

2017 IEEE International Conference on Computer Design (ICCD 2017)

**Boston, Massachusetts, USA
5-8 November 2017**



IEEE Catalog Number: CFP17ICD-POD
ISBN: 978-1-5386-2255-1

**Copyright © 2017 by the Institute of Electrical and Electronics Engineers, Inc.
All Rights Reserved**

Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limit of U.S. copyright law for private use of patrons those articles in this volume that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

For other copying, reprint or republication permission, write to IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08854. All rights reserved.

****** This is a print representation of what appears in the IEEE Digital Library. Some format issues inherent in the e-media version may also appear in this print version.***

IEEE Catalog Number:	CFP17ICD-POD
ISBN (Print-On-Demand):	978-1-5386-2255-1
ISBN (Online):	978-1-5386-2254-4
ISSN:	1063-6404

Additional Copies of This Publication Are Available From:

Curran Associates, Inc
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: (845) 758-0400
Fax: (845) 758-2633
E-mail: curran@proceedings.com
Web: www.proceedings.com

CURRAN ASSOCIATES INC.
proceedings
.com

2017 IEEE 35th International Conference on Computer Design

ICCD 2017

Table of Contents

Welcome Message	xvi
Organizing Committee	xvii
Program Committee	xviii
Keynote Talks	xxiii
Tutorials	xxvi

Best Papers Session

Estimating the Limits of CPU Power Management for Mobile Games	1
<i>Benedikt Dietich, Nadja Peters, Sangyoung Park, and Samarjit Chakraborty</i>	

Session 1A: Hardware Security I

BioChipWork: Reverse Engineering of Microfluidic Biochips	9
<i>Huili Chen, Seetal Potluri, and Farinaz Koushanfar</i>	
Subcomponent Timing-Based Detection of Malware in Embedded Systems	17
<i>Sixing Lu, Roman Lysecky, and Jerzy Rozenblit</i>	
Security Trade-Offs in Microfluidic Routing Fabrics	25
<i>Jack Tang, Mohamed Ibrahim, Krishnendu Chakrabarty, and Ramesh Karri</i>	
Side-Channel Attack on STTRAM Based Cache for Cryptographic Application	33
<i>Mohammad Nasim Imtiaz Khan, Shivam Bhasin, Alex Yuan, Anupam Chattopadhyay, and Swaroop Ghosh</i>	
Template Attack Based Deobfuscation of Integrated Circuits	41
<i>Abhishek Chakraborty, Yang Xie, and Ankur Srivastava</i>	
Neural Trojans	45
<i>Yuntao Liu, Yang Xie, and Ankur Srivastava</i>	

Session 1B: Read-Write Optimizations for Non-Volatile Memory

Adaptive Prefetching for Accelerating Read and Write in NVM-Based File Systems	49
<i>Shengan Zheng, Hong Mei, Linpeng Huang, Yanyan Shen, and Yanmin Zhu</i>	
A Cost-Efficient NVM-Based Journaling Scheme for File Systems	57
<i>Xiaoyi Zhang, Dan Feng, Yu Hua, and Jianxi Chen</i>	
TDV Cache: Organizing Off-Chip DRAM Cache of NVMM from a Fusion Perspective	65
<i>Tianyue Lu, Yuhang Liu, Haiyang Pan, and Mingyu Chen</i>	
RCTP: Region Correlated Temporal Prefetcher	73
<i>Dennis Antony Varkey, Biswabandan Panda, and Madhu Mutyam</i>	
A Shingle-Aware Persistent Cache Management Scheme for DM-SMR Disks	81
<i>Tianming Yang, Haitao Wu, Ping Huang, and Fei Zhang</i>	

