATE**. We hope that you enjoy the **DATE 04** Conference and Exhibition as much as all of us have enjoyed organising it for you, and that you find much scientific value in reading these proceedings.

Joan Figueras

General Chairman



Georges Gielen

Program Chairman



Page 12

Qualcomm Exhibit 1065

Qualcomm v. Network System Technologies - IPR2024-00343

Andrei Rădulescu, John Diehlissen, Kees Goossens, Edwin Rijpkema, and Paul Wielage
Philips Research Laboratories, Eindhoven, The Netherlands

In this paper we present a network interface for an on-chip network. Our network interface decouples computation from communication by offering a shared-memory abstraction, which is independent of the network implementation. We use transaction-based protocols to achieve hardware compatibility with existing bus protocols such as AXI, OCP and DTI. Our network interface has a modular architecture, which allows flexible instantiation. It provides both guaranteed and best-effort services via connections. These are configured via network interface ports using the network itself instead of a separate control interconnect. An example instance of this network interface with 4 ports has an area of 0.143mm² in a 0.13µm technology, and runs at 500 MHz.

Networks on chip (NoC) have been proposed as a solution to the interconnect problem for highly complex chips [2, 3, 5, 9, 12, 14, 15, 17, 21, 27]. NoCs help designing chips in several ways: they (a) structure and manage wires in deep submicron technologies [2, 3, 9, 12, 21], (b) allow good wire utilization through sharing [5, 9, 12, 21], (c) scale better than buses [14, 21], (d) can be energy efficient and reliable [2, 5], and (e) decouple computation from communication through well-defined interfaces, enabling IP modules and interconnect to be designed in isolation, and to be integrated more easily [2, 13, 21, 24].

Networks are composed of *routers*, which transport the data from one place to another, and *network interfaces* (NI), which implement the interface to the IP modules. In a previous article [21], we have shown the trade-offs in designing a cost-effective router combining guaranteed with best-effort traffic. In this paper, we focus on the other network component, the network interface.

Network interface design has received considerable attention for parallel computers [8, 25], and computer networks [6, 7]. These designs are optimized for performance (high throughput, low latency), and often consist of a dedicated processor, and large amount of buffering. As a consequence, their cost is too large to be applicable on chip.

On-chip network interfaces must provide a low-area overhead, because the size of IP modules attached to the NoC is relatively small. Designs of network interfaces with a low area have been proposed [4, 28]. However, they do not provide throughput or latency guarantees, which are essential for a compositional construction of complex SoCs.

Our NI is intended for systems on chip (SoC), hence, it must have a low area. To enable the reuse of existing IP modules, we must provide a smooth transition from buses to NoCs. A shared-

memory abstraction via transactions (e.g., read, write) ensures this. Further, we also have to provide a simple and flexible configuration, preferably using the NoC itself to avoid the need for a separate scalable interconnect.

We achieve a low-cost implementation of the NI by implementing the protocol stack in hardware, and by exploiting on-chip characteristics (such as the absence of transmission errors, relatively static configuration, tight synchronization) to implement only the relevant parts of a complete OSI stack. A hardware implementation of the protocol stack provides a much lower latency overhead compared to a software implementation. Further, a hardware implementation allows both hardware and software cores to be reused without change. (1)

Our NI provides services at the transport layer in the ISO-OSI reference model [22], because this is the first layer where offered services are independent of the network implementation. This is a key ingredient in achieving the *decoupling between computation and communication* [16, 24], which allows IP modules and interconnect to be designed independently from each other. We provide transport-layer services by defining connections (e.g. point-to-point or multicast) configured for specific properties (e.g.

We offer *guaranteed services* as they are essential for a compositional construction (design and programming) of SoC. The reasons are that they limit the possible interactions of IPs with the communication environment [12, 13], separate the IP requirements and their implementation, and make application quality of service independent of the IP and NoC implementations. Examples of such guarantees are lower bounds on throughput, and upper bounds on latency.

