



Published on DATE 2019 (<https://past.date-conference.com>)

[Home](#) > [Printer-friendly PDF](#) > [Printer-friendly PDF](#)

IP1 Interactive Presentations

Date: Tuesday, March 26, 2019

Time: 16:00 - 16:30

Location / Room: Poster Area

Interactive Presentations run simultaneously during a 30-minute slot. Additionally, each IP paper is briefly introduced in a one-minute presentation in a corresponding regular session

Label	Presentation Title Authors
IP1-1	FAULT INJECTION ON HIDDEN REGISTERS IN A RISC-V ROCKET PROCESSOR AND SOFTWARE COUNTERMEASURES Speaker: Johan Laurent, Univ. Grenoble Alpes, Grenoble INP, LCIS, FR Authors: Johan Laurent ¹ , Vincent Berouille ¹ , Christophe Deleuze ¹ and Florian Pebay-Peyroula ² ¹ LCIS - Grenoble Institute of Technology - Univ. Grenoble Alpes, FR; ² CEA-Leti, FR Abstract <i>To protect against hardware fault attacks, developers can use software countermeasures. They are generally designed to thwart software fault models such as instruction skip or memory corruption. However, these typical models do not take into account the actual implementation of a processor. By analyzing the processor microarchitecture, it is possible to bypass typical software countermeasures. In this paper, we analyze the vulnerability of a secure code from FISSC (Fault Injection and Simulation Secure Collection), by simulating fault injections in a RISC-V Rocket processor RTL description. We highlight the importance of hidden registers in the processor pipeline, which temporarily hold data during code execution. Secret data can be leaked by attacking these hidden registers. Software countermeasures against such attacks are also proposed.</i> Download Paper (PDF; Only available from the DATE venue WiFi)
IP1-2	METHODOLOGY FOR EM FAULT INJECTION: CHARGE-BASED FAULT MODEL Speaker: Haohao Liao, University of Waterloo, CA Authors: Haohao Liao and Catherine Gebotys, University of Waterloo, CA Abstract <i>Recently electromagnetic fault injection (EMFI) techniques have been found to have significant implications on the security of embedded devices. Unfortunately there is still a lack of understanding of EM fault models and countermeasures for embedded processors. For the first time, this paper proposes an extended fault model based on the concept of critical charge and a new EMFI backside methodology based on over-clocking. Results show that exact timing of EM pulses can provide reliable repeatable instruction replacement faults for specific programs. An attack on AES is demonstrated showing that the EM fault injection requires on average less than 222 EM pulses and 5.3 plaintexts to retrieve the full AES key. This research is critical for ensuring embedded processors and their instruction set architectures are secure and resistant to fault injection attacks.</i> Download Paper (PDF; Only available from the DATE venue WiFi)
IP1-3	SECURING CRYPTOGRAPHIC CIRCUITS BY EXPLOITING IMPLEMENTATION DIVERSITY AND PARTIAL RECONFIGURATION ON FPGAS Speaker: Benjamin Hettwer, Robert Bosch GmbH, DE Authors: Benjamin Hettwer ¹ , Johannes Petersen ² , Stefan Gehrer ¹ , Heike Neumann ² and Tim Güneysu ³ ¹ Robert Bosch GmbH, Corporate Sector Research, DE; ² Hamburg University of Applied Sciences, DE; ³ Horst Görtz Institute for IT Security, Ruhr-University Bochum, DE Abstract <i>Adaptive and reconfigurable systems such as Field Programmable Gate Arrays (FPGAs) play an integral part of many complex embedded platforms. This implies the capability to perform runtime changes to hardware circuits on demand. In this work, we make use of this feature to propose a novel countermeasure against physical attacks of cryptographic implementations. In particular, we leverage exploration of the implementation space on FPGAs to create various circuits with different hardware layouts from a single design of the Advanced Encryption Standard (AES), that are dynamically exchanged during device operation. We provide evidence from practical experiments based on a modern Xilinx ZYNQ UltraScale+ FPGA that our approach increases the resistance against physical attacks by at least factor two. Furthermore, the genericness of our approach allows an easy adaption to other algorithms and combination with other countermeasures</i> Download Paper (PDF; Only available from the DATE venue WiFi)
IP1-4	STT-ANGIE: ASYNCHRONOUS TRUE RANDOM NUMBER GENERATOR USING STT-MTJ Speaker: Ben Perach, Faculty of Electrical Engineering, Technion - Israel Institute of Technology, IL Authors: Ben Perach and Shahar Kvatinsky, Technion, IL Abstract <i>The Spin Transfer Torque Magnetic Tunnel Junction (STT-MTJ) is an emerging memory technology whose interesting stochastic behavior might benefit security applications. In this paper, we leverage this stochastic behavior to construct a true random number generator (TRNG), the basic module in the process of encryption key generation. Our proposed TRNG operates asynchronously and thus can use small and fast STT MTJ devices. As such, it can be embedded in low-power and low-frequency devices without loss of entropy. We evaluate the proposed TRNG using a numerical simulation, solving the Landau-Lifshitz-Gilbert (LLG) equation system of the STTMTJ devices. Design considerations, attack analysis, and process variation are discussed and evaluated. The evaluation shows that our solution is robust to process variation, achieving a Shannon-entropy generating rate between 99.7Mbps and 127.8Mbps for 90% of the instances.</i> Download Paper (PDF; Only available from the DATE venue WiFi)