Session 2A: Stochastic, Approximate, and Unary Computing

Low-Power and High-Speed Approximate Multiplier Design with a Tree Compressor	89
<i>Tongxin Yang, Tomoaki Ukezono, and Toshinori Sato</i>	
Neural Network Classifiers Using Stochastic Computing with a Hardware-Oriented Approximate Activation Function	97
<i>Bingzhe Li, Yaobin Qin, Bo Yuan, and David J. Lilja</i>	
Accurate and Efficient Stochastic Computing Hardware for Convolutional Neural Networks	105
<i>Joonsang Yu, Kyounghoon Kim, Jongeun Lee, and Kiyoung Choi</i>	
Approximate Reconfigurable Hardware Accelerator: Adapting the Micro-Architecture to Dynamic Workloads	113
<i>Siyuan Xu and Benjamin Carrion Schafer</i>	
Low Latency Approximate Adder for Highly Correlated Input Streams	121
<i>Xiaoliang Chen, Ahmed M. Eltawil, and Fadi J. Kurdahi</i>	
Power and Area Efficient Sorting Networks Using Unary Processing	125
<i>M. Hassan Najafi, David J. Lilja, Marc Riedel, and Kia Bazargan</i>	

Session 2B: Energy-Efficiency through Heterogeneity

Machine Learning-Based Approaches for Energy-Efficiency Prediction and Scheduling in Composite Cores Architectures	129
<i>Hossein Sayadi, Nisarg Patel, Avesta Sasan, and Houman Homayoun</i>	
LACore: A Supercomputing-Like Linear Algebra Accelerator for SoC-Based Designs	137
<i>Samuel Steffl and Sherief Reda</i>	
H3 (Heterogeneity in 3D): A Logic-on-Logic 3D-Stacked Heterogeneous Multi-Core Processor	145
<i>Vinesh Srinivasan, Rangeen Basu Roy Chowdhury, Elliott Forbes, Randy Widialaksono, Zhenqian Zhang, Joshua Schabel, Sungkwan Ku, Steve Lipa, Eric Rotenberg, W. Rhett Davis, and Paul D. Franzon</i>	
ABDTR: Approximation-Based Dynamic Traffic Regulation for Networks-on-Chip Systems	153
<i>Ling Wang, Xiaohang Wang, and Yadong Wang</i>	
CAGE: A Contention-Aware Game-Theoretic Model for Heterogeneous Resource Assignment	161
<i>Diman Zad Tootaghaj and Farshid Farhat</i>	
Energy-Aware Task Scheduling on Heterogeneous NoC-Based MPSoCs	165
<i>Suhaimi Abd Ishak, Hui Wu, and Umair Ullah Tariq</i>	

Session 3A: Debugging and Validation

Effective Signal Restoration in Post-Silicon Validation	169
<i>Xiaobang Liu and Ranga Vemuri</i>	
A Post-Silicon Trace Analysis Approach for System-on-Chip Protocol Debug	177
<i>Yuting Cao, Hao Zheng, Hernan Palombo, Sandip Ray, and Jin Yang</i>	
QUEBS: Qualifying Event Based Search in Concolic Testing for Validation of RTL Models	185
<i>Alif Ahmed and Prabhat Mishra</i>	
Automated Debugging of Arithmetic Circuits Using Incremental Gröbner Basis Reduction	193
<i>Farimah Farahmandi and Prabhat Mishra</i>	

Session 3B: Graph Processing and NoC Architectures

GraphTuner: An Input Dependence Aware Loop Perforation Scheme for Efficient Execution of Approximated Graph Algorithms	201
<i>Hamza Omar, Masab Ahmad, and Omer Khan</i>	
A High-Performance Deeply Pipelined Architecture for Elementary Transcendental Function Evaluation	209
<i>Jing Chen and Xue Liu</i>	
Conga: Towards Efficient Processing of Concurrent Graph Queries on Shared-Memory Machines	217
<i>Peitian Pan and Chao Li</i>	
Exploring Scalable Data Allocation and Parallel Computing on NoC-Based Embedded Many Cores	225
<i>Yuya Maruyama, Shinpei Kato, and Takuya Azumi</i>	
DAS: An Efficient NoC Router for Mixed-Criticality Real-Time Systems	229
<i>Mourad Dridi, Stéphane Rubini, Mounir Lallali, Martha Johanna Sepulveda Florez, Frank Singhoff, and Jean-Philippe Diguet</i>	