Our NoC, called *Ætheral*, offers a *shared-memory abstraction* to the IP modules. Communication is performed using a transaction-based protocol, where master IP modules issue *request messages* (e.g., read and write commands at an address, possibly carrying data) that are executed by the addressed slave modules, which may respond with a *response message* (i.e., status of the command execution, and possibly data) [23]. We adopt this protocol to provide backward compatibility to existing on-chip communication protocols (e.g., AXI [1], OCP [18], DTL [19]), and also to allow future protocols better suited to NoCs.

We provide a modular NI, which can be configured at design time. This is, the number of ports and their type (i.e., configuration port, master port, or slave port), the number of connections at each port, memory allocated for the queues, the level of services per port, and the interface to the IP modules are all configurable at design (instantiation) time using an XML description [11].

The NI allows flexible NoC configuration at run time. Each connection can be configured individually, requiring configurable NoC components (i.e., router and NI). However, instead of using

a separate control interconnect to program them, the NoC is used to program itself. This is performed through configuration ports using DTL-MMIO (memory-mapped IO) transactions [19]. The NoC can be configured in a distributed fashion (i.e., via multiple configuration ports), or centralized (i.e., via a single port).

The paper is organized as follows. In the first section, the background on NoCs and the design goals of the NoC are presented. In the second section, the NoC architecture is described, and the interface offered to the IP modules are described. In Section 3, we show that NoCs can be configured both in a distributed and in a centralized way, and we present the trade-offs between the two approaches. In Section 4, we present a modular network interface architecture, which is split into a kernel, providing core functionality, and a number of shells to extend functionality, e.g., wrappers to provide an interface to existing bus protocols, such as AXI or DTI. In this section, we also show how the NI allows NoC configuration using the NoC itself as opposed to via a separate control interconnect. In Section 5, we demonstrate the feasibility of our network interface design through a prototype implementation in a 0.13μm technology, and we conclude in Section 6.

As mentioned in the previous section, the communication services of the $\mathcal{A}$ etheral NoC are defined to meet the following goals: (a) decouple computation (IP modules) from communication (NoC), (b) provide backward compatibility to existing bus protocols, (c) provide support for real-time communication, and (d) have a low-cost implementation.

Decoupling computation from communication is a key ingredient in managing the complexity of designing chips with billions of transistors, because it allows the IP modules and the interconnect to be designed independently [16, 24]. In NoCs, this decoupling is achieved by positioning the network interfaces at the transport level [3, 21] or above in the ISO-OSI reference model [22]. At the transport level, the offered services are *end to end* between communicating IP modules, hiding, thus, the network internals, such

Backward compatibility with existing protocols, such as AXI or DTL, is achieved by using a model based on *transactions* [23]. In a transaction-based model, there are two types of IP modules: *masters* and *slaves*. Masters initiate transactions by issuing requests, which can be further split in commands, and write data (corresponding to the address and write signal groups in AXI). Examples of commands are read, and write. One or more slaves receive and execute each transaction. Optionally, a transaction can also include a response issued by the slave to the master to return data or an acknowledgment of the transaction execution (corresponding to the read data and valid signal groups in AXI).

In the Ethernet NoC, all these signals are sequentialized in request and response messages, which are supplied to the NoC, where they are transported by means of packets. Sequentialization is performed to reduce the number of wires, increasing their utilization, and to simplify arbitration. Packetization is performed by the NI and is thus transparent to the IP modules.

The Ethernet NoC offers its services on *connections*, which can be point to point (one master, one slave), multicast (one master, multiple slaves, all slaves executing each transaction), and narrowcast (one master, multiple slaves, a transaction is executed by only one slave) [23]. Connections are composed of unidirectional point-to-point *channels* (between a single master and a single slave). To each channel, *properties* are attached, such as guaranteed message delivery or not, in order or un-ordered message

delivery, and with or without timing guarantees. As a result, different properties can be attached to the request and response parts of a connection, or for different slaves within the same connection. Connections can be opened and closed at any time. Opening and closing of connections takes time, and is intended to be performed at a granularity larger than individual transactions.