- IP1-5 **ADAPTIVE TRANSIENT LEAKAGE-AWARE LINEARISED MODEL FOR THERMAL ANALYSIS OF 3-D ICS**
Speaker:
Milan Mihajlovic, University of Manchester, GB
Authors:
Chao Zhang, Milan Mihajlovic and Vasilis Pavlidis, The University of Manchester, GB
Abstract
Physics-based models for thermal simulation that involve numerical solution of the heat equation are well placed to accurately capture the heterogeneity of materials and structures in modern 3-D integrated circuits (ICs). The introduction of non-linear effects in thermal coefficients and leakage power improves significantly the accuracy of thermal models. However, this non-linearity increases significantly the complexity and computational time of the analysis. In this paper, we introduce a linearised thermal model by demonstrating that weak temperature dependence of the specific heat and the thermal conductivity of silicon-based materials has only minor effect to computed temperature profiles. Thus, these parameters can be considered constant in working temperature ranges of modern ICs. The non-linearity in leakage power is approximated by a piecewise linear least square fit and the resulting model is linearised by exact Newton's method contrary to previous works that employ either simple iterative or inexact Newton's method. The method is implemented in the context of transient thermal analysis with adaptive time step selection, where we demonstrate that it is essential to apply Newton corrections to obtain the right time step size selection. The resulting method is up to 2x faster than a full non-linear method, typically introducing a global relative error of less than 1%.
[Download Paper \(PDF; Only available from the DATE venue WiFi\)](#)
- IP1-6 **FASTCOOL: LEAKAGE AWARE DYNAMIC THERMAL MANAGEMENT OF 3D MEMORIES**
Speaker:
Lokesh Siddhu, IIT Delhi, IN
Authors:
Lokesh Siddhu¹ and Preeti Ranjan Panda²
¹Indian Institute of Technology, Delhi, IN; ²IIT Delhi, IN
Abstract
3D memory systems offer several advantages in terms of area, bandwidth, and energy efficiency. However, thermal issues arising out of higher power densities have limited their widespread use. While prior works have looked at reducing dynamic power through reduced memory accesses, in these memories, both leakage and dynamic power consumption are comparable. Furthermore, as the temperature rises the leakage power increases, creating a thermal-leakage loop. We study the impact of leakage power on 3D memory temperature and propose turning OFF hot channels to meet thermal constraints. Data is migrated to a 2D memory before closing a 3D channel. We introduce an analytical model to assess the 2D memory delay and use the model to guide data migration decisions. Our experiments show that the proposed optimization improves performance by 27% on an average (up to 66%) over state-of-the-art strategies.
[Download Paper \(PDF; Only available from the DATE venue WiFi\)](#)
- IP1-7 **ON THE USE OF CAUSAL FEATURE SELECTION IN THE CONTEXT OF MACHINE-LEARNING INDIRECT TEST**
Speaker:
Manuel Barragán, TIMA laboratory, FR
Authors:
Manuel Barragan¹, Gildas Leger², Florent Cilic³, Estelle Lauga-Larroze⁴, Sylvain Bourdel⁴ and Salvador Mir³
¹TIMA Laboratory, FR; ²Instituto de Microelectronica de Sevilla, IMSE-CNM, (CSIC - Universidad de Sevilla), ES; ³TIMA, FR; ⁴RFICLab, FR
Abstract