Session 4A: EDA with Focus on Multicore, FPGAs, and 3D

Monolithic 3D-Enabled High Performance and Energy Efficient Network-on-Chip	233
<i>Sourav Das, Janardhan Rao Doppa, Partha Pratim Pande, and Krishnendu Chakrabarty</i>	
Vulnerability-Aware Energy Optimization Using Reconfigurable Caches in Multicore Systems	241
<i>Yuanwen Huang and Prabhat Mishra</i>	
Dependency-Aware Parallel Routing for Large-Scale FPGAs	249
<i>Minghua Shen, Nong Xiao, and Guojie Luo</i>	
Design Exploration for Multiple Level Cell Based Non-Volatile FPGAs	257
<i>Ke Liu, Mengying Zhao, Lei Ju, Zhiping Jia, Chun Jason Xue, and Jingtong Hu</i>	
LDPC-Based Adaptive Multi-Error Correction for 3D Memories	265
<i>Mihai Lefter, George Voicu, Thomas Marconi, Valentin Savin, and Sorin Dan Cotofana</i>	
ILP-Based Identification of Redundant Logic Insertions for Opportunistic Yield Improvement during Early Process Learning	269
<i>Tuck-Boon Chan, Wei-Ting Jonas Chan, and Andrew B. Kahng</i>	

Session 4B: Hardware Acceleration for Neural Networks

A Dynamic Deep Neural Network Design for Efficient Workload Allocation in Edge Computing	273
<i>Chi Lo, Yu-Yi Su, Chun-Yi Lee, and Shih-Chieh Chang</i>	
DAWS: Exploiting Crossbar Characteristics for Improving Write Performance of High Density Resistive Memory	281
<i>Chengning Wang, Dan Feng, Jingning Liu, Wei Tong, Bing Wu, and Yang Zhang</i>	
Hardware Acceleration of Bayesian Neural Networks Using RAM Based Linear Feedback Gaussian Random Number Generators	289
<i>Ruizhe Cai, Ao Ren, Luhao Wangy, Massoud Pedramy, and Yanzhi Wang</i>	
Low Power Spatial Localization of Mobile Sensors with Recurrent Neural Network	297
<i>Nick Iliev and Amit Ranjan Trivedi</i>	
An FPGA-Based Coprocessor for Hash Unit Acceleration	301
<i>Abbas Fairouz and Sunil P. Khatri</i>	

Session 5A: Hardware Security II

Power Profile Equalizer: A Lightweight Countermeasure against Side-Channel Attack	305
<i>Chenguang Wang, Ming Yan, Yici Cai, Qiang Zhou, and Jianlei Yang</i>	
FSM Anomaly Detection Using Formal Analysis	313
<i>Farimah Farahmandi and Prabhat Mishra</i>	
Automatic Security Property Generation for Detecting Information-Leaking Hardware Trojans	321
<i>Chenguang Wang, Yici Cai, and Qiang Zhou</i>	
Implications of Distributed On-Chip Power Delivery on EM Side-Channel Attacks	329
<i>Ahmed Waheed Khan, Tanya Wanchoo, Gokhan Mumcu, and Selçuk Köse</i>	
Fingerprinting Field Programmable Gate Arrays	337
<i>Vinayaka Jyothi, Ashik Poojari, Richard Stern, and Ramesh Karri</i>	
Logic Obfuscation against IC Reverse Engineering Attacks Using PLGs	341
<i>Qutaiba Alasad and Jiann Yuan</i>	

Session 5B: Memory and Cache Optimizations

DAAIP: Deadblock Aware Adaptive Insertion Policy for High Performance Caching	345
<i>Newton, Sujit Kr Mahto, Suhit Pai, and Virendra Singh</i>	
Dual Dictionary Compression for the Last Level Cache	353
<i>Akshay Lahiry and David Kaeli</i>	
Jenga: Efficient Fault Tolerance for Stacked DRAM	361
<i>Georgios Mappouras, Alireza Vahid, Robert Calderbank, Derek R. Hower, and Daniel J. Sorin</i>	
SelSMaP: A Selective Stride Masking Prefetching Scheme	369
<i>Jiajun Wang, Reena Panda, and Lizy Kurian John</i>	
T2: A Highly Accurate and Energy Efficient Stride Prefetcher	373
<i>Sushant Kondguli and Michael Huang</i>	