Support for real-time communication is achieved by providing throughput, latency and jitter *guarantees*. In *Ethernet*, this is implemented by configuring connections as pipelined time-division-multiplexed circuits over the network. Time multiplexing is only possible when the network routers have a notion of synchronicity which allows slots to be reserved consecutively in a sequence of routers [13, 21]. This scheme [21] has smaller packet buffers, and, hence, has lower implementation cost compared to alternatives, such as rate-based packet switching [29], or deadline-based packet switching [20].

Throughput guarantees are given by the number of slots reserved for a connection. Slots correspond to a given bandwidth: H_k , and, therefore, reserving N slots for a connection results in a total bandwidth of $N \times H_k$. The latency bound is given by the waiting time until the reserved slot arrives and the number of routers data passes to reach its destination. Jitter is given by the maximum distance between two slot reservations.

Protocol stacks that are used in networks to implement communication services, require additional cost compared to buses. Protocol stacks are necessary in networks to manage the complexity of networks, and to offer differentiated services. The pressure to keep the protocol stack small is higher on-chip than off-chip, because the size of the IP modules attached to the NoC can be relatively small. However, for NoCs, the protocol stacks can be reduced by exploiting the on-chip characteristics (e.g., no transfer errors, short wires) [23]. In the ϵ therNet NoC, we optimize the performance and minimize the cost of the protocol stack by implementing it in hardware, rather than in software. We support this claim in Section 5.

Before the .Aetheral NoC can be used by an application, it must be configured. NoC (re)configuration means opening and closing connections in the system. Connections are set up depending on the application or the mode the system is running. Therefore, we must be able to open and close connections while the system is running. (Re)configuration can be partial or total (some or all connections are opened/closed, respectively).

Opening a connection involves setting several registers, and allocating shared resources (for more details see Section 4). In the case of the current prototype of the *EtherNet* NoC, for each pair of one master and one slave of a connection, there are 5 and 3 registers written at the master and slave network interfaces, respectively. The shared resources consist of the slots allocated to the connections. These slots can be configured using either a distributed or a centralized model.

In the distributed case, a connection can be opened/closed from multiple network interface ports. Multiple configuration operations can be performed simultaneously, however, potential conflicts must also be solved (e.g., connection configurations initiated at two configuration ports may try to reserve the same slot in a router). Information about the slots is maintained in the routers, which also accept or reject a tentative slot allocation.

In a centralized system, there is only one place that performs NoC configuration. In such a case, the slot information can be




Figure 1. NI kernel and shells

stored in the configuration module instead of the routers, which simplifies the design, and, in the case of small NoCs, may even speed up configuration. For large NoCs, however, centralized configuration can introduce a bottleneck.

In the initial prototype of the Æthereal NoC, we opt for centralized configuration, because it is able to satisfy the needs of a small NoC (around 10 routers), and has a simpler design and lower cost. We use transactions to program the NoC, both for connection registers in the NIs, and for the slot information. We present details of how NoC configuration is performed in Section 4.

4 Network Interface Architecture

The network interface (NI) is the component that provides the conversion of the packet-based communication of the NoC to the higher-level protocol that IP modules use. We split the design of the network interface in two parts (see Figure 1): (a) the *NI kernel*, which implements the channels, packetizes messages and schedules them to the routers, implements the end-to-end flow control, and the clock domain crossing, and (b) the *NI shells*, which implement the connections (e.g., narrowcast, multicast), transaction ordering for connections, and other higher-level issues specific to the protocol offered to the IP.

4.1 NI Kernel Architecture

The NI kernel (see Figure 2) receives and provides messages, which contain the data provided by the IP modules via their protocol after sequentialization. The message structure may vary depending on the protocol used by the IP module. However, the message structure is irrelevant for the NI kernel, as it just sees messages as pieces of data to be transported over the NoC.

The NI kernel communicates with the NI shells via *ports*. At each port, point-to-point connections can be configured, their maximum number being selected at NI instantiation time. A port can have multiple connections to allow differentiated traffic classes, in which case there are also *connid* signals to select on which connection a message is supplied or consumed.