The test of analog, mixed-signal and RF (AMS-RF) circuits is still considered as a matter of human creativity, and although many attempts have been made towards their automation, no accepted and complete solution is yet available. Indeed, capturing the design knowledge of an experienced analog designer is one of the key challenges faced by the Electronic Design Automation (EDA) community. In this paper we explore the use of causal inference tools in the context of AMS-RF design and test with the goal of defining a methodology for uncovering the root causes of performance variation in these systems. We believe that such an analysis can be a promising first step for future EDA algorithms for AMS-RF systems.
[Download Paper \(PDF; Only available from the DATE venue WiFi\)](#)
- IP1-8 **ACCURACY AND COMPACTNESS IN DECISION DIAGRAMS FOR QUANTUM COMPUTATION**
Speaker:
Alwin Zulehner, Johannes Kepler University Linz, AT
Authors:
Alwin Zulehner¹, Philipp Niemann², Rolf Drechsler³ and Robert Wille¹
¹Johannes Kepler University Linz, AT; ²Cyber-Physical Systems, DFKI GmbH, DE; ³University of Bremen, DE
Abstract
Quantum computation is a promising research field since it allows to conduct certain tasks exponentially faster than on conventional machines. As in the conventional domain, decision diagrams are heavily used in different design tasks for quantum computation like synthesis, verification, or simulation. However, unlike decision diagrams for the conventional domain, decision diagrams for quantum computation as of now suffer from a trade-off between accuracy and compactness that requires parameter fine-tuning on a case-by-case basis. In this work, we—for the first time—describe and evaluate the effects of this trade-off. Moreover, we propose an alternative approach that utilizes an algebraic representation of the occurring irrational numbers and outline how this can be incorporated in a decision diagram in order to overcome this trade-off.
[Download Paper \(PDF; Only available from the DATE venue WiFi\)](#)
- IP1-9 **ONE METHOD - ALL ERROR-METRICS: A THREE-STAGE APPROACH FOR ERROR-METRIC EVALUATION IN APPROXIMATE**
Speaker:
Saman Fröhlich, University of Bremen/DFKI GmbH, DE
Authors:
Saman Fröhlich¹, Daniel Grosse² and Rolf Drechsler²
¹University of Bremen/DFKI GmbH, DE; ²University of Bremen, DE
Abstract
Approximate Computing is a design paradigm that makes use of the error tolerance inherited by many applications, such as machine learning, media processing and data mining. The goal of Approximate Computing is to trade off accuracy for performance in terms of computation time, energy consumption and/or hardware complexity. In the field of circuit design for Approximate Computing, error-metrics are used to express the degree of approximation. Evaluating these error-metrics is a key challenge. Several approaches exist, however, to this day not all relevant metrics can be evaluated with formal methods. Recently, Symbolic Computer Algebra (SCA) has been used to evaluate error-metrics during approximate hardware generation. In this paper, we generalize the idea to use SCA and propose a methodology which is suitable for formal evaluation of all established error-metrics. This approach can be divided into three-stages: (i) Determine the remainder of the AC circuit wrt. the specification using SCA, (ii) build an Algebraic Decision Diagram (ADD) to represent the remainder and (iii) evaluate each error-metric by a tailored ADD traversal algorithm. Besides being the first to propose a closed formal method for evaluation of all relevant error-metrics, we are the first to ever propose formal algorithms for the evaluation of the worst-case-relative and the average-case-relative error-metrics. In the experiments, we apply our algorithms to a large and well-known benchmark set.
[Download Paper \(PDF; Only available from the DATE venue WiFi\)](#)