Session 6A: Verification and Fault Tolerance

Automatic Protocol Compliance Checking of SystemC TLM-2.0 Simulation Behavior Using Timed Automata	377
<i>Mehran Goli, Jannis Stoppe, and Rolf Drechsler</i>	
Crosstalk Free Coding Systems to Protect NoC Channels against Crosstalk Faults	385
<i>Kimia Soleimani, Ahmad Patooghy, Nasim Soltani, Lake Bu, and Michel A. Kinsky</i>	
Limited Magnitude Error Correction Using OLS Codes for Memories with Multilevel Cells	391
<i>Abhishek Das and Nur A. Touba</i>	
Yoda: Judge Me by My Size, Do You?	395
<i>Jiangwei Zhang, Donald Kline Jr., Liang Fang, Rami Melhem, and Alex K. Jones</i>	
Fast Search-Based RTL Test Generation Using Control-Flow Path Guidance	399
<i>Sonal Pinto and Michael S. Hsiao</i>	

Session 6B: Lithography and Patterning

DTCO for DSA-MP Hybrid Lithography with Double-BCP Materials in Sub-7nm Node	403
<i>Jiaojiao Ou, Xiaqing Xu, Brian Cline, Greg Yeric, and David Z. Pan</i>	
Post-Layout Perturbation towards Stitch Friendly Layout for Multiple E-Beam Lithography	411
<i>Sudipta Paul, Pritha Banerjee, and Susmita Sur-Kolay</i>	

Patterning Aware Design Optimization of Selective Etching in N5 and Beyond	415
<i>Yibo Lin, Peter Debacker, Darko Trivkovic, Ryoung-Han Kim, Praveen Raghavan, and David Z. Pan</i>	

Special Session 1: On How to Design and Manage Complex Heterogeneous Distributed Computing Systems

Scaling Neural Network Performance through Customized Hardware	
Architectures on Reconfigurable Logic	419
<i>Michaela Blott, Thomas B. Preußen, Nicholas Fraser, Giulio Gambardella, Kenneth O'Brien, Yaman Umuroglu, and Miriam Leeser</i>	
The Role of CAD Frameworks in Heterogeneous FPGA-Based Cloud Systems	423
<i>Lorenzo Di Tucci, Marco Rabozzi, Luca Stornaiuolo, and Marco D. Santambrogio</i>	
A Common Backend for Hardware Acceleration on FPGA	427
<i>Emanuele Del Sozzo, Riyadh Baghdadi, Saman Amarasinghe, and Marco D. Santambrogio</i>	
Towards Accelerating Generic Machine Learning Prediction Pipelines	431
<i>Alberto Scolari, Yunseong Lee, Markus Weimer, and Matteo Interlandi</i>	
Convolutional Neural Networks on Dataflow Engines	435
<i>Nils Voss, Marco Bacis, Oskar Mencer, Georgi Gaydadjiev, and Wayne Luk</i>	

Session 7A: LCD with Focus on Emerging Technology

Exploring STT-MRAM Based In-Memory Computing Paradigm with Application of Image Edge Extraction	439
<i>Zhezhi He, Shaahin Angizi, and Deliang Fan</i>	
Floating Point Square Root under HUB Format	447
<i>Julio Villalba-Moreno and Javier Hormigo</i>	
Read Error Resilient MLC STT-MRAM Based Last Level Cache	455
<i>Wen Wen, Youtao Zhang, and Jun Yang</i>	
Ultra-Fast SOT-MRAM Cell with STT Current for Deterministic Switching	463
<i>Behzad Zeinali, Jens K. Madsen, Praveen Raghavan, and Farshad Moradi</i>	
Pulse Ring Oscillator Tuning via Pulse Dynamics	469
<i>Aditya Dalakoti, Merritt Miller, and Forrest Brewer</i>	
DLL-Assisted Clock Synchronization Method for Multi-Die ICs	473
<i>Chia-Yuan Cheng, Shi-Yu Huang, Ding-Ming Kwai, and Yung-Fa Chou</i>	