Inside the NI kernel, there are two message queues for each point-to-point connection (one source queue, for messages going to the NoC, and one destination queue, for messages coming from the NoC). Their size is also selected at the NI instantiation time. In our NI, queues are implemented using custom-made hardware fifos.




Figure 2. Network interface kernel

and are also used to provide the clock domain crossing between the network and the IP modules. Each port can, therefore, have a different clock frequency.

Each channel is configured individually. In a first prototype of the Æthereal NI, we can configure if a channel provides time guarantees (GT) or not (we call this best effort, BE), reserve slots for GT connections, configure the end-to-end flow control, and the routing information.

End-to-end flow control ensures that no data is sent unless there is enough space in the destination buffer to accommodate it. This is implemented using credits [26]. For each channel, there is a counter (*Space*) tracking the empty buffer space of the remote destination queue. This counter is initialized with the remote buffer size. When data is sent from the source queue, the counter is decremented. When data is consumed by the IP module at the other side, credits are produced in a counter (*Credits*) to indicate that more empty space is available. These credits are sent to the producer of data to be added to its *Space* counter. In the Æthereal prototype, we piggyback credits in the header of the packets for the data in the other direction to improve NoC efficiency. Note that at most *Space* data items can be transmitted before credits are received. We call the minimum between the data items in the queue and the value in the counter *Space*, the *sendable data*.

From the source queues, data is packetized (*Pck*) and sent to the NoC via a single link. A packet header contains the routing information (NI address for destination routing, and path for source routing), remote queue id (i.e., the queue of the remote NI in which the data will be stored), and piggybacked credits.

There are multiple channels which may require data transmission, we implement a scheduler to arbitrate between them. The scheduler checks if the current slot is reserved for a GT channel. If the slot is reserved, is the GT channel has data which can be transmitted, and if there is space in the channel's destination buffer, then the channel is granted data transmission. Otherwise,

3

the scheduler selects a BE channel with data and remote space using some arbitration scheme: e.g. round-robin, weighted round-robin, or based on the queue filling.

To optimize the NoC utilization, it is preferable to send longer packets. To achieve this, we implemented a configurable threshold mechanism, which skips a channel as long as the sendable data is below the threshold. This is applicable for both BE and GT channels. To prevent starvation at user application level (e.g., due to write data being buffered indefinitely on which the IP module waits for an acknowledgment), we also provide a *flush* signal for each channel (and a bit in the message header) to temporarily override the threshold. When the flush signal is high for a cycle, a snapshot of its source queue filling is taken, and as long as all the words in the queue at the time of flushing have not been sent, the threshold for that queue is bypassed.

A similar threshold is set for credit transmission. The reason is that, when there is no data on which the credits can be piggybacked, the credits are sent as empty packets, thus, consuming extra bandwidth. To minimize the bandwidth consumed by credits, a credit threshold is set, which allows credits to be transmitted only when their sum is above the threshold. Similarly to the data case, to prevent possible starvation, we provide a flush signal to force credits to be sent even when they are below their threshold.

As credits are piggybacked on packets, a queue becomes eligible for scheduling when either the amount of sendable data is above a first threshold, or when the amount of credits is above a second threshold. However, once a queue is selected, a packet containing the largest possible amount of credits and data will be produced. Note the amount of credits is bound by implementation to the given number of bits in the packet header, and packet have a maximum length to avoid links being used exclusively by a packet/channel, which would cause congestion.

On the outgoing path, packets are depacketized, credits are added to the counter *Space*, and data is stored in the corresponding queue, which is given by a queue id field in the header.

4

For more complex types of connections, such as narrowcast or multicast, and to provide conversions to other protocols, we add shells around the NI kernel. As an example, in Figure 1, we show a NI with two DTL and two AXI ports. All ports provide point-to-point connections. In addition to this, the two DTL ports provide narrowcast connections, and one DTL and one AXI port provide multicast connections. Note that these shells add specific functionality, and can be plugged in or left out at design time according to the requirements. NoC instantiation is simple, as we use an XML description to automatically generate the VHDL code for the NIs as well as for the NoC topology.