IP1-10 **REVERSIBLE PEBBLING GAME FOR QUANTUM MEMORY MANAGEMENT**

Speaker:

Giulia Meuli, EPFL, CH

Authors:

Giulia Meuli¹, Mathias Soeken¹, Martin Roetteler², Nikolaj Björner² and Giovanni De Micheli¹

¹EPFL, CH; ²Microsoft, US

Abstract

Quantum memory management is becoming a pressing problem, especially given the recent research effort to develop new and more complex quantum algorithms. The only existing automatic method for quantum states clean-up relies on the availability of many extra resources. In this work, we propose an automatic tool for quantum memory management. We show how this problem exactly matches the reversible pebbling game. Based on that, we develop a SAT-based algorithm that returns a valid clean-up strategy, taking the limitations of the quantum hardware into account. The developed tool empowers the designer with the flexibility required to explore the trade-off between memory resources and number of operations. We present two show-cases to prove the validity of our approach. First, we apply the algorithm to straight-line programs, widely used in cryptographic applications. Second, we perform a comparison with the existing approach, showing an average improvement of 52.77%.

[Download Paper \(PDF; Only available from the DATE venue WiFi\)](#)

IP1-11 **TYPECNN: CNN DEVELOPMENT FRAMEWORK WITH FLEXIBLE DATA TYPES**

Speaker:

Lukas Sekanina, Brno University of Technology, CZ

Authors:

Petr Rek and Lukas Sekanina, Brno University of Technology, CZ

Abstract

The rapid progress in artificial intelligence technologies based on deep and convolutional neural networks (CNN) has led to an enormous interest in efficient implementations of neural networks in embedded devices and hardware. We present a new software framework for the development of (approximate) convolutional neural networks in which the user can define and use various data types for forward (inference) procedure, backward (training) procedure and weights. Moreover, non-standard arithmetic operations such as approximate multipliers can easily be integrated into the CNN under design. This flexibility enables to analyze the impact of chosen data types and non-standard arithmetic operations on CNN training and inference efficiency. The framework was implemented in C++ and evaluated using several case studies.

[Download Paper \(PDF; Only available from the DATE venue WiFi\)](#)

IP1-12 **GUARANTEED COMPRESSION RATE FOR ACTIVATIONS IN CNNs USING A FREQUENCY PRUNING APPROACH**

Speaker:

Sebastian Vogel, Robert Bosch GmbH, DE

Authors:

Sebastian Vogel¹, Christoph Schorn¹, Andre Guntoro¹ and Gerd Ascheid²

¹Robert Bosch GmbH, DE; ²RWTH Aachen University, DE

Abstract

Convolutional Neural Networks have become state of the art for many computer vision tasks. However, the size of Neural Networks prevents their application in resource constrained systems. In this work, we present a lossy compression technique for intermediate results of Convolutional Neural Networks. The proposed method offers guaranteed compression rates and additionally adapts to performance requirements. Our experiments with networks for classification and semantic segmentation show, that our method outperforms state-of-the-art compression techniques used in CNN accelerators.

[Download Paper \(PDF; Only available from the DATE venue WiFi\)](#)

IP1-13 **RUNTIME MONITORING NEURON ACTIVATION PATTERNS**

Speaker:

Chih-Hong Cheng, fortiss, DE

Authors:

Chih-Hong Cheng¹, Georg Nührenberg¹ and Hirotoshi Yasuoka²

¹fortiss - Landesforschungsinstitut des Freistaats Bayern, DE; ²DENSO Corporation, JP

Abstract

For using neural networks in safety critical domains such as automated driving, it is important to know if a decision made by a neural network is supported by prior similarities in training. We propose runtime neuron activation pattern monitoring - after the standard training process, one creates a monitor by feeding the training data to the network again in order to store the neuron activation patterns in abstract form. In operation, a classification decision over an input is further supplemented by examining if a pattern similar (measured by Hamming distance) to the generated pattern is contained in the monitor. If the monitor does not contain any pattern similar to the generated pattern, it raises a warning that the decision is not based on the training data. Our experiments show that, by adjusting the similarity-threshold for activation patterns, the monitors can report a significant portion of misclassifications to be not supported by training with a small false-positive rate, when evaluated on a test set.

[Download Paper \(PDF; Only available from the DATE venue WiFi\)](#)

IP1-14 **CHIP HEALTH TRACKING USING DYNAMIC IN-SITU DELAY MONITORING**

Speaker:

Hadi Ahmadi Balef, Eindhoven University of Technology, NL

Authors:

Hadi Ahmadi Balef¹, Kees Goossens² and José Pineda de Gyvez¹

¹Eindhoven University of Technology, NL; ²Eindhoven university of technology, NL

Abstract

Tracking the gradual effect of silicon aging on circuit delays requires fine-grain slack monitoring. The conventional slack monitoring techniques intend to measure the worst-case static slack, i.e. the slack of longest timing path. In sharp contrast to the conventional techniques, we propose a novel technique that is based on dynamic excitation of in-situ delay monitors (i.e. the dynamic excitation of timing paths that are monitored). As delays degrade, path delays increase and the monitors are excited more frequently. With the proposed technique, a fine-grained signature of delay degradation is extracted from the excitation rate of monitors. The in-situ monitors are inserted at intermediate points along timing paths to increase the sensitivity of signature to delay degradation. A new efficient monitor insertion algorithm is also proposed that reduces the number of monitors by ~2.1X compared to other works for an ARM Cortex M0 processor.