Session 7B: Power-Performance Optimization of Multicore Architecture

M2S-CGM: A Detailed Architectural Simulator for Coherent CPU-GPU Systems	477
<i>Christopher E. Giles and Mark A. Heinrich</i>	
Sharing-Aware Efficient Private Caching in Many-Core Server Processors	485
<i>Sudhanshu Shukla and Mainak Chaudhuri</i>	
Decomposed Task Mapping to Maximize QoS in Energy-Constrained Real-Time Multicores	493
<i>Lei Mo, Angeliki Kritikakou, and Olivier Sentieys</i>	
Using Application-Level Thread Progress Information to Manage Power and Performance	501
<i>Sabrina M. Neuman, Jason E. Miller, Daniel Sanchez, and Srinivas Devadas</i>	
Configurable SoC In-Situ Hardware/Software Co-Design Design Space Exploration	509
<i>Siyuan Xu, Benjamin Carrion Schafer, and Yidi Liu</i>	
Quality of Service-Aware Dynamic Voltage and Frequency Scaling for Mobile 3D Graphics Applications	513
<i>Navid Farazmand and David R. Kaeli</i>	

Session 8A: Synthesis and Security

Approximate Disjoint Bi-Decomposition and Its Application to Approximate Logic Synthesis	517
<i>Yue Yao, Shuyang Huang, Chen Wang, Yi Wu, and Weikang Qian</i>	
Timing-Abstract Circuit Design in Transaction-Level Verilog	525
<i>Steven F. Hoover</i>	
ReHLS: Resource-Aware Program Transformation Workflow for High-Level Synthesis	533
<i>Atieh Lotfi and Rajesh K. Gupta</i>	
Identifying Reversible Circuit Synthesis Approaches to Enable IP Piracy Attacks	537
<i>Samah Saeed, Nithin Mahendran, Alwin Zulehner, Robert Wille, and Ramesh Karri</i>	
Compiler-Assisted Threshold Implementation against Power Analysis Attacks	541
<i>Pei Luo, Konstantinos Athanasiou, Liwei Zhang, Zhen Hang Jiang, Yunsi Fei, A. Adam Ding, and Thomas Wahl</i>	
TAINt: Tool for Automated INsertion of Trojans	545
<i>Vinayaka Jyothi, Prashanth Krishnamurthy, Farshad Khorrami, and Ramesh Karri</i>	

Session 8B: Cloud and Storage Solutions

A Scale-Out Enterprise Storage Architecture	549
<i>Wonil Choi, Myoungsoo Jung, Mahmut Kandemir, and Chita Das</i>	
CloudShelter: Protecting Virtual Machines' Memory Resource Availability in Clouds	557
<i>Tianwei Zhang, Yuan Xu, Yungang Bao, and Ruby B. Lee</i>	
Using Disturbance Compensation and Data Clustering (DC)2 to Improve Reliability and Performance of 3D MLC Flash Memory	565
<i>Yazhi Feng, Dan Feng, Wei Tong, Yu Jiang, and Chuanqi Liu</i>	
Improving Performance of TLC RRAM with Compression-Ratio-Aware Data Encoding	573
<i>Jie Xu, Dan Feng, Yu Hua, Wei Tong, Jingning Liu, Chunyan Li, and Wen Zhou</i>	
Encoding Separately: An Energy-Efficient Write Scheme for MLC STT-RAM	581
<i>Jie Xu, Dan Feng, Wei Tong, Jingning Liu, and Wen Zhou</i>	
Quick-and-Dirty: Improving Performance of MLC PCM by Using Temporary Short Writes	585
<i>Mingzhe Zhang, Lunkai Zhang, Lei Jiang, Frederic T. Chong, and Zhiyong Liu</i>	

Special Session 2: Effective Voltage Scaling in Late CMOS Era

Applications of Deep Neural Networks for Ultra Low Power IoT	589
<i>Sreela Kodali, Patrick Hansen, Niamh Mulholland, Paul Whatmough, David Brooks, and Gu-Yeon Wei</i>	
Cross-Layer Resilience in Low-Voltage Digital Systems: Key Insights	593
<i>Eric Cheng, Jacob Abraham, Pradip Bose, Alper Buyuktosunoglu, Keith Campbell, Deming Chen, Cheng-Yong Cher, Hyungmin Cho, Binh Le, Klas Lilja, Shahrzad Mirkhani, Kevin Skadron, Mircea Stan, Lukasz Szafaryn, Christos Vezyrtzis, and Subhasish Mitra</i>	
Pre-RTL Voltage and Power Optimization for Low-Cost, Thermally Challenged Multicore Chips	597
<i>Alec Roelke, Runjie Zhang, Kaushik Mazumdar, Ke Wang, Kevin Skadron, and Mircea R. Stan</i>	
Very Low Voltage (VLV) Design	601
<i>Ramon Bertran, Pradip Bose, David Brooks, Jeff Burns, Alper Buyuktosunoglu, Nandhini Chandramoorthy, Eric Cheng, Martin Cochet, Schuyler Eldridge, Daniel Friedman, Hans Jacobson, Rajiv Joshi, Subhasish Mitra, Robert Montoye, Arun Paidimarri, Pritish Parida, Kevin Skadron, Mircea Stan, Karthik Swaminathan, Augusto Vega, Swagath Venkataramani, Christos Vezyrtzis, Gu-Yeon Wei, John-David Wellman, and Matthew Ziegler</i>	