In Figure 3, we show an example of a narrowcast shell. Narrowcast connections are connections between one master and several slaves, where each transaction is executed by a single slave selected based on the address provided in the transaction [23]. Narrowcast connections provide a simple, low-cost solution for a single shared address space mapped on multiple memories. It implements the splitting/merging of data going to/coming from these memories.

We implement the narrowcast connection as a collection of point-to-point connections, one for each master-slave pair. Within a narrowcast connection, the slave for which the transaction is destined is selected based on the address (*Conn* block). The address range assigned to a slave is configurable in the narrowcast module. To provide in-order response delivery, the narrowcast must also keep a history of connection identifiers of the transactions including responses (e.g., reads, and acknowledged writes), and the length of these responses. In-order delivery per slave of request messages is already provided by the point-to-point connections.

When a slave using a connectionless protocol (e.g., DTL) is connected to a NI port supporting multiple connections, a multi-connection shell must be included to arbitrate between the connections. A multi-connection shell (see Figure 4) includes a scheduler to select connections from which messages are consumed, based e.g., on their filling. As for the narrowcast, the multi-connection shell has a connection id history for scheduling the responses.




Figure 8. NI configuration

In Figures 5 and 6, we show a master and slave shells that implement a simplified version of a protocol such as AXI. The basic functionality of such a shell is to serialize commands and their flags, addresses, and write data in request messages, and to de-serialize messages into read data, and write responses. Examples of the message structures (i.e., after serialization) passing from NI shells and NI kernel are shown in Figure 7. In full-fledged master and slave shells, more blocks would be added to implement e.g., the unbuffered writes at the master side, and read linked, write conditional at the slave side.

4.3 NI Configuration

As mentioned in Section 3, in our prototype Ethernet NoC, we opt for centralized configuration. This means that there is a single configuration module that configures the whole NoC, and that slot tables can be removed from the routers. Consequently, only the NIs need to be configured when opening/closing connections.

NIs are configured via a configuration port (CNIP), which offers a memory-mapped view on all control registers in the NIs. This means that the registers in the NI are readable and writable by any master using normal read and write transactions.

Configuration is performed using the NoC itself (i.e., there is no separate control interconnect needed for NoC configuration). Consequently, the CNIPs are connected to the NoC like any other slave (see CNIP at NI2 in Figure 8). At the configuration module Cfg's NI, we introduce a configuration shell (Config Shell), which, based on the address, configures the local NI (NI1), or sends configuration messages via the NoC to other NIs. The configuration shell optimizes away the need for an extra data port at NI1 to be connected to the NI1's CNIP.

In Figure 9, we show the necessary steps in setting up a connection between two modules (master B and slave A) from a configuration module (Cfg). Like for any other memory-mapped register, before sending configuration messages for configuring the B to A connection at NI2, a connection from Cfg to NI2's CNIP must be set up. This connection is opened in two steps corresponding to the request and response channels. First, the request channel to the NI2's CNIP is set up by writing the necessary registers in NI1 (Step 1 in Figure 9). Second, we use this channel to set up (via the NoC) the response channel from NI2's CNIP to Cfg (Step 2). The three shown messages are delivered and executed in order at NI2. The last of them sets requests an acknowledgment message to confirm that the channel has been successfully set up.

After the configuration connection has been set up, the remote NI2 can be configured to set up a channel from B to A. For config-




Figure 9. Connection configuration example

uring NI2 (B's NI), the previously set up configuration connection is used. For configuring NI1, the NI1's configuration port is accessed directly via Config Shell. First, the channel from the slave module A to the master module B is configured at NI1 (Step 3). Second, the channel from the master module B to the slave module A is configured (Step 4) through messages to NI2.

5 Implementation

In the previous section, we describe a prototype of a configurable NI architecture. In this section, we discuss the synthesized area and speed figures for the network interface components: NI kernel, narrowcast, multichannel and configuration shells, and master and slave shells for a simplified version of DTL.