[Download Paper \(PDF; Only available from the DATE venue WiFi\)](#)

IP1-15 **PCFI: PROGRAM COUNTER GUIDED FAULT INJECTION FOR ACCELERATING GPU RELIABILITY ASSESSMENT**

Speaker:

Paolo Rech, UFRGS, BR

Authors:

Fritz Previlon, Charu Kalra, Devesh Tiwari and David Kaeli, Northeastern University, US

Abstract

Reliability has become a first-class design objective for GPU devices due to increasing soft-error rate. To assess the reliability of GPU programs, researchers rely on software fault-injection methods. Unfortunately, software fault-injection process is prohibitively expensive, requiring multiple days to complete a statistically sound fault-injection campaign. Therefore, to address this challenge, this paper proposes a novel fault-injection method, PCFI, that reduces the number of fault injections by exploiting the predictability in fault-injection outcome based on the program counter of the soft-error affected instruction. Evaluation on a variety of GPU programs covering a wide range of application domains shows that PCFI reduces the time to complete fault-injection campaigns by 22% on average without sacrificing the accuracy.

[Download Paper \(PDF; Only available from the DATE venue WiFi\)](#)

IP1-16 CHARACTERIZING THE RELIABILITY AND THRESHOLD VOLTAGE SHIFTING OF 3D CHARGE TRAP NAND FLASH

Speaker:

Weihua Liu, Huazhong University of Science and Technology, CN

Authors:

Weihua Liu¹, Fei Wu¹, Meng Zhang¹, Yifei Wang¹, Zhonghai Lu², Xiangfeng Lu³ and Changsheng Xie¹

¹Huazhong University of Science and Technology, CN; ²KTH Royal Institute of Technology, SE; ³Beijing Memblaze Technology Co., Ltd., CN

Abstract

3D charge trap (CT) triple-level cell (TLC) NAND flash gradually becomes a mainstream storage component due to high storage capacity and performance, but introducing a concern about reliability. Fault tolerance and data management schemes are capable of improving reliability. Designing a more efficient solution, however, needs to understand the reliability characteristics of 3D CT TLC NAND flash. To facilitate such understanding, by exploiting a real-world testing platform, we investigate the reliability characteristics including the raw bit error rate (RBER) and the threshold voltage (V_{th}) shifting features after suffering from variable disturbances. We give an analysis of why these characteristics exist in 3D CT TLC NAND flash. We hope these observations can guide the designers to propose high efficient solutions to the reliability problem.

[Download Paper \(PDF; Only available from the DATE venue WiFi\)](#)

IP1-17 HIDDEN DELAY FAULT SENSOR FOR TEST, RELIABILITY AND SECURITY

Speaker:

Giorgio Di Natale, CNRS - TIMA, FR

Authors:

Giorgio Di Natale¹, Elena Ioana Vatajelu², Kalpana SENTHAMARAI KANNAN² and Lorena Anghel³

¹LIRMM, FR; ²TIMA, FR; ³Grenoble-Alpes University, FR

Abstract

In this paper we present a novel hidden-delay-fault sensor design and a preliminary analysis of its circuit integration and applicability. In our proposed method, the delay sensing is achieved by sampling data on both rising and falling clock edges and using a variable duty cycle to control the range of the sensed delay fault. The main advantage of our proposed method is that it works at nominal frequency, can cover a wide range of delay faults and it is versatile in its applicability. It can be used (i) during testing to perform user-defined hidden-delay-fault test, (ii) for reliability degradation estimation due to process, environmental variations and ageing, and (iii) in security to detect the insertion of Trojan horses that alter the path delay.