Special Session 3: Spin-Computing: Lower the Barrier between Memory and Logic

Stochastic Switching of SHE-MTJ as a Natural Annealer for Efficient Combinatorial Optimization	605
--	-----

Yong Shim, Akhilesh Jaiswal, and Kaushik Roy

Energy Efficient In-Memory Binary Deep Neural Network Accelerator with Dual-Mode SOT-MRAM	609
---	-----

Deliang Fan and Shaahin Angizi

Programmable Stateful In-Memory Computing Paradigm via a Single Resistive Device	613
--	-----

Wang Kang, He Zhang, Peng Ouyang, Youguang Zhang, and Weisheng Zhao

Session 9A: Architecture and Microarchitecture Optimizations

NUPLet: A Photonic Based Multi-Chip NUCA Architecture	617
---	-----

Janibul Bashir and Smruti R. Sarangi

RAPS: Restore-Aware Policy Selection for STT-MRAM-Based Main Memory under Read Disturbance	625
--	-----

Armin Haj Aboutalebi and Lide Duan

BACM: Barrier-Aware Cache Management for Irregular Memory-Intensive GPGPU Workloads	633
---	-----

Yuxi Liu, Xia Zhao, Zhibin Yu, Zhenlin Wang, Xiaolin Wang, Yingwei Luo, and Lieven Eeckhout

Efficient Tagged Memory	641
-------------------------------	-----

Alexandre Joannou, Jonathan Woodruff, Robert Kovacsics, Simon W. Moore, Alex Bradbury, Hongyan Xia, Robert N.M. Watson, David Chisnall, Michael Roe, Brooks Davis, Edward Napierala, John Baldwin, Khilan Gudka, Peter G. Neumann, Alfredo Mazzinghi, Alex Richardson, Stacey Son, and A. Theodore Markettos

Improving Execution Time of Parallel Programs on Large Scale Chip Multiprocessors with Constant Average Power Processing	649
--	-----

Kramer Straube, Christopher Nitta, Raj Amirtharajah, Matthew Farrens, and Venkatesh Akella

Effective Optimization of Branch Predictors through Lightweight Simulation	653
--	-----

Chaobing Zhou, Libo Huang, Tan Zhang, Yongwen Wang, Chengyi Zhang, and Qiang Dou

Session 9B: Novel Architecture with 3D and Flash Memory

CooECC: A Cooperative Error Correction Scheme to Reduce LDPC Decoding Latency in NAND Flash	657
<i>Meng Zhang, Fei Wu, Yajuan Du, Chengmo Yang, Changsheng Xie, and Jiguang Wan</i>	
Fast, Ring-Based Design of 3D Stacked DRAM	665
<i>Andrew J. Douglass and Sunil P. Khatri</i>	
Memory-Bounded Randomness for Hardware-Constrained Encrypted Computation	673
<i>Nektarios Georgios Tsoutsos, Oleg Mazonka, and Michail Maniatakos</i>	
Exploiting Process Variation for Read Performance Improvement on LDPC Based Flash Memory Storage Systems	681
<i>Qiao Li, Liang Shi, Yejia Di, Yajuan Du, Chun J. Xue, and Edwin H.M. Sha</i>	
A Design-for-Test Solution for Monolithic 3D Integrated Circuits	685
<i>Abhishek Koneru, Sukeshwar Kannan, and Krishnendu Chakrabarty</i>	
Author Index	689