We have synthesized an instance of a NI kernel with a STU of 8 slots, and 4 ports having 1, 1, 2, and 4 channels, respectively, with all queues being 32-bit wide and 3-word deep. The queues are area-efficient custom-made hardware fifos. We use these fifos instead of RAMs, because we need simultaneous access at all NI ports (possibly running at different speeds) as well as simultaneous read and write access for incoming and outgoing packets, which cannot be offered with a single RAM. Finally, for the small queues needed in the NI, multiple RAMs have a too large area overhead. Moreover, the hardware fifos implement the clock domain boundary allowing each NI port to run at a different clock frequency. The router side of the NI kernel runs at a frequency of 500 MHz, which matches our prototype router frequency [21], and delivers a bandwidth toward the router of 16 Gbit/s in each direction. The synthesized area for this NI kernel instance is 0.11 mm² in a 0.13μm technology.

Narrowcast and multi-connection shells have an area of 0.004 mm² and 0.007 mm², corresponding to 4% and 6% of the NI kernel area. The DTL shells are very small, 0.005 mm² and 0.002 mm² for the master and slave ports, corresponding to 5% and 2% of the NI kernel area, respectively. (This is also due to the fact that not all of the DTL functionality has been implemented.) The configuration shell, which provides a simplified DTL-MMIO interface to configure the NoC, has an area of 0.01 mm².

Summing up, for an example NI with 4 ports, one for configuration (one channel to which the configuration shell is attached), two masters (one offering narrowcast), and one slave (multichannel), the total area is 0.11 + 0.01 + 2 × 0.005 + 0.004 + 0.002 + 0.007 = 0.143 mm².

The latency introduced by our current NI is 2 cycles in the DTL master shell (due to serialization, as part of packetization), 0 to 2 in the narrowcast and multichannel shells (depending on the NI instance), and between 1 and 3 cycles in the NI kernels (as data needs to be aligned to a 3 word fit boundary), and 2 clock cycles for clock domain crossing. Additional delay is caused by the arbitration, but we do not include this in the NI latency overhead, as it needs to be performed anyway (also in the case of a bus, arbitration is performed).

The resulting latency overhead introduced by our NI is between 4 and 10 cycles, which is pipelined to maximize throughput. The latency overhead of a software implementation of the protocol is much larger (e.g., 47 instructions for packetization only [14]). A hardware implementation allows both legacy software and hardware task implementations to be used without change.

6 Conclusions

In this paper, we describe a network interface architecture which offers high-level services at a low cost. Our network interface provides a shared-memory abstraction, where communication is performed using read/write transactions. We offer, via connections, high-level services, such as transaction ordering, throughput and latency guarantees, and end-to-end flow control. These connections are configurable at runtime via a memory-mapped configuration port. We use the network to configure itself as opposed to using a separate control interconnect for network configuration.

Our network interface has a modular design, composed of kernel and shells. The NI kernel provides the basic functionality, including arbitration between channels, transaction ordering, end-to-end flow control, packetization, and a link protocol with the router. Shells implement (a) additional functionality, such as multichannel and narrowcast connections, and (b) adapters to existing protocols, such as AXI or DTL. All these shells can be plugged in or left out at design time according to the needs. This is done using an XML description of the network, which is used to automatically generate the VHDL code for the network interfaces, as well as for the network topology.

We show an instance of our network interface, which shows that the cost of implementing our protocol stack in hardware is small (0.143mm² in a 0.13μm technology, running at 500 MHz). Our hardware protocol stack implementation provides a very low protocol overhead of 4 to 6 cycles, which is much lower than a software stack implementation.

In conclusion, we provide an efficient network interface offering a shared-memory abstraction, high-level services (including guarantees), which allows runtime network configuration using the network itself.