[Download Paper \(PDF; Only available from the DATE venue WiFi\)](#)

IP1-18 EFFECT OF DEVICE VARIATION ON MAPPING BINARY NEURAL NETWORK TO MEMRISTOR CROSSBAR ARRAY

Speaker:

Wooseok Yi, POSTECH, KR

Authors:

Wooseok Yi, Yulhwa Kim and Jae-Joon Kim, Pohang University of Science and Technology, KR

Abstract

In memristor crossbar array (MCA)-based neural network hardware, it is generally assumed that entire wordlines (WLs) are simultaneously enabled for parallel matrix-vector multiplication (MxV) operation. However, the error probability of MxV in a memristor crossbar array (MCA) increases as the resistance ratio (R -ratio) of a memristor decreases and the resistance variation and the number of simultaneously activated WLs increase. In this paper, we analyze the effect of R -ratio and variation of memristor devices on read sense margin and inference accuracy of MCA-based Binary Neural Network (BNN) hardware. We first show that only a limited number of WLs should be enabled to ensure correct MxV output when the R -ratio is small. On the other hand, we also show that, if the resistance variation becomes higher than a certain level, simultaneous activation of large number of WLs produces the higher accuracy even when R -ratio is small. Based on the analysis, we propose the Accuracy Estimation (AE) factor to find the optimal number of word lines that are simultaneously activated.

[Download Paper \(PDF; Only available from the DATE venue WiFi\)](#)

IP1-19 ACCURATE WIRELENGTH PREDICTION FOR PLACEMENT-AWARE SYNTHESIS THROUGH MACHINE LEARNING

Speaker:

Daijoon Hyun, KAIST, KR

Authors:

Daijoon Hyun, Yuepeng Fan and Youngsoo Shin, KAIST, KR

Abstract

Placement-aware synthesis, which combines logic synthesis with virtual placement and routing (P&R) to better take account of wiring, has been popular for timing closure. The wirelength after virtual placement is correlated to actual wirelength, but correlation is not strong enough for some chosen paths. An algorithm to predict the actual wirelength from placement-aware synthesis is presented. It extracts a number of parameters from a given virtual path. A handful of synthetic parameters are compiled through linear discriminant analysis (LDA), and they are submitted to a few machine learning models. The final prediction of actual wirelength is given by the weighted sum of prediction from such machine learning models, in which weight is determined by the population of neighbors in parameter space. Experiments indicate that the predicted wirelength is 93% accurate compared to actual wirelength; this can be compared to conventional virtual placement, in which wirelength is predicted with only 79% accuracy.

[Download Paper \(PDF; Only available from the DATE venue WiFi\)](#)

IP1-20 A MIXED-HEIGHT STANDARD CELL PLACEMENT FLOW FOR DIGITAL CIRCUIT BLOCKS

Speaker:

Yi-Cheng Zhao, National Tsing Hua University, TW

Authors:

Yi-Cheng Zhao¹, Yu-Chieh Lin¹, Ting-Chi Wang¹, Ting-Hsiung Wang², Yun-Ru Wu², Hsin-Chang Lin² and Shu-Yi Kao²

¹National Tsing Hua University, TW; ²Realtek Semiconductor Corp., TW

Abstract

In this paper, we present a mixed-height standard cell placement flow for digital circuit blocks. To our best knowledge, commercial tools currently do not support this type of flow in a fully automated manner. In our placement flow, we leverage a commercial placement tool and integrate it with several new point tools. Promising experimental results are reported to demonstrate the efficacy of our placement flow.

[Download Paper \(PDF; Only available from the DATE venue WiFi\)](#)

IP1-21 AGGRESSIVE MEMORY SPECULATION IN HW/SW CO-DESIGNED MACHINES

Speaker:

Simon Rokicki, INRIA, FR

Authors:

Simon Rokicki, Erven Rohou and Steven Derrien, IRISA, Rennes, FR

Abstract

Single-ISA heterogeneous systems (such as ARM big.LITTLE) are an attractive solution for embedded platforms as they expose many performance and energy consumption trade-offs directly to the operating system. Recent works have demonstrated the ability to increase their efficiency by using VLIW cores, supported through Dynamic Binary Translation (DBT). Such an approach exposes even more heterogeneity while maintaining the illusion of a single-ISA system. However, VLIW cores cannot rival with Out-of-Order (OoO) cores when it comes to performance. One of the reason is that OoO cores heavily rely on speculative execution. In this work, we study how it is possible to take advantage of memory dependency speculation during the DBT process. More specifically, our approach builds on a hardware accelerated DBT framework, which enables fine-grained dynamic iterative optimizations. This is achieved through a combination of hardware and software, following the principles of co-designed machines. The experimental study conducted demonstrates that our approach leads to a geo-mean speed-up of 20% while keeping the hardware overhead low.

[Download Paper \(PDF; Only available from the DATE venue WiFi\)](#)