References

- [1] ARM. *AMBA AXI Protocol Specification*, June 2003.
- [2] L. Benini and G. De Micheli. Powering networks on chips. In *Proc. ISSS*, 2001.
- [3] L. Benini and G. De Micheli. Networks on chips: A new SoC paradigm. *IEEE Computer*, 34(1):78–80, 2002.
- [4] P. Bhovani and R. Mahapatra. Interfacing cores with on-chip packet-switched networks. In *Proc. VLSI Design*, 2003.
- [5] E. Bolotin et al. QNoC: QoS architecture and design process for network on chip. *Journal of Systems Architecture*, 49, Dec. 2003.
- [6] T. Callahan and S. C. Goldstein. NFFDY: A low overhead, high throughput network interface. In *Proc. ISCA*, 1995.
- [7] A. Chien et al. Design challenges for high-performance network interfaces. *IEEE Computer*, 31(1):42–44, 1998.
- [8] D. J. Culler et al. *Parallel Computer Architecture: A Hardware/Software Approach*. Morgan Kaufmann Publishers, 1999.
- [9] W. J. Dally and B. Towles. Route packets, not wires: On-chip interconnection networks. In *Proc. DAC*, 2001.
- [10] O. P. Ganguly et al. Understanding video pixel processing applications for flexible implementations. In *Proc. EuroMicro DSD*, 2003.
- [11] S. Gonzalez-Petrucci et al. Cost-performance trade-offs in networks on chip: A simulation based approach. In *Proc. DATE*, 2004.
- [12] K. Goossens et al. Networks on silicon: Combining best-effort and guaranteed services. In *Proc. DATE*, 2002.
- [13] K. Goossens et al. Guaranteeing the quality of services in networks on chip. In J. Nuzni, H. Teuhonen, J. Isoaho, and A. Jantsch, editors, *Networks on Chip*, pages 61–82. Kluwer, 2003.
- [14] P. Guernier and A. Gresser. A generic architecture for on-chip packet-switched interconnections. In *Proc. DATE*, 2000.
- [15] F. Karim et al. An interconnect architecture for networking systems on chip. *IEEE Micro*, 22(5), 2002.
- [16] K. Krut'z et al. System-level design: Orthogonalization of concerns and platform-based design. *IEEE Trans. on CAD of Integrated Circuits and Systems*, 19(12):1523–1543, 2000.
- [17] S. Kumar et al. A network on chip architecture and design methodology. In *Proc. RVLSI*, 2002.
- [18] OCP International Partnership. *Open Core Protocol Specification 2.0 Release Candidate*, 2003.
- [19] Philips Semiconductors. *Device Transaction Level (DTL) Protocol Specification Version 2.2*, July 2002.
- [20] J. Rexford. *Tailoring Router Architectures to Performance Requirements in Cut-Through Networks*. PhD thesis, Univ. Michigan, 1999.
- [21] E. Rijpkema et al. Trade offs in the design of a router with both guaranteed and best-effort services for networks on chip. In *Proc. DATE*, 2003.
- [22] M. T. Rose. *The Open Book: A Practical Perspective on OSI*. Prentice Hall, 1990.
- [23] A. Radulescu and K. Goossens. Communication services for networks on chip. In S. Bhattacharyya, E. Depeytere, and J. Teich, editors, *Domain-Specific Embedded Multiprocessors*. Dekker, 2003.
- [24] M. Sgroi et al. Addressing the system-on-a-chip interconnect issues through communication-based design. In *Proc. DAC*, 2001.
- [25] P. Steenkiste. A high-speed network interface for distributed-memory systems: Architecture and applications. *ACM Trans. on Computer Systems*, 15(1):75–109, 1991.
- [26] A. S. Tienbaum. *Computer Networks*. Prentice Hall, 1996.
- [27] D. Wiklund and D. Liu. Switched network on chip for hard real time embedded systems. In *Proc. IPDPS*, 2003.
- [28] C. A. Zefrenko, M. E. Krut'z, L. Cerro, and A. A. Svan. A study on communication issues for system-on-chip. In *Proc. SRCI*, 2002.
- [29] H. Zhang. Service disciplines for guaranteed performance service in packet-switching networks. *Proc. of the IEEE*, 83(10):1374–1396, 1995.